A Pilot of a QoS-Aware Wireless Back-Haul Network for Rural Areas

Philipp Batroff[1], George Ghinea[2], Thorsten Horstmann[1], Karl Jonas[1], and Jens Moedeker[1]

[1] Fraunhofer FOKUS, Sankt Augustin, Germany
{philipp.batroff,thorsten.horstmann, karl.jonas,jens.moedeker}@fokus.fraunhofer.de
[2] Brunel University, London, United Kingdom
george.ghinea@brunel.ac.uk

Abstract. Rural areas in emerging regions often lack affordable broadband Internet connectivity, which limits the access to, for example, knowledge, government services or education. The major limiting factors are the Capital Expenditure (CAPEX) and the Operational Expenditure (OPEX) related to traditional wireless carrier equipment, its relatively large energy footprint and the vast but sparsely populated areas to be covered. Since in many rural regions access to a power grid is not available or highly instable, ensuring a 24/7 operation of cell site is a very costly task. To address those issues, we have developed a carrier-grade heterogeneous back-haul architecture in order to complement, extend or even replace traditional operator equipment. Our Wireless Back-Haul (WiBACK) network technology provides wireless back-haul coverage while building on cost-effective and low-power equipment. In this paper we present a pilot scenario in Maseru, Lesotho, where an entrepreneur starts out with three eKiosk/VoIP sites with the goal to cover large parts on the city of Maseru. Using a testbed resembling the initial deployment scenario and identical hardware as planned for Maseru, we validate the self-configuration mechanisms, evaluate their performance in cases of node failures and show that the remaining network can quickly be reorganized.

Keywords: Heterogeneous Wireless Mesh, QoS, MPLS, IEEE 802.21.

1 Introduction

In the last years Wireless Mesh Networks (WMNs) have been a hot topic of interest for commercial operators as well as researchers in the academic world. Their potential to reduce OPEX tremendously by providing a resilient and fault-tolerant network due to Self-configuration and Self-management features, while requiring less deployment cost compared to traditional operator networks is one of the major advantages. This particularly applies if cost-effective network technology is used, e.g. IEEE 802.11[1]. However, in order to be considered as a carrier-grade network which eventually is deployed by an operator to connect customers, WiBACK networks must accomplish the same requirements in terms

R. Popescu-Zeletin et al. (Eds.) AFRICOMM 2011, LNICST 92, pp. 96–105, 2012.

of Quality of Service (QoS), availability and reliability as regular operator networks. These include the support of triple-play services as nowadays regularly used and expected by todays customers. Given that Voice-over-IP (VoIP) might generate comparatively requirements and load on non-collision free wireless links such as IEEE 802.11 EDCA, the allocation of available link and spectrum resources must be strictly managed by the network. Thus, our WiBACK architecture provides a Topology Management Function (TMF) as well as a Capacity Management Function (CMF). Whereas the first optimizes the scare radio spectrum resources by controlling which frequency is used on which link, the latter is in control of allocating the available network capacity to best accommodate user QoS-traffic demands. Additionally, in order to allow for service continuity the CMF quickly reacts on link failures or fluctuations reported by the monitoring component.

The WiBACK architecture itself can be used in manifold use-cases and allows for providing customized solutions which might range from a single-hop long distance wireless link to a region-wide multi-hop back-haul network connecting large urban and rural areas equally. Particularly to address the constraints of rural Africa the WiBACK architecture is designed to maintain a low energy footprint so that WiBACK nodes can be powered with alternative power sources such as solar and wind and hence be easily deployed even in areas without a stable power grid.

Our heterogeneous multi-radio WiBACK is inspired by the work of the EU FP7 CARrier grade wireless MEsh Network (CARMEN) [2] project and adopts its centrally managed cross-layer concept as well as the general concepts of IEEE 802.21 such as command and event services as well as the media abstraction paradigm. Moreover, in order to allow for effective QoS differentiation WiBACK relies on Multi Protocol Label Switching (MPLS)-based Traffic Engineering (TE) and a model to describe wireless channels. For user data transportation between two arbitrary nodes in the network MPLS Label-Switched Paths (LSPs) with dedicated per-hop resource allocation called *Pipes* are established.

The rest of this paper is structured as follows: First we present a prototype deployment scenario for rural areas in Maseru, Lesotho followed by a detailed description of the WiBACK solution. We then show validation and evaluation results regarding the TMF. Concluding, we will summarize our work and an outlook on future work is given.

2 Prototype Deployment Scenario

Larger and rather densely populated areas in developed or emerging countries, can be considered to offer Internet or telephone access to the majority of the population, albeit at highly varying levels regarding the available bandwidth and the cost structure.

In the vast and sparsely populated regions on the other hand, connectivity is often poor or even not available at all. The main reasons for this are the rather high CAPEX and OPEX of traditional operator equipment, unreliable or non-existent connections to a power grid and the lack of trained personnel. Combined with a rather low amount of revenue per customer, larger scale deployments are often not economically feasible.

Hence, to increase the end-user connectivity in rural areas, the CAPEX and mainly the OPEX must be lowered. One option to lower the CAPEX of a deployment scenario is to consider alternative wireless technologies to build the back-haul network. As aforementioned, possible options might be IEEE 802.11 or 802.16 based hardware, which can be adapted to provide high throughput over long distances. It is crucial though, to ensure that requirements such as strict QoS-enforcement and predictable behavior under high load situations can be met.

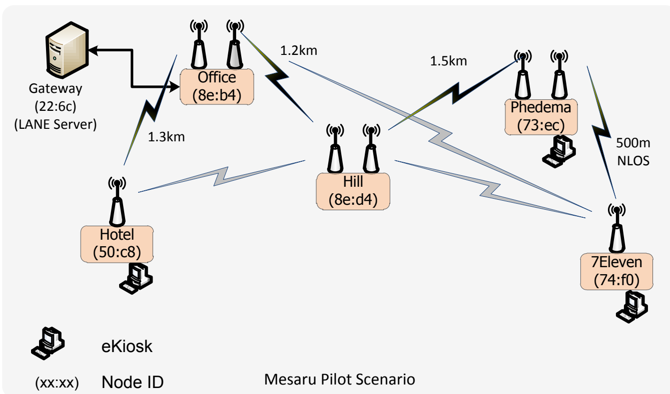

Fig. 1. The initial pilot in Maseru, Lesotho consists of five outdoor WiBACK nodes and one indoor node acting as a WiBACK and eKiosk management node

In this paper we outline a planned prototype deployment scenario in Maseru, Lesotho. A local entrepreneur is planning a small eKiosk businesses with the goal to provide reliable voice and data services. The initial setup, as depicted in figure 1 consists of five wireless nodes and one gateway which will have a connection to a local ISP. The first eKiosk systems are to be deployed at the *7eleven* and *Phedema* as well as at the *Hotel* site. In contrast the *Office* and *Hill* nodes are simple repeater nodes only forwarding traffic. Moreover, the *Gateway* and *Office* nodes will be connected via an Ethernet cable whereas the connections between all other nodes will be realized through Line of Sight (LOS) Wifi links with a range of approximately 1.2 km to 1.5 km. Parallel to the implementation arrangements for Lesotho we setup a smaller testbed with the same topology but smaller distances at our premises in Sankt Augustin for initial testing. All nodes are equipped equally in terms of hardware and number of Wifi devices in order to allow for an easy replacement of nodes in case of a hardware failure and to avoid extensive pre-configuration. A node determines the proper configuration by itself during the initial bootstrap phase. As antennas directional flat panels are used to support the black colored links in figure 1. However, communication via the grey colored links is also possible although with sub-optimal link budget. It should be noted, that the network does not aim at remaining constant in terms of size and nodes but rather to constantly grow, and by establishing more eKiosk systems to cover larger parts of Maseru.

3 WiBACK Approach

The architecture of WiBACK can be divided into the data and the control plane. The control plane is used to setup, manage and maintain the network nodes. Therefore the key concepts of the IEEE 802.21 standard [3] are adapted to support all challenges of wireless network management. While the standard covers seamless handover between heterogeneous technologies the included concepts of media abstraction can be easily extended. As depicted in Figure 2, the Interface Management Function (IMF) is the central messaging component in the control plane at each node. It extends the functionality of the IEEE 802.21 Media Independent Handover Function (MIHF) and supplies an integrative interface to higher level modules. These modules can use a common set of technology agnostic primitives to communicate with each other or with MAC adapters, located logically below the IMF. While the technology specific MAC adapters conduct the management of the radio interfaces, the modules provide all necessary functionalities to setup and manage the network like topology discovery, route calculation or monitoring. For this, the IMF utilizes the messaging services defined by IEEE 802.21. The media independent messaging mechanism can be used by the modules via the *MOD_SAP*[1] and the *AI_SAP*[2] for exchanging commands locally and remotely, either between modules or between modules and MAC adaptors.

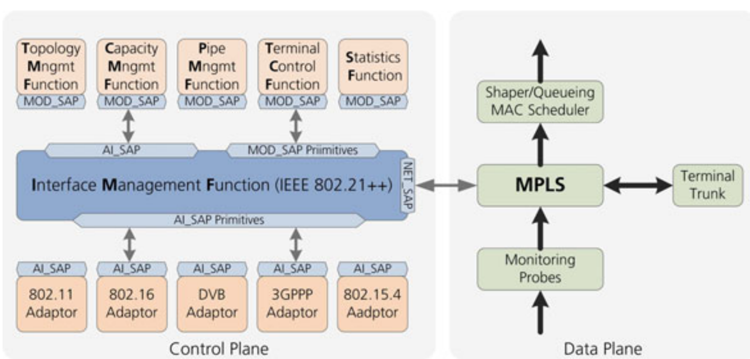

Fig. 2. The WiBACK architecture

[1] IMF Module Service Access Point.
[2] IMF Abstract Interface Service Access Point.

The WiBACK network management is based on a centralized approach. Following the concept of a centralized stateful Path Computation Element (PCE) [4], authoritative *Master* nodes controls the resource allocation and routing state of a set of *Slave* nodes in their administrative area. Backup *Master* nodes can be used to keep the network area operational if a primary *Master* node fails. Thereby the discovery and routing setup process can be abbreviated in such a case. In contrast to well known routing protocols such Open Shortest Path First (OSPF) [5] or Optimised Link State Routing (OLSR) [6] a centralized management offers the opportunity to perform network wide optimizations. Requirements like global radio planning or assigning the overall network capacity to best match payload demands can be fulfilled by the *Master* nodes by using monitoring information regarding the link states and resource allocation. In typical distributed link state protocols, like for example Open Shortest Path First - Traffic Engineering (OSPF-TE) [7], these information have to be kept up-to-date and consistent among all nodes locally to make coherent TE routing decisions. Especially wireless links in an unlicensed spectrum often suffer from volatile conditions so that distributed protocols may be inconsistent or not even converge at all, see [8]. In these scenarios a centralized approach addresses the described challenges much better because only the *Master* nodes need to be kept up-to-date and they can calculate resource allocation based on consistent state.

The centralized approach requires a communication path between a *Master* node and each of its *Slaves* to enable the message exchange for the Command, Event and Information service of the IMF. Therefore *Management Pipes* are setup so that the modules and MAC adapters can communicate to each other on the control plane. On the data plane *Data Pipes* can be established between two arbitrary WiBACK nodes. Such *Data Pipes* can be distinguished on their specific QoS demands such as *Best Effort* or *VoIP*. To fulfill the QoS requirements, monitoring the link state is a major task at each node. Thus within the data plane a monitoring component performs fast per-packet link and LSP measurements. Accumulated data will be pushed to the *Statistics Function* where they are further processed, enabling timely reactions to link degradations or failures can be triggered. For example with the MPLS Fast Reroute (FRR) feature WiBACK pipes can be protected against link failure. Furthermore TE processes to ensure QoS assurances are supported by the usage of MPLS.

The WiBACK network management can be divided into two time scales. While the Topology Management Function (TMF) manages nodes, radio interfaces and spectrum resources on a slowly time scale in a range of minutes, the Capacity Management Function (CMF) assigns the available capacity to resource requests and particularly reacting to capacity changes due to link status fluctuations at a faster time scale in the dimension of seconds. The CMF operates on a set of logical links which is a subset of all physical links managed by the TMF. Both modules push pipe state into the network through the Pipe Management Function (PMF) which modifies existing pipe resource allocations or tear down pipes.

3.1 Topology Management Function

For topology discovery the TMF uses a ring-based approach starting at a *Master* node by setting up his own radio interfaces. To achieve an optimal radio configuration, the local capabilities as well as the ambient spectrum are assessed by a passive channel utilization analysis. On successful completion the *Master* starts sending WiBACK beacons on all configured, active interfaces to inform neighbor nodes (*Slaves*) about its availability.

After the bootstrap phase the *Slave* nodes start scanning on all administratively permitted channels for WiBACK beacons sent by a *Master* node or already associates *Slave* nodes. The scanning process is executed periodically to allow the slave node to react on the scan results. If one or multiple WiBACK beacons from other connected nodes are detected, the *Slave* will try to associate with the highest rated neighbor, determined by *Signal Quality* and *hop distance* to the gateway.

The decision making in the association process is based on local knowledge only and might not be an optimal choice considering overall network topology or other TMF optimization criteria. Therefore the TMF *Master* can reject the actual association request and offer alternative nodes or interfaces to associate with.

The currently used optimization criteria is to establish point-to-point links whenever possible as well as choosing the least occupied channel. To minimize channel interference, a separation of at least 60 MHz for IEEE 802.11 radios using 20 MHz channel bandwidth will be ensured.

In the event of a node or link failure the *Master* will be notified, and the affected links/nodes will be marked as *down* in the topology. If other members of the network are afflicted by that, the *Master* will try to reconnect the afflicted nodes according to his optimization criteria. Whenever a *Slave* detects a connection problem to its *Master*, it will stop broadcasting Beacons and jump back into the *beacon scan* mode and attempt a new association.

3.2 Capacity Management Function

Once a new *Slave* is successfully associated with the WiBACK network, the TMF computes the optimal channel configuration out of all available adjacent WiBACK nodes. The set of *ASSIGNED* links is then pushed to the CMF to establish *data pipes* to or in between nodes to manage capacity allocation. The CMF implementation is a central, stateful PCE using Media Independent Handover (MIH)-style primitives for the messaging of the Path Computation Element Protocol (PCEP). For each link the CMF keeps track of the actual allocated resources and the available resources. In order to maintain an up-to-date state of the network wide resource allocation and the link state under his control, the CMF can subscribe to various events like LINK_STATUS_CHANGED or *pipe* related events.

3.3 Pipe Management Function

PMF implements Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)-style LSP setup and tear down signalling using source-routed IMF messages. Additionally existing LSP resources can be reallocated or altered if needed. Both regular downstream-assigned and upstream-assigned multicast LSPs are supported. PMF uses MIH messages for pipe setup and teardown and allocates resources described via the *TrafficSpecififcations* Type-Length-Value (TLV) at each outgoing interface along the path. This information is used to monitor and enforce the proper QoS-handling of an LSP and Media Access Control (MAC) layer resources by configuring traffic shapers, IEEE 802.16 service flows or IEEE 802.11e queuing parameters.

Classifier Rules are used and maintained to determine which payload from an *Edge Interface* is sent via a specific *Pipe*. These rules can be stateless such as, i.e. a typical IPv4/IPv6 five-tuple rule, or stateful in order to allow a more complex matching. Once a *Pipe* has been set up, the rules can be changed, edited or removed at runtime. A duplex connection between two WiBACK nodes can only be established by configuring a pair of *Pipes*, since each pipe is a unidirectional resource.

3.4 Terminal Control Function

Terminal Control Function (TCF) is an optional component of the WiBACK architecture which provides the functionality required to directly connect User Terminals (UTs) to WiBACK nodes. Depending on the implementation, TCF may provide multiple services such as UT detection and hand-over control as well as local capacity management to match UT traffic demand. The main task is to keep UT/*Pipe* bindings up to date by maintaining proper *Classifier Rules*.

TCF may be implemented to support seamless terminal mobility via for instance, integration with Proxy Mobile IP (PMIP). If mobility is not a major concern, for example, due to rather fixed or nomadic UT usage patterns, less complex approaches such as our QoS-aware LAN Emulation (QLANE)-style mechanism may be implemented.

4 Validation

During the last months the implementation in Maseru has been delayed by unexpected additional costs for waveband licensing. This is still under negotiation with the Lesotho Communication Authority and might break the business case for the local entrepreneur. Therefore we decided to present in this section the obtained results in our outdoor testbed which was build to match the planned Maseru scenario as closely as possible. The main difference is that the distances between nodes are rather in the range of 10...20m instead of 1...2km. In order to provide similar connectivity among the nodes, we have used directional antennas and adopted the transmit power accordingly. Despite those efforts the overall connectivity is higher compared to what we expect for the actual deployment site, but the higher quality links are the same as planned for Maseru.

4.1 TMF Validation

We have evaluated if the TMF properly detects and configures the six-node topology assuming a blackout scenario where all nodes are switched on at the same time. With the current parameterization regarding channel scan and channel analysis timing, it took between three and five minutes for TMF to completely configure the topology and to install the additional *Best Effort* and *VoIP data pipes*. Figure 3 depicts the link and channel configuration after a completed discovery process where TMF has properly detected and configured point-to-point links and assigned channels maintaining a minimum of 60 MHz center frequency separation.

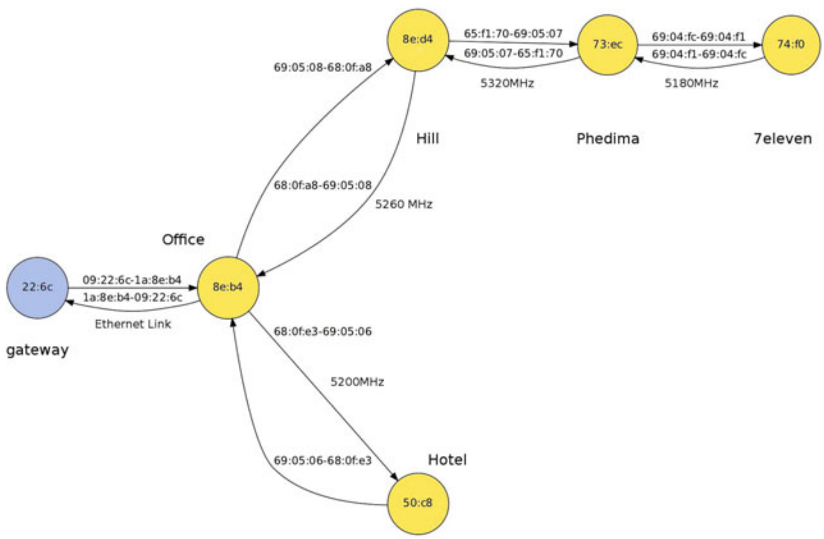

Fig. 3. The TMF Master at node 6c:22 (*Gateway*) has successfully discovered and configured the network by forming point-to-point links

4.2 TMF Evaluation

We have validated the recovery procedure of the Topology Management Function in case of node failures to verify that TMF properly reconnects all remaining nodes. As a second test we've performed an initial evaluation of the recovery functionality to quantify the down times of parts of the network depending on the failing node. For this we disconnected each of the nodes ten times and measured the duration until the complete network was recognized and connected again (Table 1) independently of the discovered topology. Though the network usually finds the same arrangement based on the given environmental circumstances.

The result shows that the network always reconnects whereat the duration varied distinctively. On the one hand this leads back to the periodic scanning behaviour (about 40s, depending on the possible frequencies) of all disconnected nodes, on the other hand the duration mainly depends on the disconnected node and the count of nodes on the path behind. E.g. disconnecting the *Hill* (8e:d4) also cut off the *Phedima* (73:ec) and *7eleven* (74:f0), based on the ring-based algorithm described in section 3.1, the *Hill* get reconnected than *Phedima* and finally *7eleven*.

Table 1. Measured times to complete network recovery after temporary disconnecting a single mesh node

Node	arg/stddev	Samples								
8e:b4	234.5 ± 33.0	246.2	275.3	219.3	245.9	177.3	233.3	284.0	209.3	219.9
8e:d4	213.8 ± 35.0	243.7	193.9	242.4	196.7	208.5	174.4	279.6	175.3	209.4
50:c8	46.1 ± 4.1	41.1	45.4	49.3	51.0	50.5	49.9	41.4	44.6	41.4
73:ec	46.0 ± 6.6	58.0	36.1	42.1	44.6	45.9	43.3	42.8	54.5	46.4
74:f0	60.9 ± 5.5	58.5	63.9	53.4	59.6	59.1	68.2	69.1	61.8	54.2

5 Conclusion and Future Work

We have described the use case of our WiBACK pilot in Maseru, Lesotho, our WiBACK architecture and have shown that it addresses the technical requirements while supporting a rather low cost implementation. Using a testbed resembling the actual deployment scenario, we have verified that our Topology Management Function properly detects and configures the given topology in order to provide collision free wireless links among the involved nodes. The performed measurements exhibit the duration of network reconfiguration periods as a reaction to node failures. With the conservative parameterization it took between 46 seconds and 235 seconds to reconfigure the network.

Future work will focus on tuning the periodical channel scanning mechanism in order to reduce the reconnection times. Further within the framework of the *SolarMesh*[3] project we will exploit how far energy-awareness can lower the overall energy footprint or improve the systems independence from a power grid.

Acknowledgement. This work has been funded by the Federal Ministry of Education and Research of the Federal Republic of Germany (Förderkennzeichen 01 BU 1116, SolarMesh - Energieeffizientes, autonomes großflächiges Sprach- und Datenfunknetz mit flacher IP- Architektur). The authors alone are responsible for the content of this paper.

References

1. IEEE standard for information technology-telecommunications and information exchange between systems-local and metropolitan area networks-specific requirements - part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications, IEEE Std 802.11-2007 (Revision of IEEE Std 802.11-1999), pp. C1–1184 (December 2007)

[3] http://www.solarmesh.de

2. Banchs, A., Bayer, N., Chieng, D., de la Oliva, A., Gloss, B., Kretschmer, M., Murphy, S., Natkaniec, M., Zdarsky, F.: Carmen: Delivering carrier grade services over wireless mesh networks. In: Proc. IEEE 19th International Symposium on Personal, Indoor and Mobile Radio Communications PIMRC 2008, September 15-18, pp. 1–6 (2008)
3. IEEE standard for local and metropolitan area networks- part 21: Media independent handover, IEEE Std 802.21-2008, pp. c1–301 (January 2009)
4. Farrel, A., Vasseur, J.-P., Ash, J.: A Path Computation Element (PCE)-Based Architecture. RFC 4655 (Informational) (August 2006)
5. Moy, J.: OSPF Version 2. RFC 2328 (Standard). Updated by RFC 5709 (April 1998)
6. Clausen, T., Jacquet, P.: Optimized Link State Routing Protocol (OLSR). RFC 3626 (Experimental) (October 2003)
7. Katz, D., Kompella, K., Yeung, D.: Traffic Engineering (TE) Extensions to OSPF Version 2. RFC 3630 (Proposed Standard). Updated by RFC 4203 (September 2003)
8. Bernardos, C.J., Fitzpatrick, J., Kuo, F.-C., Kretschmer, M., Lessmann, J., Niephaus, C., de la Oliva, A., Robitzsch, S., Zdarsky, F.: Carrier-grade wireless mesh networks: D3.4-unicast and multicast routing specification and analysis (December 2009), http://www.ict-carmen.eu/wp-uploads/2009/D3.4.pdf