Diagnosing Channel Issues Using GTP Protocol Messages in LTE Core Networks

Saman Feghhi*, John Keeney†, Liam Fallon†, Douglas J. Leith*
*School of Computer Science & Statistics, Trinity College Dublin, Ireland.
{feghhis, Doug.Leith}@tcd.ie
†Network Management Lab, Ericsson, Athlone, Ireland.
{John.Keeney, Liam.Fallon}@ericsson.com

***Abstract*—For some time now the volume of data traffic in Mobile Telecommunications Networks has far outweighed voice traffic and most users are more concerned about the quality of their data connection than their voice calls. Adequate traffic service requires analysis of traffic flows to identify network issues affecting users' services. This can be achieved by placing network analytics probes throughout the network, but it would be better to have a small number of probes at just key locations in the Mobile Core Network. We show how to use core network traffic analysis to identify which user owns and which base station hosts each traffic flow, particularly where users and their connected traffic flows may dynamically switch between base stations or radio technologies.**

I. Introduction

Modern telecommunications networks need constant monitoring and management to operate effectively. Traditional Operating Support Systems (OSS) constantly gather information from network nodes to analyse the operation of the network and the services that use the network. In almost all current Mobile Telecommunications Networks (*e.g.* 3G and LTE), data-traffic volumes are much higher than voice traffic volumes. In addition to collecting performance events and counters from network nodes there is potential to examine the characteristics of the data traffic to diagnose issues in the network.

The challenge when dealing with current mobile networks is the enormous volume of user data and control messages that are being transmitted throughout the network. Given the scale of messages which must be processed in real-time for effective network analysis, collecting data from all nodes and interfaces in the network is clearly not plausible.

The solution is to find a minimal set of key interfaces from which to gather measurements that characterise the live status of the network and the services being delivered to subscribers. To this end we describe a method to collect information about the traffic of users to analyse the performance of the mobile network in handling that traffic. This data is gathered by inserting the minimum number of network traffic probes at key positions in the Core Network (CN), the part of the network that operates between the network operator's Radio Access Network (RAN) and the public Internet.

In Section II we review the definition of GPRS Tunnelling Protocol (GTP), used to control and carry user-plane traffic in mobile core networks. We also provide an overview of Core Network nodes, links and topology. Section III explains our approach for network measurement collection, and the challenges we encountered. In section IV we assess our approach on a real medium-sized 3G/LTE mobile network. Section V discusses current and future applications of our method.

This work was partly supported by Science Foundation Ireland under Grant No. 11/PI/1177.

978-1-5090-0223-8/16/$31.00 © 2016 IEEE

II. Mobile Core Network Topology

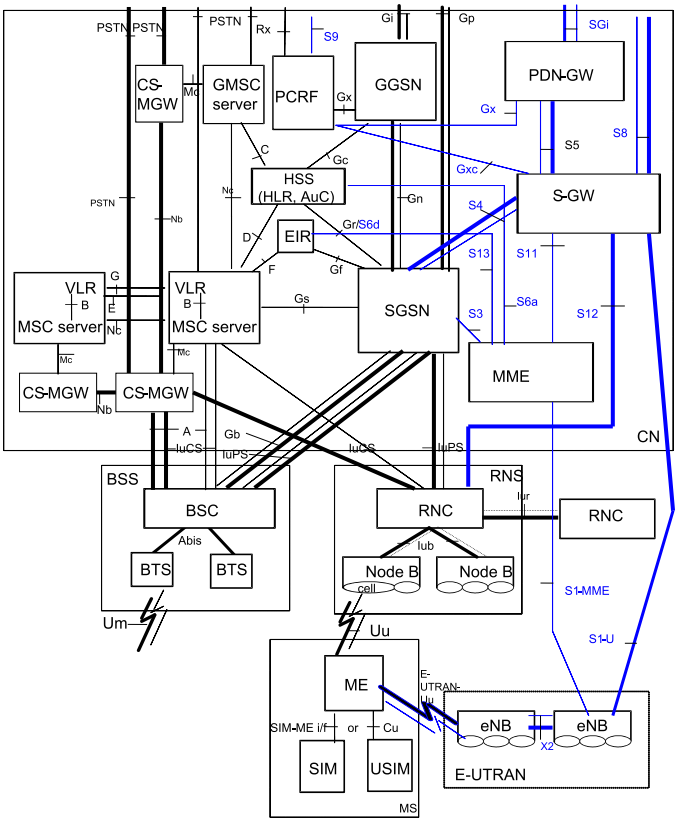

Fig. 1: Illustration of 2G/3G/LTE Mobile Core Network [1]

This section provides an overview of GTP and mobile core networks for both 3G and LTE networks. Fig. 1 shows a general Mobile Core Network and the important nodes and links. Mobile User Equipment (UEs or MSs), *e.g.* phones or modems, communicate over radio with base stations (3G NodeBs and LTE eNodeBs). The base stations then connect to the mobile core network that delivers their traffic to internet gateways (LTE: Packet Gateway (PGW or PDN-GW), 3G:

Gateway GPRS Support Node (GGSN)). While the physical path between the UE and the internet gateways may vary depending on UE movement *etc*, all data traffic is tunnelled over various nodes and links between the gateway and the UE using the GPRS Tunnelling Protocol (GTP) to create a stable virtual link between the two nodes. These tunnels are often referred to as "bearers" in LTE or "contexts" in 3G.

A. 3G Core Network

3G networks consist of two major parts: the traditional circuit-switched network that handles voice calls, and the packet-switched network which handles data communications. Traffic coming from a UE first arrives over radio to whatever base station (NodeB) the UE is currently connected to. Each NodeB can serve one or more radio cells or sectors providing radio coverage in its surrounding area. Pools of NodeBs are then controlled by a Radio Network Controller (RNC) which coordinates the NodeBs and delivers the traffic to an SGSN (Serving GPRS Support Node). The SGSN then routes packets based on charging and policy information received from the PCRF (Policy and Charging Rules Function) through an appropriate GGSN internet gateway.

B. LTE Core Network

In an LTE core network (Evolved Packet Core) RNC functions are performed in the eNodeB base station. LTE supports only packet data so data packets are sent from the eNodeB to the Serving Gateway (S-GW). The MME (Mobility Management Entity) node is responsible for coordinating UE mobility and also for establishing and maintaining all traffic bearers from UEs to an internet gateway (PDN-GW). Most LTE core network nodes maintain interconnections with 3G core network nodes to support backwards compatibility and 2G/3G/LTE handovers, especially since traditional circuit switched voice calls are not natively supported in LTE and fall back to 2G/3G.

C. GTP Protocol

GTP (GPRS Tunnelling Protocols) are a group of protocols defined for QoS-specific traffic tunnelling between base stations and core network nodes (Fig. 2), seamlessly maintaining data and connections as UEs move between base stations. This approach maintains the same IP address for the UE regardless of its location, its mobility or the physical links/path between UEs and the internet gateway. We discuss the two main protocols: GTP-U for user traffic between UEs and the gateway, and GTP-C for control traffic to coordinate bearers and GTP tunnels (Fig. 3) between different core nodes, thus abstracting the physical nodes required to route traffic. The QoS characteristics of bearers are defined by the radio-specific settings applied to those bearers, with different QoS templates available for selection when a bearer is set up or modified.

All control and data packets travelling through GTP tunnels are routed using TEIDs (Tunnel Endpoint IDs), unique identifiers within each mobile network. TEIDs are generated and maintained using GTP-C messages. Fig. 2 shows the GTP protocol stacks for GTP-C and GTP-U packets. It is important to note that all user-plane traffic is encapsulated inside GTP-U packets. There are a variety of GTP-C control messages to control GTP tunnels, which usually come in Request/Response pairs. These GTP-C messages include:

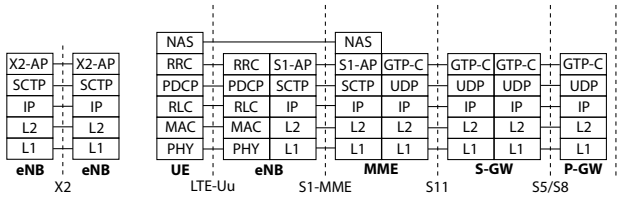

(a) GTP-C Protocol Stacks

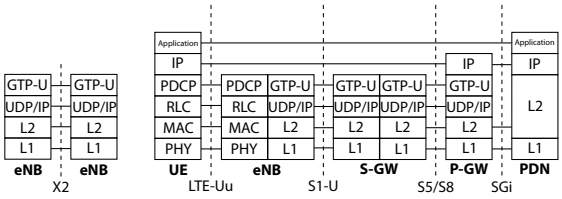

(b) GTP-U Protocol Stacks

Fig. 2: LTE protocol stacks for control and data packets. The stacks show the data transmission on Uu, X2, S1, S11, S5/S8 and SGi. The protocol stacks for 3G are similar.

LTE Create/Delete Session Request/Response: Session control messages are sent during attach/detach of the UE to/from the mobile network. They may also be sent when a UE switches between base stations that are served by different S-GWs. A Create Session Request sends the UE's IMSI (a unique identifier for each UE), which can be later used to match TEIDs to individual UEs. Once a session is closed with a Delete Session pair, the TEIDs are released to use for a different tunnel (bearer).

LTE Create/Delete Bearer Request/Response: During each attach, at least one best effort EPS bearer (default bearer) is created for each user. EPS bearers are created with Create Bearer message pairs and closed with Delete Bearer pairs. Additional EPS bearers (best-effort or dedicated) may also be established as the UE demands, each extra bearer being identified by a 4-bit EBI (EPS Bearer ID). Every bearer on a UE is assigned a separate IP address, so a UE will have an IP address for every EPS bearer it uses. These UE IP addresses are used by the PDN-GW to route the user's traffic in and out of the core network, with all traffic to/from each UE IP address tunnelled using GTP-U based on TEIDs.

3G Create/Delete PDP Context Request/Response: Data sessions in 3G are very similar to those in LTE, except tunnels are called 'PDP Contexts' and PDP Context ids are called NSAPIs. During 3G-LTE handovers EBIs are automatically translated to/from NSAPIs.

LTE Modify Bearer Request/Response: Whenever some aspect of a bearer is modified, a pair of Modify Bearer messages are sent, for example as part of a handover between two cells in LTE which may or may not be controlled by the same MME, or as part of an LTE-3G handover, or when the QoS requirements for a bearer need to be changed.

3G Update PDP Context Request/Response: Similar 3G messages to Modify Bearer messages in LTE.

III. Proposed Method

The first and most important step to identify which core network traffic flows go through each base station is to first

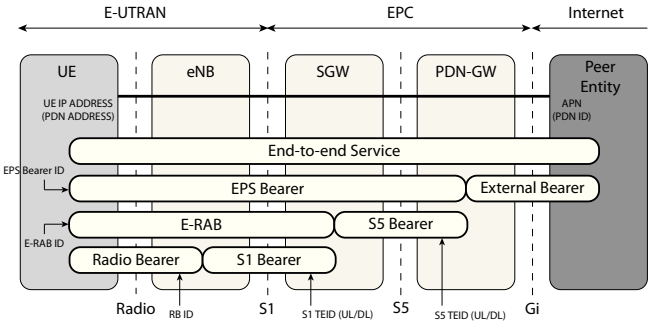

Fig. 3: LTE Network Bearers

identify the UEs that use the flows and therefore which user is using each GTP bearer. When a UE attaches to the network, a default bearer is established that routes traffic for that UE through the network. Additional bearers with different QoS parameters may also be requested by the UE, where each new bearer will be identified by a different EBI/NSAPI. For each bearer, a different IP address is assigned to the UE, so a single UE may have multiple IP addresses. These IP addresses persist through handovers and are valid for the lifetime of the UE's attachment to the network. These IP addresses can be used to identify the traffic generated by each UE. Further, it is possible to use GTP messages to identify which base station is serving the UE at any point of time. It is important to note that as the UE moves throughout the network these TEID to IP address mappings must be updated.

Irrespective of whether the network is 3G or LTE the initial procedure is as follows: A new bearer is created using 'Create PDP Context/Bearer Request/Response' containing the IMSI, identifying the subscriber associated with a bearer and therefore the UE. The TEIDs and IP addresses of intermediate nodes are also exchanged during this procedure. From the outer IP packet carrying the GTP messages the 'source' IP addresses identify the eNodeB (LTE) or the RNC (3G), while the 'destination' IP address identifies PDN-GW (LTE) or GGSN (3G) which handles the traffic from UE. These IP addresses associate a base station or gateway with given TEIDs.

A. 3G/GPRS Network

Activation Procedure: In 3G, the only GTP message containing a unique identifier for a UE is 'Create PDP Context', where the IMSI of the UE is sent from SGSN to GGSN on the Gn interface as part of the attach procedure. The message contains the source TEID and IP address, with which the GGSN may respond to the traffic coming from the SGSN. Until the SGSN receives the response, it can accept packets from the GGSN with the current TEID and IP address, but it can only start sending these packets when the uplink TEID and IP is received in a matching 'Create PDP Context Response'. After that the default bearer is established and traffic flows in both directions.

Intra-RNC handover: During handover, if the source and destination channels/cells/sectors are hosted by the same NodeB or different NodeBs handled by the same RNC, the handover is handled in the NodeB/RNC without changing any tunnel characteristics *i.e.* TEIDs and IP addresses remain unchanged.

Inter-RNC handover: When a handover requires changing to a different RNC, if the source and target RNCs are connected to the same GGSN an 'Update PDP Context Request/Response' pair is sent on luPS with message elements containing the source and target TEIDs and GGSN IP addresses. This information can be used further to associate GTP-C messages and PDUs to/from the UE on the target bearer. If the source and target GGSNs are different then the sessions needs to be completely closed using 'Delete Session Request/Response' on the luPS link. Depending on the type of handover (soft or hard) a 'Create Session Request/Response' is sent either before or after the old session is closed. These 'Create Session' messages are sent using the new TEIDs and IP addresses and also contains the UE's IMSI so can be used to update its flows' network addresses.

3G-LTE (Inter-RAT) handover: A UE may switch Radio Access Technology (RAT) from 3G to LTE if an LTE cell is in range or coverage improves. Whenever a UE initiates such a handover, a 'Create Session Request/Response' pair is sent on the S11 interface between the MME and target S-GW using the same TEID and PDN IP addresses as the old session in 3G, allowing the UE to be associated with the new session. The UE can also be recognised if the messages contains the IMSI and assigned IP address of the UE.

B. LTE Network

Activation Procedure: A 'Create Bearer Request/Response' is sent from the PDN-GW to the S-GW on S5/S8 interface and from the S-GW to the MME on the S11 interface as part of dedicated bearer activation procedure. Both messages contain the TEIDs and IP addresses for the S1-U and S5/S8 tunnels. Further, similar to 3G a 'Create Session Request/Response' is also sent as part of the activation procedure which contains the TEIDs and IP-addresses along with the UE's IMSI.

Intra-eNodeB handover: LTE-LTE handover between channels/cells hosted by the same eNodeB is handled by the eNodeB without any GTP messages.

Inter-eNodeB handover (X2 handover): If the eNodeBs are neighbours and are associated with the same MME then neighbouring eNodeBs can handle the handover natively using the X2 interface, and a 'Modify Bearer Request/Response' is sent over the S11 and S5/S8 which contains the TEIDs and IP addresses of the new bearer. Similar to 3G a 'Create Session Request/Response' message pair may also be sent on the same interfaces using the old TEIDs, which will contain the new TEIDs and IP addresses, and may also contain UE's IMSI that helps identifying the UE. The old session then expires either before or after the new session is created depending on whether the handover is a soft or hard handover.

Inter-eNodeB handover (S1 handover): Where an X2 connection does not exist between to two eNodeBs, or the eNodeBs are associated with different MMEs, then an S1-based handover is required. S1 handovers are treated them same as X2 handovers and as before the TEIDs and IP addresses of the old bearer and the new bearer can be used to match the flows to the UE.

LTE-3G (Inter-RAT) handover The UE initiates an inter-RAT handover when the LTE coverage is no longer satisfactory due to either lack of equipment or congestion. As part of this handover 'Modify Bearer Request/Response' messages are sent on the S11 interface between the MME and the S-GW and on the S5/S8 interface between the S-GW and the PDN-GW.

An 'Update PDP Context Request/Response' message pair can also be seen on the luPS interface, which uses the same TEIDs and IP addresses from the original LTE session.

C. Tracing Network Flows to Identify Congestion

Since it is now possible to track flows from all UEs, the challenge is to determine which base station is serving a UE at any point of time, and so bundle the UE traffic loads according to the base stations currently hosting those UEs. For LTE, this information can be obtained using the UDP PDU packets that carry the GTP messages. The source IP address of the UDP/IP header that represents the eNodeB that sends or receives the UDP packet is used so that the throughput of each eNodeB can be estimated without knowing the origin of the flows.

The situation is more complex in 3G. The source IP address no longer belongs to the NodeB but to the RNC which controls a set of NodeBs. Communication between an RNC and its child NodeBs are not seen in the core network. With this approach it is possible to estimate the load of each RNC but it is not possible to further partition this load for the NodeBs controlled by that RNC. Therefore, with 3G, it is not possible to rely only on GTP messages on the core network interfaces. However in any RAN OSS system there will be other readily available data sources to identify which base station hosts which UE such as monitoring events from SGSN/MME, cells and UEs, which, when correlated with the GTP data makes it possible to localise individual flows to particular NodeBs.

While data traffic is most likely to cause congestion in a base station it is not the only traffic through the base station. Other traffic includes voice call traffic, SMS traffic, and management traffic. This traffic is not transported using GTP. Therefore a challenge remains to trace and track this traffic, which currently remains a considerable but relatively deterministic traffic volume. However with relative growth in data traffic and the growth of VoLTE (Voice over LTE) and internet services as an alternative to traditional circuit switch voice communications, this proportion of the traffic that is not transported over GTP is rapidly decreasing.

IV. Analysis of a Real Mobile Network

We used the method described in §III to analyse the GTP traffic from a real mobile network. We used a number of deep-packet inspection (DPI) probes to analyse GTP-C traffic on key interfaces in a segment of a medium-sized 3G/4G mobile network. Using the information from GTP headers, we were able to analyse the behaviour of user traffic in the network.

A. Experimental Setup

We used probes to collect GTP control messages over the S1 and X2 and equivalent 3G interfaces. For the purposes of this paper, the analysis was scoped to one hour of traffic collected between 20:00 and 21:00 local time, peak hour for data traffic. The traffic was collected on both 3G and LTE interfaces and was then classified by criteria such as UE, cell, and (e)NodeB using our method. We then interpreted this information to analyse the network behaviour. In this section we present some statistics derived from that hour of inspected data.

	3G	LTE
IMSIs	189638	390583
IP addresses	234787	479233
UTRAN/EUTRAN Cells	3830	1
Service Areas	18139	1
Routing Areas	23	1
Tracking Areas	1	1
Location Areas	1	1

TABLE I: 3G & 4G nodes seen in Network Segment in 1 hour

Network Type	Create	Update	Delete
3G	511589	53209	91047
LTE	769329	19519365	747464

TABLE II: GTP-C messages seen in Network Segment in 1 hour

B. Statistics and Results

Using information from *create session* messages, we evaluated the number of active users on the network. Further we used *update session* messages to track users in the network and in turn count the active cells as well as Service, Routing, Tracking, and Location Areas. Table I summarises these statistics.

The number of *create*, *update* and *delete* messages for 3G and LTE network are summarised in table II. We noticed that the number of LTE session messages are relatively higher than those of 3G, which is due to a much higher usage of LTE than 3G for data traffic with this operator. This can be confirmed by looking at the number of IMSIs and IP addresses that appeared in LTE GTP messages which are almost twice those in 3G.

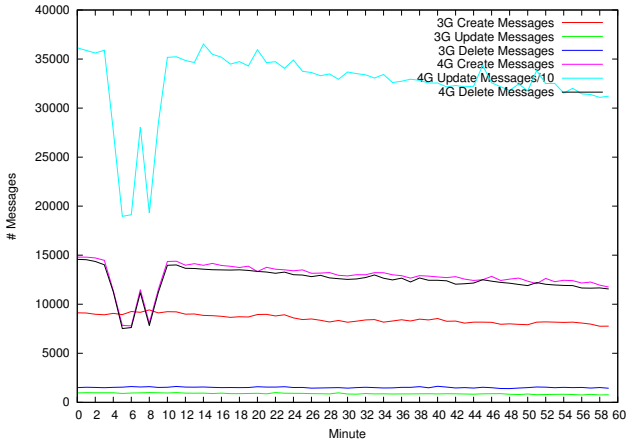

Fig. 4: Message volumes seen in Network Segment in 1 hour

Fig. 4 illustrates the volume of GTP-C session messages observed over the one hour period examined.[1] The number of LTE update messages is a factor of 10 greater than 3G update messages because, in LTE, the eNodeB nodes and the SGW cooperate to perform many of the functions carried out between the RNC and NodeBs in 3G.

We also observed the number of active 3G/LTE bearers for each bearer type. In 3G, a Network (Layer) Service Access Point Identifier (NSAPI) identifies a PDP content. In LTE, an EPS Bearer Identifier (EBI) identifies a bearer. The NSAPI and EBI both have 4 bits, allowing a maximum of 16 PDP

[1]The unusual drop in LTE message volume at Minute 5 was due to an unidentified issue in the LTE network resolved at Minute 10.

NSAPI/EBI Value	# of NSAPIs (3G)	# of EBIs (LTE)
5	505398	722431
6	2215	35609
7	527	476
8	454	3
9	438	-
10	421	-
11	413	-
12	376	-
13	406	-
14	375	-
15	566	-

TABLE III: Bearer Types for Created Sessions in 1 hour

contexts/bearers to be established for each UE. The default bearer for each UE has a NSAPI/EBI value of 5.

These identifiers should not be confused with QoS Class Identifier (QCI), an integer from 1 to 9 indicating different QoS performance characteristics for each IP packet. In particular, packets from bearers of different EBI/NSAPIs may have similar QCI values. As expected, most bearers in both 3G and LTE networks have the default value of 5, assigned to UEs when they attach to the network. Other values are used when secondary bearers with different QoS requirements are initiated. 99% of sessions in 3G have the default NSAPI value but in LTE, which has support for dedicated bearers for services like video that require different QoS characteristics, the number of sessions with the default EBI drops to 95%.

3G SGSN IP addresses seen	860
3G handovers seen by SGSN	49003
LTE eNodeB IP addresses seen	189
LTE inter eNodeB handovers seen	11161425

TABLE IV: IP Addresses and Handovers Observed in 1 hour

Table IV summaries the number of gateway nodes and handovers observed using our approach. We used GTP message data to identify the host cells and NodeBs of UEs, and in turn detect handovers each time UEs switch between different nodes. The data represents only the number of nodes and handovers that appeared in the target network segment over the one hour observation period, so the actual number of nodes and handovers may be much higher than what we observed.

One can observe that the number of LTE handovers are much higher than 3G handovers. However, this does not imply that there are more LTE handovers in a mobile network than 3G handovers, but that mobility and LTE handovers is more visible in GTP messages than 3G, given that the majority of 3G handovers are intra-SGSN which would not produce control messages on the GTP control interfaces we monitored.

Observing how volumes of handovers change over time in a network can give a good picture of network performance and even network saturation over time. Figure 5 illustrates volumes of 3G and LTE handovers over the observation period. We noticed a rapid increase in the number of handovers in LTE as opposed to a more predictable descending behaviour for 3G. These differing trends may be explained by considering that the data is captured in the evening when there is less mobility and Inter-SGSN handovers are less likely to occur; we did not monitor the more localised intra-RNC handovers.

An analysis of handovers can produce valuable insights into network status. For example, a sudden handover volume

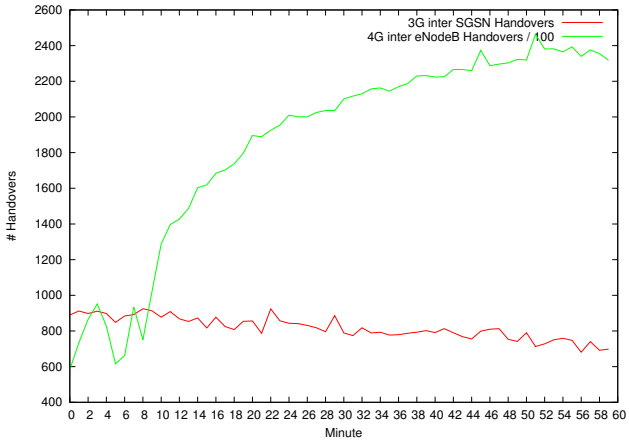

Fig. 5: Observed 3G & LTE Handovers in 1 hour

increase may suggest a network disruption. This can be further investigated by examining the volume of inter-RAT handovers and comparing the volume of 3G and LTE handovers.

While handover information gives useful insights, deductions depend heavily on the characteristics of the observation data *e.g.* which nodes and interfaces are monitored, which areas are involved, *etc*. As a simple example, if we focus on the data for a certain cell by filtering on the base stations/sectors involved, an increase in the volume of handovers compared to the other cells may suggest that a considerable number of UEs reside at the edge of that cell and are continuously ping-ponging between different base stations/sectors. Further decisions can then be made to change antenna tilt angles, antenna power or inter-cell handover thresholds for those cells to provide better coverage for these areas. Similarly, an increase in inter-RAT handovers in a certain cell or base station may suggest a local network issue *e.g.* weak reception, *etc*. which can be resolved by investigating the base stations involved or increasing the number of sectors for better reception.

V. Other Applications & Related Work

In addition to GTP-specific analysis, standard IP traffic analysis of both the GTP IP traffic and the encapsulated IP traffic can be very informative. Transmission delay and packet loss can be caused by many issues in mobile networks. Base station saturation is a common case, where UEs are forced to switch to nearby cells or downgrade to 3G or 2G until better quality connections are available. A less common case is where backhaul links are unreliable, for example where a microwave link to a base station suffers interference. While reductions in traffic quality exhibit similar symptoms at the UE, the cause of issues is difficult to determine without collecting extensive measurements from all the nodes in the affected region.

When a base station is operating close to its maximum capacity, delays and packet losses may be due to saturated radio channels. We can use this assumption to detect congested base stations. In order to estimate a base station's load it is first necessary to detect flows travelling in/out of it. An obvious näive approach is to monitor all links in the network, but it makes sense to minimise the number of network probes required. We have proposed a method to estimate the load on each base station using only GTP protocol information. By

analysing only GTP traffic it is possible to identify individual traffic flows between UEs and an internet gateway. When the traffic is associated with particular users and base stations, then standard traffic analysis can be used to identify transmission delay, packet loss, *etc*., and so infer if the base station is suffering congestion. Specifically we can achieve this by monitoring the GTP traffic only on the S1, S11 and S5/S8 connections in LTE and equivalent links in 3G, as shown in Fig. 1.

An important application of the proposed method is to detect corrupted links in the backhaul network. Using the aggregated flows for each node, one can estimate the current operating bandwidth of each base station. If the base station is operating below its capacity, given that inner core links are usually reliable (or unreliability is easy to detect), any delay or packet loss suggests an issue on the link between the base station and the core network. The operator can then narrow down the investigation to candidate links and fix them if required. In some networks many of these connections can be over microwave links sensitive to interference, or over links not directly controlled by the operator (*e.g.* on customer sites).

As suggested in [2], handover information can be used to estimate user mobility, UE population at each cell and even to deduce an approximate network topology at any given time. Further, base station saturation can be detected by comparing its operating bandwidth to its capacity. Mobile operators can use these insights to decide to perform targeted equipment upgrades in certain areas, add new sectors and/or base stations, adjust cell sizes, or adjust antenna tilt angles. This information can also guide load balancing on base stations, S-GWs or even P-GWs without a need for expensive surveys or drive tests [3].

Real-time congestion detection is the key to successful congestion control. While decentralized methods like OWAMP [8] requires additional functionality on nodes and injection of extra traffic on channels which reduces the overall throughput of the network. Using the method proposed here enables congestion detection in real-time by inspecting traffic traversing through network links without a need for extra equipment or traffic.

The accumulated bandwidth of each user can also be obtained from the summarized information of flows for each user. The network can use this information to calculate the bandwidth occupation percentage for each user which can be further used to establish resource allocation fairness.

A general model for finding minimal probing nodes is introduced in [4] but not for networks carrying GTP traffic. In [5] the emphasis is on GTP-C traffic analysis and the authors minimise the number of probes by placing them on the S1-U, S11 and Gn interfaces. While our focus is on correlating between user plane and control messages to provide more comprehensive measurements to analyse network performance, we take a similar approach to probe placement.

Most other approaches for Core Network analysis (*e.g.* [7]) require the collection of measurements at each base station. This either requires extensive probing or access to a monitoring interface or performance data at each base station, which in turn requires significant infrastructure to gather and process this data. The need for EPC probing is also discussed in [6], where probing points are divided into categories depending on the method that is used and the information they provide. Our described method focuses on the interfaces and methods that are categorised as DPI probing in [6].

The approach in [2], which in some ways is similar to that presented here, collects summarized information for each network entity (*e.g.* user, sector, base station site, group of sites *etc*.) by probing the inner core interfaces. However that approach suggests probing a long list of interfaces to capture almost every important control message traversing the network links. Our method aims to minimize the number of probe points and limit them to S1-MME, S1-U, S11, S5/S8 (and possibly S4) in LTE and equivalent links in 3G.

VI. Conclusion

In this work we have discussed how user traffic is transported over modern mobile telecommunications network between UEs and the core network's gateway to the public internet. Users' data connections are maintained as tunnelled GTP bearers/contexts between the gateways and the base stations (or RNCs) hosting those users as they move around the network and switch between radio technologies. We demonstrate how to gather and maintain information about these individual flows, and as a result support various analytics tasks to monitor and understand the behaviour of the core network. We have shown that to analyse the network, it is not always necessary to collect data from all nodes and interfaces in the network, and a minimum set of interfaces is sufficient to collect enough information to perform significant analysis tasks.

References

[1] 3GPP. 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Network architecure, 2014. TS-23002.

[2] S.K. Kovvali, M. Vutukuru, C.W. Boyle, R. Ghai, J. Hutchins, T. Abou-Assali, Y. Zhang, N. Rana, T.V.G. Araveti, R. Sirisikar, et al. RAN Analytics, Control And Tuning Via Multi-Protocol, Multi-Domain, And Multi-RAT Analysis, January 24 2013. US Patent App. 13/555,787.

[3] L.E. Li, Z.M. Mao, and J. Rexford. Toward Software-Defined Cellular Networks. In *Software Defined Networking (EWSDN), 2012 European Workshop on*, pages 7–12, Oct 2012.

[4] M. Natu and A. S. Sethi. Probe Station Placement for Robust Monitoring of Networks. *Journal of Network and Systems Management*, 2008.

[5] A. Sultan, F. Benbadis, V. Gauthier, and H. Afifi. Mobile Data Network Analysis Platform. In *Proceedings of the 6th International Workshop on Hot Topics in Planet-Scale Measurement*, HotPlanet '15, pages 13–18, New York, NY, USA, 2015. ACM.

[6] B. Wehbi, J. Sankala, and E.M. de Oca. Network Monitoring Challenges in the Evolved Packet Core. In *Future Network Mobile Summit (FutureNetw), 2012*, pages 1–8, July 2012.

[7] V. Yanover and Z. NUSS. Method and Apparatus for Load Management in Cellular Communication Networks, August 28 2014. US Patent App. 14/187,009.

[8] Y. Zaki, T. Weerawardane, S. Hauth, E. Wallmeier, and C. Gorg. Intelligent Traffic Enforcement for LTE Backhaul. In *Personal Indoor and Mobile Radio Communications (PIMRC), 2013 IEEE 24th International Symposium on*, pages 3077–3082, Sept 2013.