A Layer-2 Approach for Mobility and Transport in the Mobile Backhaul

Nuutti Varis, Jukka Manner
Department of Communications and Networking
Aalto University, Finland
nuutti.varis, jukka.manner@aalto.fi

Johanna Heinonen
Nokia-Siemens Networks
Espoo, Finland
johanna.heinonen@nsn.com

ABSTRACT

GTP tunnels are used to transport user data packets in LTE networks over the mobile backhaul between the base stations and core network elements. Maintaining the tunnel as the user equipment moves requires heavy signaling. Moreover, the protocol stack effectively adds headers and thus introduces overhead. In this paper, we design and analyze an LTE architecture and mobility management solution that is based on layer-2 switching. We remove the GTP tunnel and its associated IP/UDP tunnel, and embed user IP packets directly in Ethernet frames within the mobile backhaul. Mobility management within an access network is distributed. We show that the design is very efficient and would simplify the future LTE evolution.

1. INTRODUCTION

The traffic capacities of the forthcoming radio access networks (LTE and LTE-A) are expected to increase significantly. This is driven by the huge increase in mobile data usage, predicted by vendors, such as Cisco, Ericsson and NSN, to increase by over 10 fold in five years. As the capacity increases, the cell sizes will decrease, leading to much more frequent handovers. Furthermore, the expected, and already experienced, strong growth of traffic in data networks [1] is tied to high pressure on transport costs – this evolution is challenging.

Also standardization in 3GPP is moving towards converged packet services. Therefore traditional transport network technologies are not sufficient in the mobile backhaul. Mobile operators need to find alternative solutions – one possibility as a transport network technology will be Ethernet and/or Carrier Ethernet [2].

Ethernet/Carrier Ethernet is already used as a transport in the mobile backhaul, but it is worth examining the possibilities to implement mobility management functions in the link layer without user specific tunnels as an alternative to the current IP-based solution. Mobility solution without tunneling in 3GPP was already discussed in 2005 [3] but 3GPP did not select that option. Standard Ethernet in the access network might not be optimal because of its poor scalability characteristics. A new Ethernet bridging protocol, the IETF-driven TRILL is an attempt to solve some of the problems of the standard STP-based Ethernet.

The contribution of this paper is an analysis of an alternative solution to packet transport and mobility management in the LTE backhaul. We focus mostly on the access network but also touch the core network. We design a technique to use TRILL with an additional extension using Distributed Hash Tables (DHT), and employ Carrier Ethernet in the mobile backhaul in 3GPP release 8 networks. We show that our scheme provides lower overhead and a simplified implementation compared to the current standard. Handover related signaling seen by the 3GPP core network elements is decreased as well. However, the solution does not affect the signaling load caused by activating and deactivating bearers.

2. MOBILITY MANAGEMENT IN 3GPP R8

3GPP release 8 specifications 23.401 and 23.402 [4] [5] introduce a packet only flat network architecture called Evolved Packet System (EPS). Enhanced Node B (eNB) elements in the radio access network are connected to the evolved packet core (EPC) via the S1 interface. Control plane is connected to the Mobility management entity (MME) element and user plane to the serving gateway (S-GW) element. The packet data network gateway (PDN-GW) allocates IP addresses for the user equipments (UEs) and it has an interface to external IP networks. The S-GW and PDN-GW can be implemented in the same physical network element.

A UE has to register itself to the network to access services. The registration is done during the attach procedure, at the same time the IP connectivity for the UE is enabled by establishing a default bearer. In addition one or more dedicated bearers can be activated. Data flows mapped to the same EPS bearer have the same QoS, i.e. the same scheduling policy, queue management and traffic shaping policy. Traffic flow templates in the UE and PDN-GW map traffic to the right bearer. User data is thereby always tunneled inside the EPC.

GPRS tunneling protocol (GTP) is used to implement EPS bearers in the S1 interface. The fundamental point of

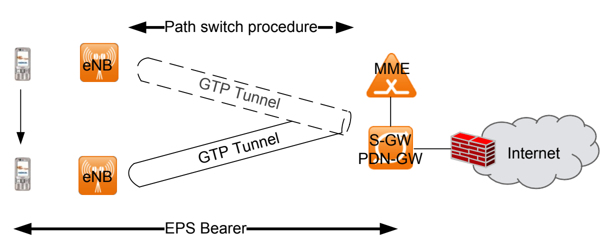

Fig. 1 Handover procedure in 3GPP R8 networks

GTP is to tunnel the user packets from the PDN-GW to the right eNB. The GTP tunnel(s) are set up when the attach procedure is successful. During handovers the eNB side endpoint of the tunnel will change. The new tunneling endpoint identifier (TEID) is signaled to the core network during a path switch procedure (Figure 1).

GTP is a simple protocol that has proven to fill the needs it was designed for in the late 90's. However, it has two clear downsides. First of all, as the tunnel needs to be re-created after every handover, the signaling load due to the handovers increases when the cell sizes become smaller. Secondly, GTP works over IP and UDP, thereby it introduces in IPv4 networks an overhead of 36 bytes to every user packet (56 bytes in IPv6). We have done an analysis of the packet flows from a Finnish mobile operator network and the average packet size is around 250 bytes. Thus, the use of GTP introduces 14% or 22% overhead, in IPv4 or IPv6 respectively.

3. BASIC ETHERNET-BASED MOBILITY

Our work investigates possibilities to implement the mobility of the user equipments in L2 instead of L3 in 3GPP Release 8 network architecture [6]. L2 Ethernet connectivity is assumed in the mobile backhaul between the base stations (eNBs) and mobile core network elements. The changes we design to the specifications are mostly implemented in the user plane; control plane protocol stack, UEs and the air interface are not modified; this is an important distinction with regard to deployment of our scheme. Our design introduces significant changes to the 3GPP release 8 specifications: the GTP based bearer model is implemented with Ethernet. As a consequence the QoS implementation requires alternative solution as well.

A fundamental problem in the IP in mobility is that the IP address identifies the node and fixes the location in the network topology. This is why 3GPP specifies the GTP tunneling protocol, to avoid changing the IP address of the user device after each handover.

If routing and packet forwarding is done in L2, the main purpose of the GTP protocol becomes obsolete. Therefore in Ethernet mobility the user data packets are embedded directly into the Ethernet frames and the transport IP/UDP/GTP layers are not used; this saves tens of bytes per packet. In our scheme L2 MAC addresses are used as locators and IP addresses only identify the users. In Ethernet mobility a natural choice for the locator of the UE is the eNB MAC address. This is updated during the handovers. The identifier is the UE IP address allocated for each user by the PDN-GW during the attach procedure.

Because user specific GTP tunnels are not used, explicit path switch procedure during handovers is not needed and the signaling to the core network is reduced. Only the new MAC address of eNB, has to be propagated to the network.

The 3GPP release 8 network elements and L2 based mobile backhaul is presented in Figure 2. Here the carrier core network has a few Provider Edge (PE) switches that connect access network through Customer Edge (CE) switches.

Ethernet mobility utilizes VLANs to separate different traffic types and define the QoS treatment. In order to make the solution scalable VLAN stacking with 802.1ad frames is used (Figure 2). Outer VLAN (S-VLAN) is used to separate traffic destined to different eNB groups. If several operators share the network they can have S-VLANs of their own. The target is that handovers are rare between the groups. Inner VLANs (C-VLANs) are used to separate different traffic types and QoS classes. 3GPP signaling towards the MME element has a C-VLAN of its own.

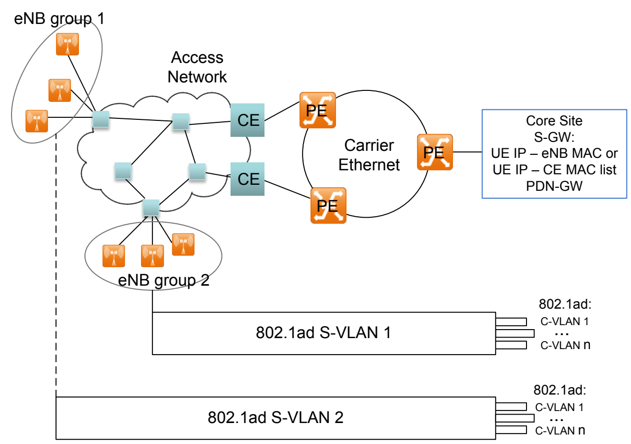

Fig. 2 Ethernet and VLANs in the mobile network

QoS classes have their own C-VLANs. Each tunnel (C-VLAN) carries data flows of all UEs located under the same eNB group. Data packets arriving to the eNB must be mapped to the right radio bearer, which is user specific. It is assumed that the UE IP address is unique and constant, and is used for the mapping. If this is not the case, the identifier may be any sufficiently unique identifier that can be tied to an individual UE. It is assumed that a certain user can have only one data bearer per QoS class and therefore C-VLAN separates the bearers of a certain user.

A further benefit of using Ethernet VLAN tags is that the carrier network can switch packets based on VLAN tags. This is very beneficial in terms of scaling the network segments and minimizing any state stored on the switches.

In our scheme, we use the MAC address of eNBs as the locator for the user, which means that each handover still forces some update to happen about the new location. Effectively this becomes a change in the IP-MAC bindings, and can be implemented simply as a gratuitous ARP after handover; we do not require routing change signaling through the carrier core, in particular since there is no GTP tunnel that needs to be re-initialized. However, the predictive handovers and X2 interface bring some challenges to the solution.

To avoid signaling every handover through the core network up to the MME, we employ mapping of addresses on the CE switches. The S-GW sends user data with the destination MAC address of the CE switch, configured during the attach procedure. The CE processes the packet, uses the IP address as the identifier and builds a new Ethernet frame with the destination MAC of the eNB serving the UE. Vice versa on the uplink, the eNB sends packets to its CE, that then sends the packet to the S-GW with its own MAC as the source.

The benefit of our scheme is that nodes in the core do not see MAC addresses of any eNB, which keeps the amount of state reasonable. As discussed later, all handovers under the same CE and access network remain unnoticed by the S-GW; it always sends packets to the same CE. Our design uses an extended TRILL-based technology in the access network, and the core is based on legacy Carrier Ethernet.

4. TRILL

The TRILL protocol [7] is an evolutionary approach to Ethernet bridging, engineered for the next generation Ethernet bridges, called routing bridges (RBridges). The protocol is a significant change from existing Spanning Tree Protocol (STP) [9] based approaches, solving forwarding table state explosion by use of a link-state protocol (IS-IS [11]). TRILL also solves the broadcast storm issues that STP based approaches may have during topology changes. TRILL routes Ethernet frames based on a link-state protocol, run between all RBridges in an Ethernet domain, called a *campus*; TRILL switches are identified by a unique *nickname*. The link-state protocol computes Shortest Path First (SPF) trees in the campus, allowing RBridges to optimally route traffic to every other RBridge leading to better link utilization throughout the domain. Many traditional Ethernet characteristics have been retained in the TRILL protocol, including the plug and play nature of the Ethernet, the presumption of a very static network switch topology as well as Ethernet frame flooding. TRILL is incrementally deployable, partitioning STP based domains into smaller ones as the size of the TRILL campus grows. TRILL frames can also travel STP based clouds during transit without issues. A further benefit is that TRILL includes Equal Cost Multipath Routing (ECMP) that allows efficient use of the whole network for data transport. The downside of TRILL is that it introduces new Ethernet headers, thus increasing the MAC layer overhead on the link. Yet, these headers do not lower the actual MAC layer data payload, i.e., the amount of space available for the IP packets.

While TRILL improves on the broadcasting nature of the current Ethernet switching technology and eliminates broadcast storms, more aggressive elimination of broadcasting could be accomplished at the cost of increased resource usage and processing overhead on a per RBridge basis. Additionally, our vision requires mobility support directly on the Ethernet layer, and thus in TRILL.

The plug and play nature and the location agnostic addressing of the Ethernet is ideally suited for a clean slate approach to UE mobility, further emphasized by the fact that the broadcast nature of the Ethernet as well as some of the current IP layer technologies may cause inefficient energy usage in devices with limited power.

In addition to standard TRILL routing, we use Consistent Hashing [8] (or synonymously, Distributed Hash Table, DHT) mechanism to extend TRILL with a link layer location service. The TRILL DHT location service is used to unicast Ethernet frames destined to unknown end-stations within the TRILL campus, instead of the conventional flooding mechanism. Additionally, the location service is also used to enable end-station mobility directly on the link layer. Furthermore, the DHT mechanism can be used to create a distributed ARP service for the whole TRILL campus, further reducing the amount of broadcast traffic flowing through the intermediate TRILL nodes [12]. Note that this does not eliminate all broadcast traffic; multicast IP traffic, some DHCP/BOOTP traffic and other multicast packets/frames will still be flooded using normal TRILL multi-destination frame delivery mechanisms.

5. TRILL-BASED MOBILITY MANAGEMENT IN ACCESS NETWORKS

As every UE is uniquely identified by an IP(v4 or v6) address in the operator network, our TRILL DHT location service can use the IP address information as a unique identifier for mapping the location information of the UE.

Our DHT design for TRILL uses a one-hop DHT instance for every customer VLAN identifier in the TRILL campus. Every RBridge may advertise to be a "bucket" (i.e. a device responsible for storing information, namely, <key, value> tuples) for a DHT instance as well as to indicate that the node is computing a DHT instance for a VLAN identifier. The DHT instance represents a view of the overall set of available buckets for a VLAN identifier on each RBridge, facilitated by the link state protocol. The buckets act both as a storage node for the information, as well as a specialized frame relaying device, when the information stored is location based. Typically, the key is a unique identifier (such as a UE IP address) that can be mapped to a single network entity, while the value can be whatever is being indexed in the DHT instance [12].

The TRILL DHT location service lookup process is done by an RBridge on every ingressed native Ethernet frame with a unicast destination address, for which the RBridge cannot find a corresponding entry from the local MAC address learning table. The unicast destination address is hashed and a DHT instance is queried for a bucket, for which the hash value belongs to. The bucket corresponds to a specific RBridge within the TRILL campus that has previously advertised its willingness to store location information for network entities. The frame is then encapsulated as a normal TRILL Data frame and sent towards the relaying RBridge.

When the CE device, running our DHT-based TRILL protocol, receives a packet from the S-GW, it notes the destination UE IP address, and looks up the right TRILL egress node, based on the UE IP. If there is no corresponding eNB MAC address for the UE IP, the CE must use the DHT location service to relay the frame towards the eNB. The CE also stores the <IP, eNB MAC> mapping for uplink packets, i.e., for an uplink packet to IP address X, switch to MAC address from the incoming packet (e.g. the S-GW's address). This operation is presented in Figure 3.

For uplink, eNB builds an Ethernet frame towards the previously stored source (e.g. S-GW) and transmits using TRILL encapsulation. The CE receives the packet, removes the TRILL encapsulation and transmits to the core using as destination address the S-GW.

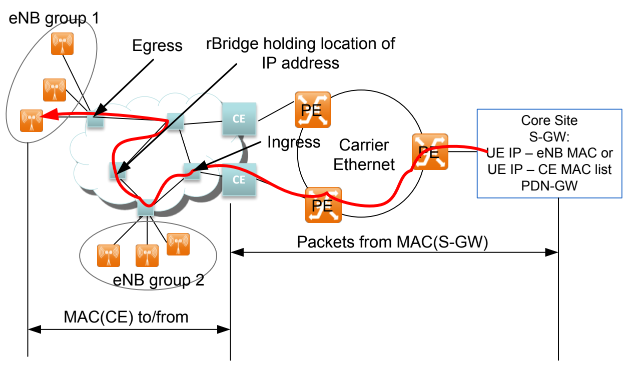

Fig. 3 Packet transport in the Ethernet-based network

There are several advantages to using an IP based location service in the access network: ARP can be directly utilized as a location update trigger from eNBs; as an attachment process completes, a gratuitous ARP message for the attached UE will begin a location update in the TRILL DHT location service.

Furthermore, the use of UE IP address restricts the MAC learning table state explosion effect everywhere in the access network; at worst, each edge RBridge (including the CE nodes) will see eNB amount of MAC addresses, instead of an UE amount. Thirdly, the CE nodes can leverage the uniqueness property of IP addresses to create an efficient address translation scheme, where an IP address is bound to an eNB MAC address. This address translation scheme allows the operator core to only see CE MAC addresses for UE originated traffic, further reducing state information that has to be stored on the mobile core network elements.

Limiting the location awareness inside the access network requires moderate changes: we introduce our TRILL extension into the switches that are directly connect to the eNBs. In addition, the eNB nodes require changes, e.g., they need to emit gratuitous ARP frames, change the UE context information, and map packets to the bearer. TRILL nodes do not need to do frame buffering during handovers, as the functionality is already present in eNB nodes.

The technical overview on how the IP-based TRILL DHT location service works is given below. The TRILL DHT location service state is updated whenever a UE attaches itself to an eNB, facilitated by the eNB emitting a gratuitous ARP message. The first-hop RBridge on the path will extract addressing- and location information from the gratuitous ARP message, and lookup the RBridge where the location information is stored, based on the UE IP address. The updated location information is sent to the relaying RBridge, which performs the following operations:

- Insert or update the location information in the local copy of the DHT information database.
- If the RBridge nickname of the location information changed (e.g. the end-station moved to an eNB that is behind another RBridge), signal the new location information (both the new RBridge nickname as well as the new eNB MAC address) back to the old RBridge.

Additionally, if the end-station changed eNB so that the new eNB after the handover process is behind another RBridge, the first-hop RBridge for the original eNB must also do the following:

- If the RBridge receives a TRILL Data frame destined to itself, containing the destination IP address of the newly moved UE, it must signal to the ingress RBridge of the TRILL Data frame that the UE has moved behind a new eNB (e.g. propagate the new location information to the ingress RBridge)

The end-result of the state reparation process is that each RBridge containing stale information about the location of the moved UE will be updated with the new location information as long as they actively try to send Ethernet frames to the UE.

Furthermore, the gratuitous ARP message can be broadcast to the access network using a tightly segmented VLAN identifier, which causes the CE nodes to update their <UE IP, eNB MAC> bindings based on the information contained in the message. The broadcasting must be tightly controlled to avoid MAC table state explosion on RBridges where the information contained in gratuitous ARP messages would be irrelevant. As we are presuming that eNB can buffer payload packets during handover process, the <UE IP, eNB MAC> bindings in CE nodes need not to be updated; stale information in a CE node will cause the frames to be forwarded to the old eNB, which buffers- and upon handover succeeding, forwards the buffered frames to the new eNB. As CE is also a TRILL DHT enabled node, the handover process will trigger a state repair, updating the location information of the moved UE in CE.

Figure 4 presents a high level sequence diagram for the state repair process in the access network when a UE performs a handover between eNB A and eNB B with adjacent RBridges A and B. The control plane signaling messages are depicted by dashed arrows, while user plane (data payloads) uses continuous and emphasized arrows. The example presumes that the eNB nodes are directly connected to two separate RBridges in the same access network. The state repair process is started when a UE begins the handover process to eNB B, facilitated by S1 signaling initiated by eNB A. The 3GPP messages are not interpreted by the TRILL protocol, they are simply IP packets transferred between eNBs.

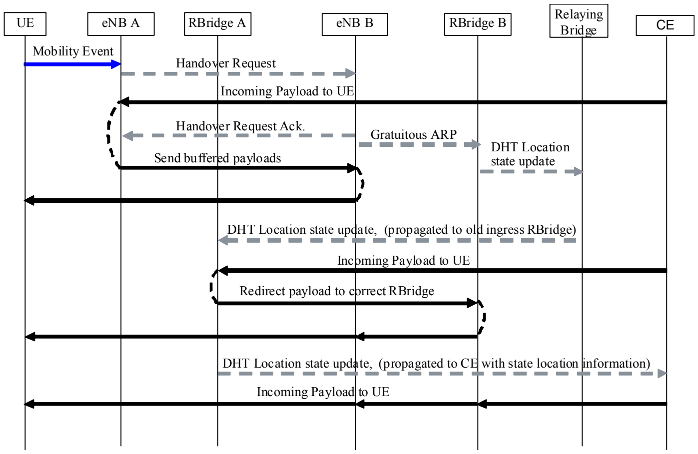

Fig. 4 Location state repair after mobility event

6. MOBILITY MANAGEMENT IN THE MOBILE BACKHAUL

The intra-access network mobility concept sets some requirements for the mobile core network. The mobile core network elements must know for each eNB the access network it belongs to and the CE devices through which the eNB is reachable. TRILL nodes inside the access network in turn must be aware of UE IP address – eNB mappings as well as the location of the eNB. The mobile core and access networks must exchange information explicitly as follows:

- The S-GW must learn the CE MAC addresses through which the traffic to the UE is delivered. An attach request message contains the cell global ID which uniquely identifies the eNB. The CE devices through which the eNB is reachable can be configured to the S-GW. A more dynamic way is to simply store the source MAC address of the attach message on the S-GW; this message comes from the CE with its MAC address and thus implicitly indicates the access network where the UE is located.

- After an attach procedure the eNB behind which the UE is located must inform the access network about the new UE IP address and its location by sending a gratuitous ARP message. To make it more reliable the message should be repeated. The gratuitous ARP message will trigger a location update in the access network using our TRILL signaling messages. Similarly the new location is informed to the access network after each handover and tracking area update. All handovers within a TRILL access network remain unnoticed by the mobile core network nodes..

- If the access network changes, the CEs through which the UE is reachable will change. During tracking area update (TAU) the S-GW can learn this information from the TAU request message; during handovers the target eNB sends a location update to the S-GW. This frame has a new source CE MAC address and the S-GW can update the location of the UE.

If the UE performs a handover across access networks, the signaling needs to go through the S-GW and MME. After the handover is complete, subsequent handovers within the new access network remain invisible to the nodes in the mobile core network.

7. SIMULATIONS

Localized handovers in Ethernet mobility reduce the signaling load in the S1 interface compared to the standard 3GPP Release 8 networks. The share of handovers in total signaling load increases when the cell sizes become smaller and thus the benefits of Ethernet mobility become notable.

This can be illustrated with a simplified mobility management model where the UEs are randomly distributed into the simulation area and the movement of the UE is realized by randomly selecting one of its neighboring cells. Time to the next movement is also randomized. A handover occurs during a cell change if the UE is active. A tracking area update occurs if an idle mode UE changes the tracking area. Each UE has a mobility management state (attached/detached) and a connection management state (active/idle). Mobility management timer tells the time to be attached and connection management timer the time to be active.

The signaling procedures modeled are attach, detach, service request (idle->active), S1 release (active->idle), handover and tracking area update. The signaling load is measured by calculating the amount of 3GPP signaling messages sent over the S1 interface during the simulation time. In standard 3GPP Release 8 networks attach and tracking area updates have three messages over the S1 interface, other signaling procedures are assumed to have two. In Ethernet mobility GTP tunnels do not exist and therefore attach and TAU procedures have only two messages over the S1 interface. Handovers are localized inside the access network.

Time-based simulation was implemented with Matlab. During one hour simulation time the state of each UE was checked every second. The signaling procedures, state and cell changes were collected. The results are presented in Figure 5. It can be seen that the benefits of Ethernet mobility are meaningful when the UEs change the cells rapidly (left side figure) which means that the handover rate seen by each cell (right side figure) is high. The simulation covered only signaling seen by mobile core network elements.

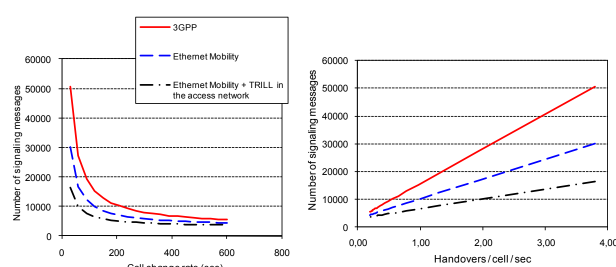

Fig. 5 Simulations of the Ethernet-based solution

Furthermore, as noted earlier, our Ethernet-based packet transport removes the GTP-tunneling, saving tens of bytes for each packet. With our data showing an average IP packet size of 250 bytes, this means a saving of 14% or 22% with IPv4 or IPv6, respectively.

In 3G networks the radio network controller (RNC) hides handovers from the core network elements. With LTE, all the handovers are visible to the gateways, which increases the signaling load towards these elements. The solution presented in this paper can hide handovers from core network in a similar way as the RNC does in 3G networks.

If a CE element has a 100Gbps link towards core network elements, one access network domain can have maximum 1000 eNBs assuming that the capacity of each eNB is 100Mbps. If the maximum number of subscribers under one base station is 2000 from which 5-10% are active, the amount of users making regular handovers in a domain is 100 000 – 200 000. It is estimated that voice users make approximately one handover in every 10 seconds. If all users behave like voice users, one access network domain can hide 600 000 – 1 200 000 handovers per minute. However, in practice this figure is smaller due to the fact that data users do not move the same way as voice users do.

8. CONCLUSIONS

Ethernet mobility requires changes to the 3GPP Release 8 network architecture. User data packets are routed in L2 and therefore the GTP-U protocol is not used. Our design enables localized handovers because our TRILL extension can be used to hide the handovers from the mobile core network. As a consequence the signaling load to the mobile core network is decreased. This is an advantage because a flat network architecture combined with small cell sizes is doing the opposite. However, the solution does not affect the signaling load caused by activating/deactivating bearers. In the user plane the protocol overhead is decreased because of missing IP/UDP/GTP layers.

Further benefits of our approach are that since eNBs are not visible to the core network switches outside the access network, the core can be used by a number of operators, enabling a shared backbone infrastructure to lower the operational costs. TRILL in the access network further reduces the state requirements and enables efficient multipath switching within the access network, including multiple CEs connected to the core network while requiring minimal changes to LTE elements, such as the eNB. Moreover, the access network enables efficient UE to UE data communication due to location awareness; however, the 3GPP specification does not support this at the moment.

The downside of our architecture is that it is a clean-slate approach, meant to introduce a more efficient and cost-effective EPC. Secondly, some signaling complexity is now offloaded towards the access network, causing additional overhead in access network control protocol processing and complexity. Thus, our future work will target deployment scenarios and a deeper analysis of the implications of an Ethernet-based solution to an LTE mobile network.

9. ACKNOWLEDGMENTS

This work has been performed in the framework of the CELTIC project CP7-011 MEVICO. The authors would like to acknowledge the contributions of their colleagues.

10. REFERENCES

[1] CiscoVisual Networking Index: Global Mobile Data Traffic Forecast Update, 2010-2015. http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html

[2] Metro Ethernet Forum. http://metroethernetforum.org.

[3] Nokia. Considerations for system architecture evolution. Technical Report SRJ0020/S2-051260, 3GPPTSG-SA WG2 Meeting 46, May 2005.

[4] General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access (Release 8); 3rd Generation Partnership Project, TS 23.401 v8.3.0, September 2008.

[5] Architecture enhancements for non-3GPP accesses (Release 8); Technical Specification Group Services and System Aspects; 3rd Generation Partnership Project, TS 23.402 v8.3.0, September 2008.

[6] Heinonen, J. "Ethernet Mobility", Licentiate Thesis, Helsinki University of Technology. 2009.

[7] Perlman, R. et al, Rbridges: Base Protocol Specification, The Internet Engineering Task Force, Internet-Draft (work in progress).

[8] Krager, D. et al, Consistent hashing and random trees: distributed caching protocols for relieving hot spots on the World Wide Web, STOC'97: Proceedings of the twenty-ninth annual ACM symposium on Theory of computing, 1997.

[9] Jeffree, T. et al, IEEE Standards for Local and metropolitan area networks – Media Access Control (MAC) Bridges, IEEE Computer Society, Std 802.1D-2004.

[10] Jeffree, T. et al, IEEE Standards for Local and metropolitan area networks – Virtual Bridged Local Area Networks, IEEE Computer Society, Std 802.1Q-2005.

[11] Intermediate System to Intermediate System intra-domain routeing information exchange protocol, ISO/IEC 10589, Second Edition, 2002.

[12] Varis, N. and Manner, J., Minimizing ARP Broadcasting in TRILL, GLOBECOM '09 Workshops, IEEE Computer Society, December 2009.