An OpenFlow Based Cellular Backhaul Network

Subir Varma

1.0 Introduction

Cellular backhaul networks are assuming increasing importance as the amount of data traffic flowing over cellular networks increases. The issue that usually gets the most attention with respect to these networks, is that of bandwidth capacity. This is being addressed by operators, as they deploy technologies such as fiber and high capacity point to point wireless links in their backhaul networks. However even with the increased network capacity, there are some problems with the way these networks are architected that prevents them from fully supporting the applications and traffic patterns that one is likely to see in future cellular networks. In particular, even though all traffic is carried over packet networks, the use of tunneling limits the flexibility with which the mobile devices can send traffic to each other, and also with caches or CDNs that may exist in the backhaul network.

The objective of this report is to propose an alternative OpenFlow based architecture for LTE based cellular backhaul networks. By using OpenFlow as the control plane, we will show how some of the problems with the current backhaul network architecture can be overcome, thus resulting in a network that enables more flexible traffic flows. This architecture leaves the control interfaces on the mobile devices unchanged, and replaces the existing control interfaces on the Base Stations and Service Gateways by OpenFlow. The current control node in the LTE network, called the MME, becomes an application in the OpenFlow controller.

The rest of this report is organized as follows: In Section 2 we give a short description of the LTE backhaul architecture, in Section 3 we point out the problems with this architecture, in Section 4 we introduce an OpenFlow based LTE backhaul architecture with both Layer 2 and Layer 3 based designs, in Section 5 and 6 we show how the proposed architecture can be used to support mobility and intra-backhaul communications respectively, in Section 7 we briefly describe some future evolutions of the proposed architecture by making the connection with the work being by the Distributed Mobility Management group in the IETF and finally in Section 8 we talk about future directions for the architecture.

2.0 LTE Cellular Backhaul Architecture

2.1 The Physical Network Topology

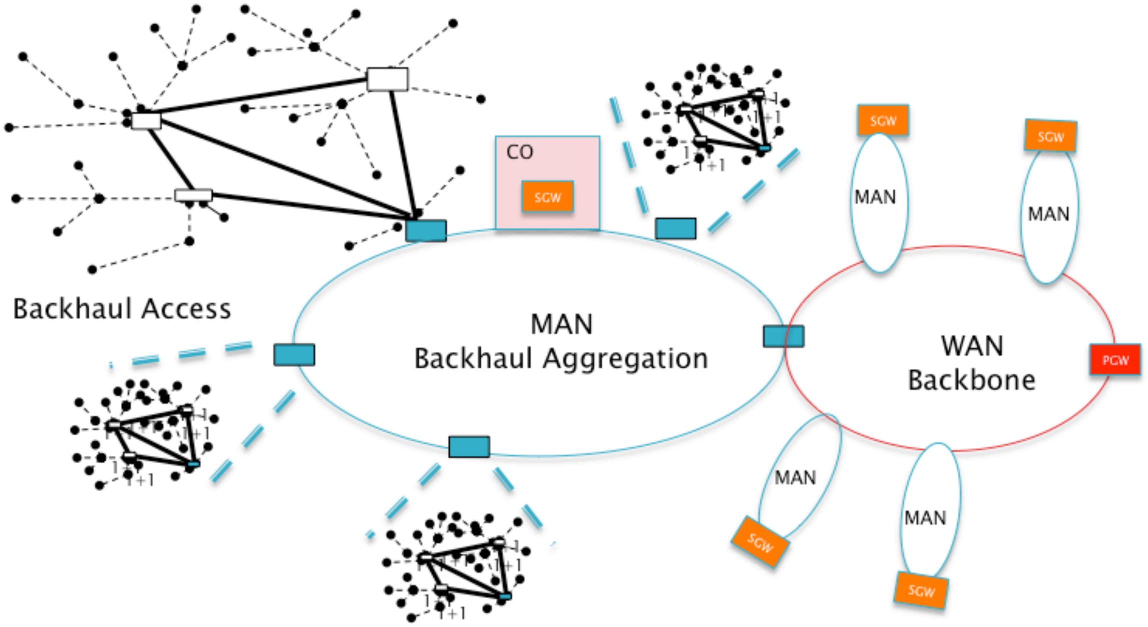

Figure 1: A Typical Cellular Backhaul Network

Figure 1 shows the physical topology of a typical cellular backhaul network. The backhaul network can be divided into three parts, namely the Access, Aggregation and Backbone. The access network is shown towards the left of the figure, and in this example it consists of point to point wireless links. These links can be aggregated using a ring or mesh topology as shown in the figure, which can either use wireless or fiber. The access network in turn is backhauled using a metropolitan area network (MAN), which is typically fiber based. There may be one or more Central Offices (COs) in the metro area, where the operator can place nodes such as Service Gateways (SGW). The MANs are connected to a regional or national fiber ring as shown, and nodes such as the PDN Gateway (PGW) are located in larger COs (node shown in red in Fig 1), also called National Data Centers.

2.2 Higher Layer Protocols Used in the Backhaul Network

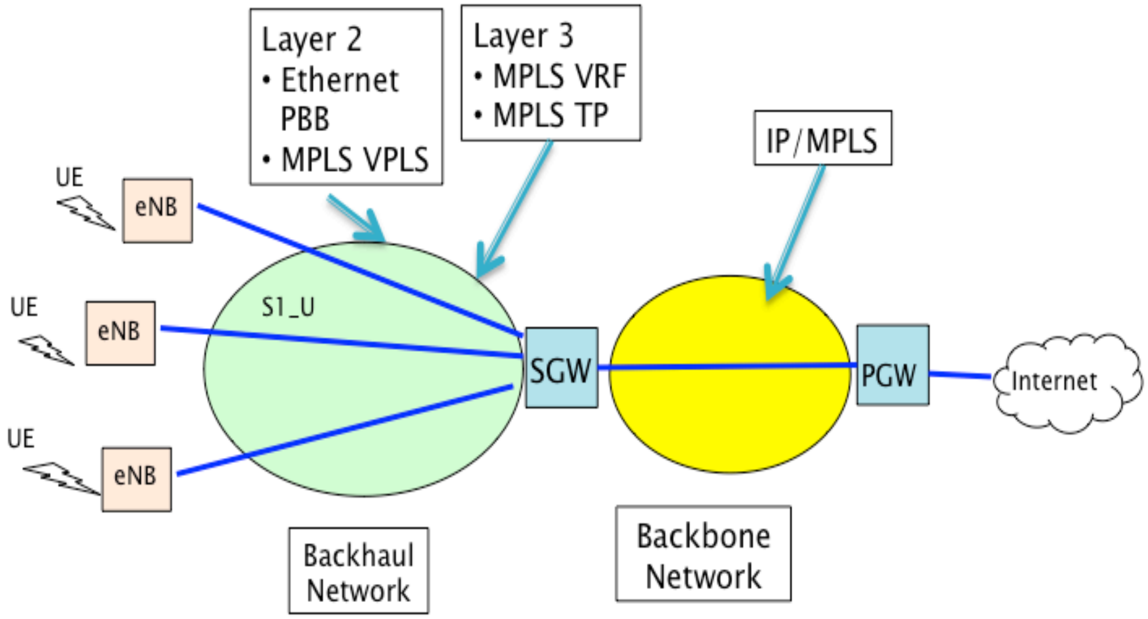

Figure 2: Higher Layer Protocols used in the Backhaul

Traditionally the backhaul network was dominated by TDM circuit switching, but with the advent of 3G and 4G networks, operators have begun to deploy various packet switching technologies in the backhaul. A popular choice has been Multi-Protocol Label Switching (MPLS), it main attraction being its ability to carry TDM traffic (for legacy 2G networks) and ATM traffic (for 3G networks) by using Pseudo Wire Emulation (PWE3) technology. MPLS Label Switched Paths (LSPs) are set up across the backhaul using configuration management tools, with QoS differentiation provided using DiffServ classes mapped to the MPLS EXP field. In order to simultaneously support multiple operators and/or multiple services across the backhaul, MPLS can be deployed either at Layer 2 using the Virtual Private LAN Service (VPLS) protocol, or Layer 3 using the Virtual Routing Function (VRF) protocol. Note that Ethernet is used as the underlying transport to carry MPLS in this architecture. In order to provide resiliency to link or node failures, operators have the option to use the MPLS Fast ReRoute (FRR) protocol. However FRR requires the configuration of a large number of backup LSPs, hence most operators rely on the failure recovery mechanisms available in the underlying Ethernet layer, such as Ethernet Ring Protection Switching (ITU G.8032).

There are two alternatives to using MPLS in the backhaul, that operators have also begun to deploy. The first one is a Layer 2 protocol called Ethernet Provider Backbone Bridging (PBB), also known as Ethernet MAC in MAC. A motivation for preferring Ethernet PBB over MPLS is to simplify the backhaul network (and thus reduce its cost) and also since the requirement to carry 2G TDM or 3G ATM traffic is becoming less of an issue as these older networks are phased out. In these networks packets are switched using statically configured MAC tables. The other protocol that is starting to be used in the backhaul is MPLS-Transport Protocol (MPLS-TP). This protocol uses MPLS framing over Ethernet, and utilizes statically configured LSPs. Its biggest difference as compared to regular MPLS is in its failure recovery mechanisms. This is done using in-line OAM packets that are sent at short intervals though the LSPs, thus allowing the system to quickly react to failure events.

As shown in Figure 2, the backbone network being made up of very capacity fiber rings, uses the regular MPLS protocol.

2.3 LTE Tunnel Overlays Across the Backhaul

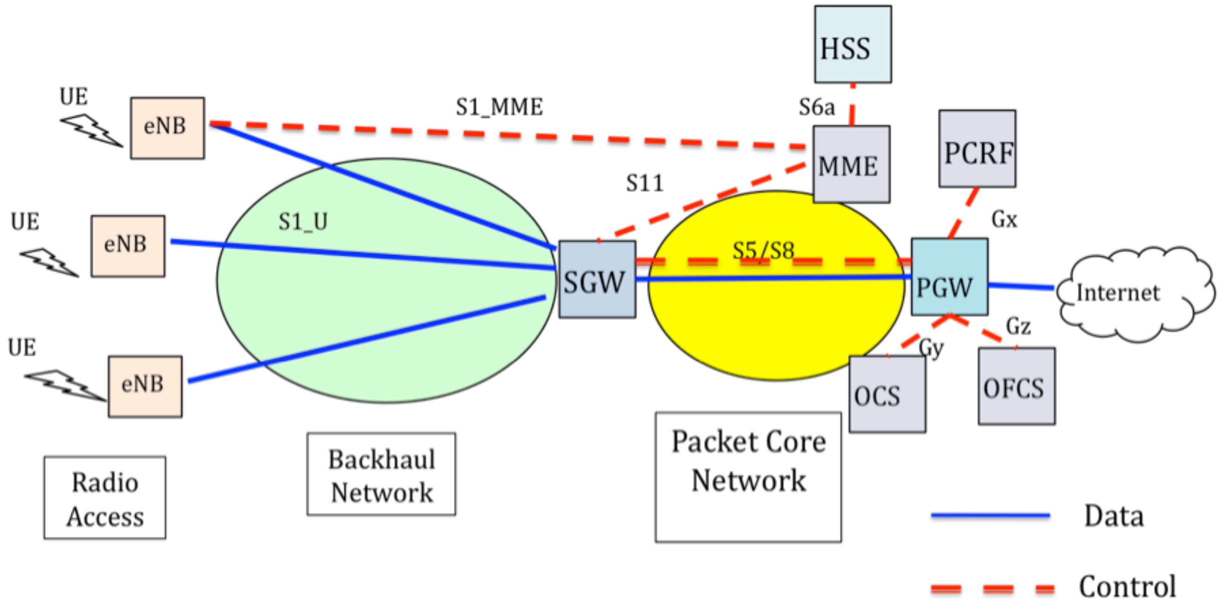

Figure 3 LTE EPC based Tunnels in the Cellular Backhaul

The backhaul network described in previous sections cannot be used to natively carry the IP packets to the UEs (or mobile devices). This is due to the fact that the UEs are moving around from eNB (or Base Station) to eNB, and packets that are addressed to that UE have to be constantly re-routed though the backhaul after taking this mobility into account. Note that regular IP or MPLS provides no support for mobility based re-routing, and the solution that LTE (and other wireless networks) have been using to solve the mobility problem is based on the use of dynamic tunnels across the backhaul network. These tunnels form an overlay network on top of the backhaul network, and LTE has defined new control protocols to do operations such a tunnel setup and teardown, a well as re-routing of tunnels in response to mobility.

Figure 3 shows the structure of the LTE backhaul network, with all the overlay network data and control nodes. The MME, PGW and SGW are the most important nodes in this architecture and their functions are described below:

- The Mobility Management Entity (MME): This is the main control node in the LTE network and is responsible for signaling the various other nodes to setup and teardown tunnels. The MME also signals the SGW to switch tunnels in response to UE mobility (using the S11 signaling interface). In addition to its tunnel related signaling functions, the MME also acts as a go-between for the UE to HSS signaling during UE network entry and authentication (using the S1-MME and S6a signaling interfaces). The MME is also responsible for paging idle mode UEs after new data is detected over the S1-MME interface.
- The PDN Gateway (PGW): This nodes functions as the fixed anchor point for all traffic going to an UE. It advertises the ability to reach the IP addresses that have been allocated to the UEs (also called home IP address) to the rest of the world, as a result all traffic addressed to the UEs ends up at the PGW. The PGW then tunnels the IP packet by encapsulating it in a 3GPP defined protocol called GTP, and sends it to the SGW. The PGW is also responsible for interfacing with the Policy function (through the PCRF node) and Billing functions (through the OCS and OFCS nodes) in the network, and for enforcing QoS rules on traffic flowing through it.
- The Service Gateway (SGW): The SGW serves as an intermediate point on the tunnel system from the PGW to the eNB. On receiving a packet over the GTP tunnel from the PGW, it extracts it and sends it over a second GTP tunnel that goes all the way to the eNB that the UE is attached to. The operator typically splits up all the eNBs within the network into several groups, and assigns a SGW to serve each group, such that any eNB is not allowed to be connected to two distinct SGWs. The SGW serves as an anchor point for UE mobility within it group, such that the GTP tunnel between the PGW and SGW is left unchanged and only the tunnel between the SGW and the eNB needs to be switched on user mobility (this is usually referred to as micro-mobility initiated tunnel switching). If the UE moves between groups, then the PGW also needs to switch its tunnel to the new SGW (this is referred to as macro-mobility initiated tunnel switching).

In addition to these three nodes, LTE also requires the following: The Home Subscriber System (HSS) node: This node is responsible for storing user subscription information as well as the shared secret that is used during user authentication.

The Policy Control Resource Function (PCRF): This node is responsible for enforcing dynamic QoS based policies for traffic flowing though the PGW. In particular, if the UE initiates an application that requires guaranteed QoS, then the application can signal its requirements to the PCRF, which then commands the PGW to set the tunnels with the required QoS characteristics.

Online Charging System (OCS) and Offline Charging System (OFCS): These nodes take part in billing, and interface with the PGW to acquire the necessary traffic information needed to carry out their functions.

3.0 Shortcomings of the Current LTE Architecture

In this section we briefly three different problems that arise as a result of the way in which the LTE tunnel overlay network has been designed. In general these issues arise not just in LTE but also in other networks that use a similar mechanism to support mobility, including WiMAX, UMTS etc. In the next and subsequent sections we show how these problems can be avoided if OpenFlow is used as the control protocol in the backhaul network.

- The Routing Problem: Referring back to Fig 3, the routing problem arises due to the fact that all traffic going to an UE has to pass through the PGW. In a typical LTE implementation, there are not more than a handful of PGWs covering the entire country, hence it is quite likely that the closest PGW to an UE may very well be in the neighboring state. This leads to a very sub-optimal routing path in the following cases:
    - UE to UE communications: If the UEs are located close to each other geographically, then it does not make sense to route their traffic through a PGW located far away. In addition to common applications such as voice traffic, many emerging applications such as Machine to Machine generate traffic that have a peer to peer characteristic.
    - Local Caching and Content Delivery Network (CDN) Traffic: As a result of the current architecture, any cache or CDN node has to located north of the PGW. In order to take better advantage of the fact that he controls the backhaul network, the service provider would like to locate these nodes closer to the edge of the network where the users are located. With the increase in importance of cloud based application delivery, optimally placing the caching and CDN resources in the network will be a critical issue in the future.
- Tunneling overhead: The GTP tunneling protocol that LTE uses adds about 36 bytes of overhead to each IP packet: 28 bytes for UDP/IP + 8 bytes for the GTP field. A previous study [?] has estimated the GTP overhead to be 14%, given an average packet size of 250 bytes. The overhead for IPv6 is even higher at 22%.
- Single point of failure: Note that all traffic going to an UE has to pass through the PGW, and given a large network with hundreds of thousands of devices,

the PGW becomes a very vulnerable node the failure of which can bring down the entire network.

- Scalability Problem: The PGW also constitutes a scalability problem due to the centralized architecture, as the size of the mobile nework continues to grow.

The problems with centralized tunnel based architectures are well known in the industry and have been pointed out by others [?]. The IETF has established the Distributed Mobility Management (DMM) group to investigate alternative designs that don't have these shortcomings. We will discuss these in more detail in a later section, where we will show that the proposed OpenFlow based approach is superior to the proposals that we have seen so far.

4.0 OpenFlow Based LTE Backhaul: Phase 1

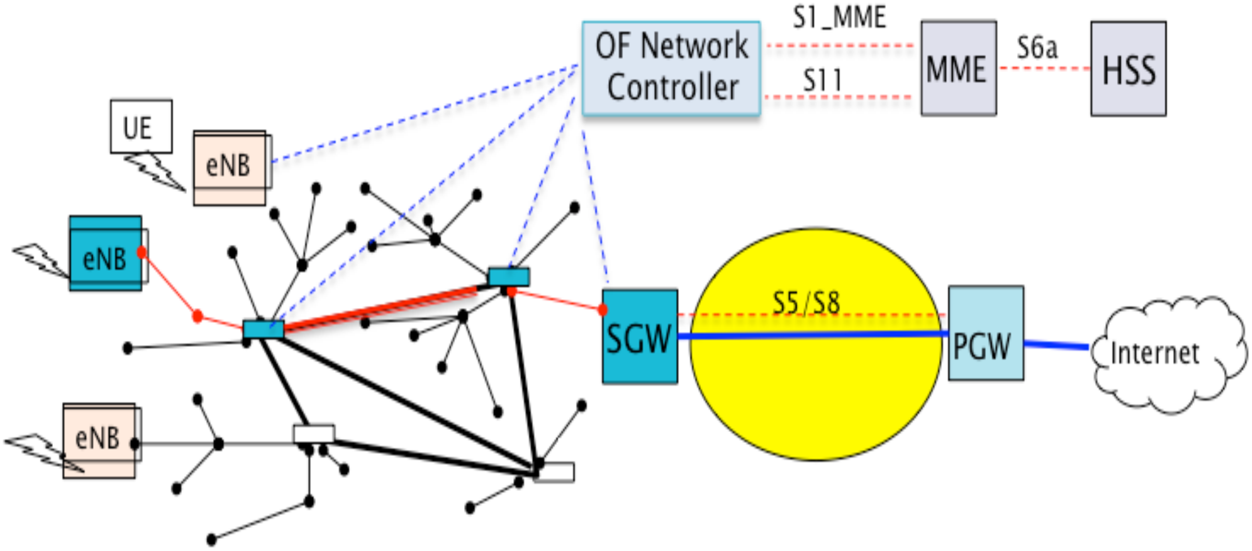

Figure 4: Proposed OpenFlow based LTE Backhaul Architecture

The Phase 1 of the proposed OpenFlow based LTE backhaul architecture (Fig 4), is characterized by the following:

- The backbone portion of the network, between the PGW and SGW is left un-touched. Note that this part of the network is characterized by a high capacity regional or national optical fiber rings, with high end routers supporting millions of flows. In general, the current generation of OpenFlow switches and controllers may not have the capacity to handle these types of networks currently.
- An OpenFlow controller is introduced to control the backhaul portion of the network, between the SGW and the eNBs. It establishes control connections to each node in the backhaul network, including the eNBs and SGW. This means that the S1-MME interface beween the eNB and the MME and the S11 interface between the SGW and the MME are replaced by OpenFlow.
- The S1-AP control interface between the UEs and the MME is left unchanged, so that no changes are required in the LTE software installed on the devices.
- The OpenFlow controller acts as a proxy for the backhaul network, for the MME. It establishes S1-MME and S11 connections to the MME, which may be internal interfaces if the MME is an application running in the controller.
- The IP address of the UE serves as its identifier, while the identity of the eNB it is attached to (MAC or IP address), serves as its location identifier.

- The OF Controller snoops on the control message traffic to obtain the information needed to discover the current location (i.e. eNB it is attached to) of each UE and to set up the data path between the SGW and the eNB.
- There is no tunneling required to move packets through the backhaul if a Layer 2 architecture is used (see Section 4.1). The OpenFlow controller installs the switching rules in the backhaul nodes that are needed to establish paths between the SGW and each eNB, and also in-between eNBs (to support the X1 connection). Note that these paths may be pre-defined by the controller, and hence don't have to be set-up every time a UE enters a network or moves around between eNBs.
- The native IP packet can be encapsulated either in Layer 2 Ethernet MAC headers (further described in Section 4.1), or in Layer 3 VxLAN or NvGRE headers (described in Section 4.2).

In the next two subsections we describe the Layer 2 and Layer 3 backhaul architectures in more detail.

4.1 Layer 2 Based Backhaul Architecture

The Layer 2 based backhaul architecture is characterized by the following:

- All packet forwarding within the backhaul network is based on Layer 2 Ethernet MAC addresses. In order to differentiate between packets from multi-operators and also QoS requirements within the backhaul network, the VLAN Q in Q (IEEE 802.1ad) header can be used.
- The OpenFlow controller maintains a table that maps the IP Address of the UE to the Ethernet MAC address of the eNB it is attached to. This mapping is initially setup when the UE first enters the network (or is woken up from idle mode), and then modified as the UE moves around. The controller is able to get this information by snooping on the control messages flowing from the UE to the MME over the S1-AP interface.
- In the downlink, the SGW decapsulates the IP packet from the GTP tunnel, and based on the information in its flow table (which maps its IP address to the eNB it I attached to), it attaches the appropriate Ethernet MAC header to the packet. The Ethernet Source MAC Address is set to that of the SGW, while the Destination MAC Address is set to that of the eNB.
- If the UE moves to another eNB, then the OpenFlow controller re-programs the flow table in the SGW using the OpenFlow interface, such that the destination Ethernet MAC address is now set to that of the target eNB.
- The OpenFlow controller sets up the forwarding rules in all the internal backhaul switches based on the destination MAC address in the Ethernet header. Note that these rules can be setup in advance of the actual traffic flows, so that the controller does not have to communicate with any of the

internal switches when an UE enters or leaves the network (or changes eNBs). All that the controller has to do in these cases, is to change the flow table rules in the SGW.

- In the uplink, the eNB receives the IP packet from the UE over the air interface, and depending upon its destination IP address, either sends it to the SGW or to an internal destination within the backhaul network itself. This determination is programmed into the flow tables in the eNB by the OpenFlow controller in the following way: If the destination IP address exists in a network that is internal to the backhaul, then the controller supplies the eNB the appropriate destination MAC address for that node. If not, then the destination MAC is set to that of the SGW, which on receiving the packet encapsulates it in GTP and forwards it on to the PGW. Note that this flexibility to switch the packet to an "internal" destination is not available in the current network architecture. This ability enables a number of very useful enhancements to the backhaul architecture, which are further explored in Section (?).

4.2 Layer 3 Based Backhaul Architecture

It is also possible to establish Layer 3 based forwarding in the backhaul network, using the (private) IP addresses of the eNB and other nodes, rather than their MAC addresses. This can be done by encapsulating the IP packet in a Layer 3 envelope, using protocols such as VxLAN or NvGRE. Note that since the UDP/IP/VxLAN header takes up 36 bytes, there will be no reduction in header overhead as compared to the GTP case.

The Layer 3 based architecture is identical to the one described in the previous section for Layer 2, except for the following:

- The IP address of the eNBs is used as the location identifier for the UEs. The OpenFlow controller uses this IP address to program forwarding rules in the flow tables.
- The 3 Byte VxLAN ID is used to differentiate between QoS clases and multiple operators using the same backhaul.

Note that this architecture has some similarities to the one proposed for controlling intra-data center communications by Nicera. In both cases the OpenFlow controller is responsible for routing the packet to the appropriate tunnel, and also appending the correct VxLAN ID to the packet.

5.0 Support for Mobility

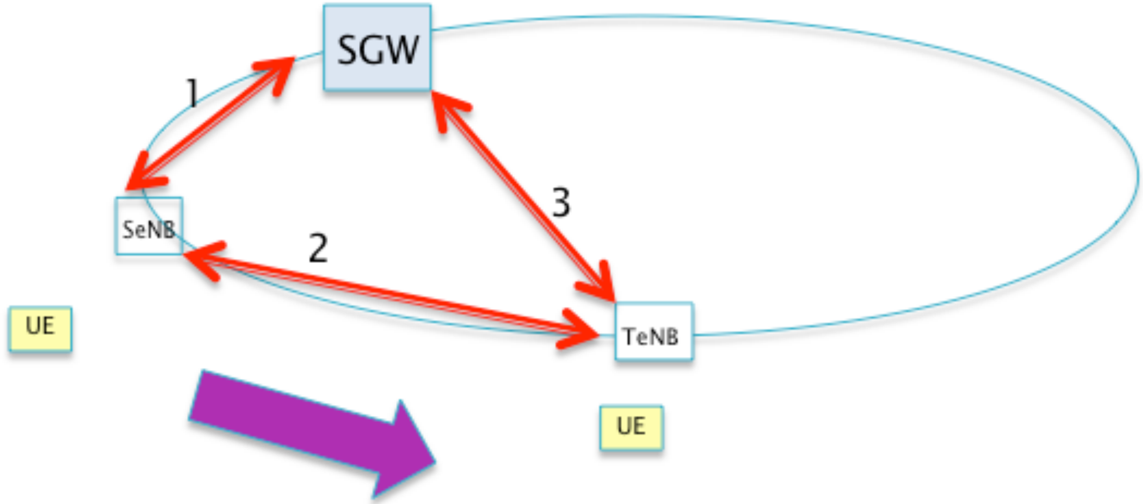

Figure 5

In order to avoid dropping packets during handovers, the LTE mobility management protocol provides a Handover Preparation phase which takes place while the UE is still attached to the source eNB (step 1 in Fig 5) and involves the following:

- The source and target eNBs exchange control information that facilitate the quick entry of the UE in the target eNB airlink
- The source eNB starts to tunnel downlink packets to the target eNB, rather than transmitting them over the airlink.

These communications are carried out over the temporary X2 connection between the source and target eNBs (step 2 in Fig 5). Once the UE is fully connected to the target eNB, the SGW switches its GTP tunnel to the UE under the direction of the MME, so that it terminates at the target eNB and the old tunnel is torn down (step 3 in Fig 5).

All these procedures carry over to the OpenFlow controlled backhaul network as well. The controller can pre-define the X2 paths between neighboring eNBs, so that they don't have to set up at the time of handover. Also, instead of switching the GTP tunnel at the SGW using the S11 interface, the controller changes the destination MAC address that the SGW uses, in order to switch the path to the target eNB.

6.0 Support for Local Caching/CDNs

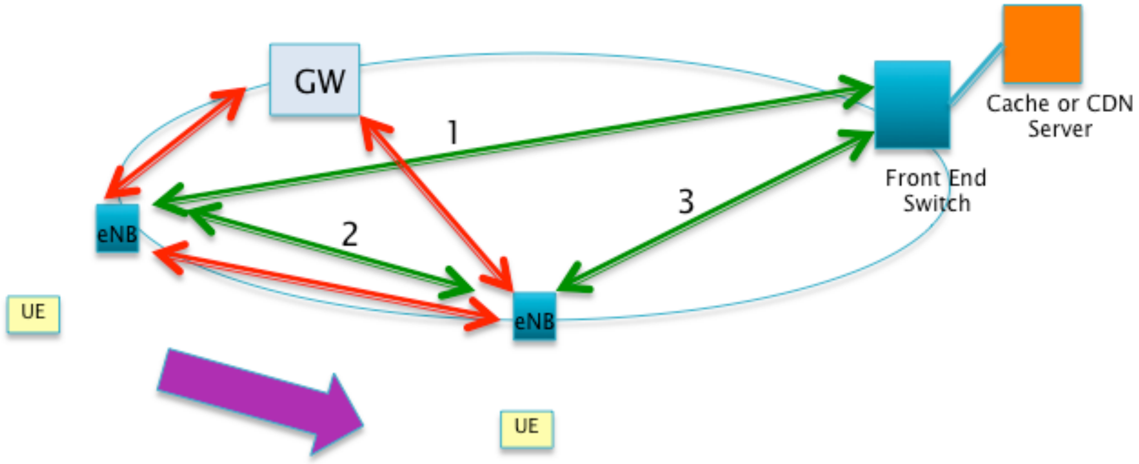

Figure 6: Local Communications within the Backhaul Network

The OpenFlow based control of packet forwarding within the backhaul network enables the UEs to establish communications with other UEs and nodes in the backhaul network, without having to go through the PGW. Establishing these paths within the backhaul network using OpenFlow is fairly straightforward, a more interesting problem is that of changing these paths in response to UE mobility and achieving lossless handovers.

As shown in Fig 6, the server is located behind an OpenFlow controlled switch that is responsible for forwarding the IP packet to the eNB to which the UE is attached. If the UE changes eNB while the server session is still in progress, then the packet forwarding is changed in 2 steps as shown in Fig 6. In the first step, while the UE is still attached to the source eNB, downlink packets from the server are forwarded to the target eNB over the temporary X2 connection. After the handover is complete, the OpenFlow controller changes the destination address at the Front End Switch, to switch the server packers directly to the target eNB.

The same process described above can be followed for UE to UE communications.

Note that thanks to OpenFlow, the operator has full visibility into all intra-backhaul communications, which can be subject to billing or other quota limits by the operator. This is not the case with some proposals that also allow intra-backhaul communications, such as the Distributed Mobility Management work in the IETF.

7.0 OpenFlow based Backhaul: Phase 2

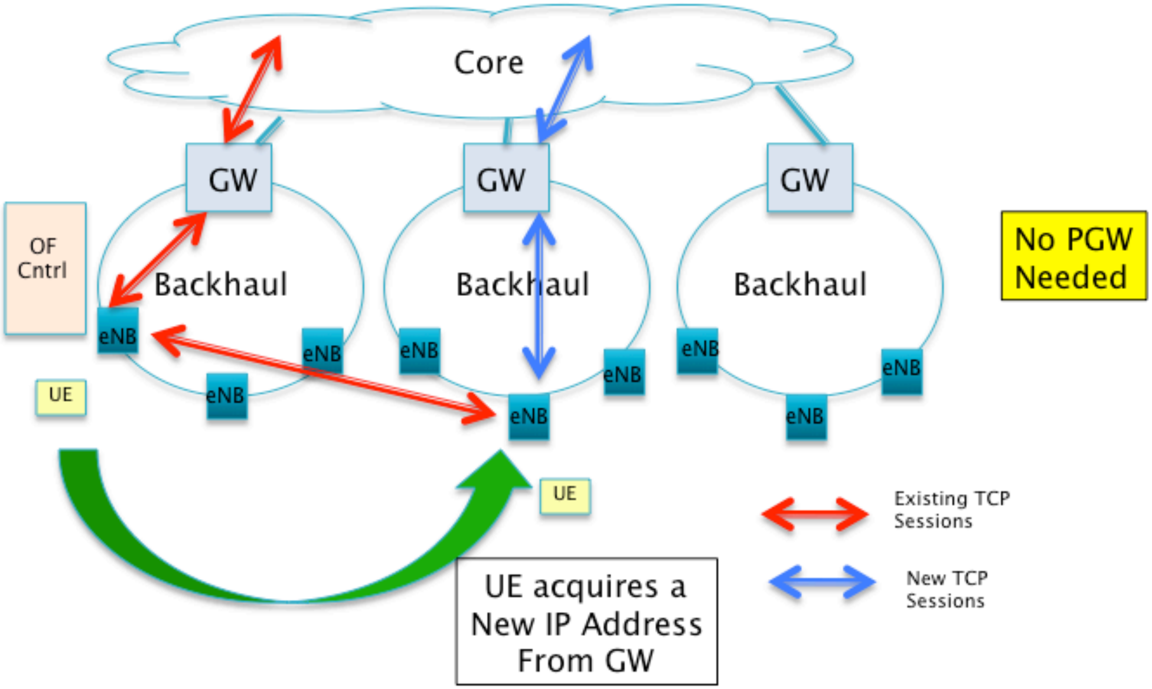

Figure 7

The architecture described in Section 4 does not address the issue of PGW scalability and single point of failure. In this section we briefly outline some ideas that can be used to solve this problem.

As shown in Fig 7, we propose a new Gateway node, that can be thought of as the SGW and PGW combined into one node. Just like the PGW, this node advertises the reachability to certain IP networks to the rest of the Internet, and like the SGW, it acts like a gateway for geographically contiguous number of eNBs.

When a UE moves within the area covered by a single GW, then the mobility management is the same as described in Section 5. However if the UE moves to an eNB that is in the domain of another GW, then the UE acquires a new IP address that belongs to the network supported by the target GW. However it does not give up its old IP address right away, but continues to use it as long as there are active sessions using that address. The traffic on these sessions is tunneled to the source eNB, and the tunnel is maintained as long as there is or more of these sessions still active. However all new sessions us the new IP address, and use the target GW for their traffic.

The basic idea described here, that of changing IP addresses on mobility, is the same as that being proposed in the IETF Distributed Mobility Management (DMM) Group [?]. However, while the DMM Group is proposing to change IP addresses whenever the UE changes eNB, we are restricting the address change only to the case when the UE crosses GW domains. This seems to be a more practical design, since changing IP addresses is a time consuming process.

Note that the design proposed here changes the signaling software on the UE, as well as the address allocation and tunneling mechanisms, in contrast to Phase 1 which left these elements unchanged. In order to be deployed in a practical LTE network, these ideas will have to be adopted by the 3GPP standards body, that is responsible for LTE specifications.

8.0 Future Work

The following additional areas need to be addressed to complete the Phase 1 architecture:

- QoS in the Backhaul: We need to make sure that all the QoS constructs supported in the LTE backhaul, are also supported in the OpenFlow controlled network.
- Handling Network Link or Node Failures: The objective here is to make sure that the disruption due to the failure is less than 50 ms, which is also the goal of protocols such as MPLS-TP currently being proposed for the backhaul.
- Policy and Billing Architecture: The policy and billing architecture needs to extended in order to take care of the intra-backhaul network traffic.

References

1. "OpenFlow Switch Specification", Version 1.3, Available at www.opennetworking.org, Apr. 2012.