Caching in flat mobile networks: design and experimental analysis

Fabio Giust[□], Gerald Kunzmann[‡], Daniele Munaretto[†], Carlos J. Bernardos[□], Bessem Sayadi[§]
[□] University Carlos III of Madrid, Spain, E-mail: {fgiust,cjbc}@it.uc3m.es
[‡] DOCOMO Communications Laboratories Europe, Munich, Germany, E-mail: kunzmann@docomolab-euro.com
[†] Department of Information Engineering, University of Padova, Italy, E-mail: munaretto@dei.unipd.it
[§] Alcatel-Lucent Bell Labs, Paris, France, E-mail: bessem.sayadi@alcatel-lucent.com

Abstract—Mobile network operators are aiming at strategies to accommodate the expected huge traffic growth with limited impact on the core network load. Enhancing the mobile architecture, along with deploying content delivery networks, represent promising areas where operators are putting their efforts to reduce costs while meeting the service guarantees.

In this article we present a novel mobile architecture for 5G networks which benefits from the joint design of distributed mobility management and mobile video content delivery networks. We first discuss the architectural aspects of this combination and then present our validation and performance results from a real Linux-based prototype implementation.

I. Introduction

In the last years, we have witnessed the explosion of a "mobile revolution", mainly driven by the massive market penetration of powerful mobile handsets, and the deployment of faster and heterogeneous radio access technologies. Technical reports, such as [1], show that mobile traffic will keep increasing in future years, and video media content will be the major driver: it will account for over 66% of total mobile data traffic by the end of 2017.

However, current architectures for mobile and cellular networks are not suitable to accommodate the vast amount of data generated by mobile users. Indeed, they are structured in a centralized and hierarchical way, clearly differentiating the Radio Access Network (RAN) and the packet core, wherein users' data flows must traverse both. If we take the 3GPP's Evolved Packet System (EPS) as the reference architecture (see Fig. 1), traffic generated in the RAN is conveyed by means of tunneling to the Evolved Packet Core (EPC), formed by (among others) the Serving Gateway (S-GW) and the Packet Data Network Gateway (P-GW). The P-GW aggregates the traffic from several edge networks (i.e., from the S-GWs) and acts as gateway between the operator's internal network and external IP networks. From a mobility perspective, data traffic is permanently anchored at the P-GW, which is in charge of moving the data tunnel end-point to the corresponding S-GW following the terminal's movements.

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7-ICT-2009-5 under grant agreement n. 258053 (MEDIEVAL project) and FP7/2007-2013 under grant agreement n. 317941 (iJOIN project).

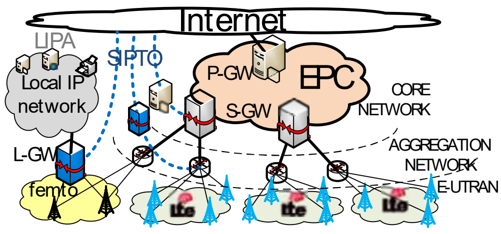

Fig. 1. Simplified 3GPP's Evolved Packet System architecture.

The 3GPP is already looking at solutions to unbind the IP connections from the P-GW, such as the Local IP Access and Selected IP Traffic Offload (LIPA-SIPTO) [2] and the LIPA Mobility and SIPTO at the Local Network (LIMONET) [3]. LIPA enables direct connectivity to local IP resources from residential/enterprise access through a Local Gateway (L-GW), without traversing the core. SIPTO allows the data plane to by-pass the core network, and it is available from macro cell access through a local breakout at or above RAN level, as well as from residential/enterprise access through the L-GW. LIMONET provides mobility support for LIPA and SIPTO within the residential/enterprise access network. In addition, the mobile architecture is further enhanced by employing tools for Network Function Virtualization (NFV), like those specified by the ETSI NFV group, that allow for instance to deploy virtual network entities, or even to virtualize the whole EPC (virtual EPC - vEPC).

Besides the efforts to design a flatter network, content caching is seen as an additional valuable resource to reduce traffic: [4] and [5] explore the benefits for a 3G and a 4G system respectively, concluding that up to two thirds of mobile traffic can be reduced. The study in [6] surveys the current strategies for mobile Content Delivery Networks (mCDN), showcasing which techniques are most suitable with respect to the caches' location. According to the survey, a mature object-oriented caching technique, like URL-based web caching, finds a natural application with caches co-located or placed by the P-GW, being it the point where packets are no longer encapsulated, and a plain HTTP proxy server would accomplish the task. In this context, authors of [7] propose to place caches at the P-GW for a centralized deployment,

978-1-4799-8088-8/15/$31.00 ©2015 IEEE

or at the L-GW for a distributed one, and then they explore an application level content delivery optimization based on a smart scheduler. However, caching at the P-GW still implies data traversing the P-GW/S-GW links, hence to cut down traffic in the core caches should be placed at a lower level, e.g., at the S-GW or in the RAN nodes. Unfortunately, this aspect requires more sophisticated caching techniques that perform a deeper monitoring of the data flows. These methods are usually referred as Redundancy Elimination (RE) and the objective is to remove arbitrary duplicated byte patterns from the traffic flowing within a portion of the network, in this case the EPC. Interested readers might find more details in [8], where a flow monitoring tool has been deployed in a real operator's EPC, along with a RE technique for TCP flows. Enabling object-oriented caching at the RAN level would definitely be a simple and valuable resource for operators.

The IETF is exploring Distributed Mobility Management (DMM) as a brand new paradigm for flat networks in 5G mobile architectures. DMM targets pushing the mobility anchors to the edge, flattening the network, and letting the IP connections to be bound to access routers rather than to a core entity. In this way IP flows can be easily offloaded from the core, thus limiting the potential congestion there. Among the various proposals emerged so far, the network-based DMM solution derived from Proxy Mobile IPv6 (PMIPv6) appears to be the most promising [9].

In this article, we propose a flat mobile network architecture inspired by the DMM paradigm, where a centralized anchor and gateway is no longer in charge of routing traffic and handling mobility for a large amount of subscribers, but instead mobility-enabled access routers accomplish such functions. In addition, we focus on the benefits brought by a smart combination of a mobile video content delivery system built on top of our flat architecture. We believe that this type of solution provides a quite significant improvement in terms of load reduction in the packet core and communication latency. Our solution is first described, explaining the basics of how to combine DMM and mCDN solutions. Then we validate it and assess its performance, by reporting on experiments conducted over a network prototype employing real equipment.

II. DMM-ENABLED VIDEO CACHING ARCHITECTURE FOR FUTURE NETWORKS

In our architecture, illustrated in Fig. 2, mobility-enabled access routers (MARs) possess a link to the Internet to deviate the data plane (like the local breakout points in the EPS), as well as a link to the core network for the control plane. Also, a pool of IPv6 prefixes belongs exclusively to each MAR, from which the MAR assigns one to each Mobile Node (MN) connected to its access links. A MAR is then the MN's first IP hop and default router. After delegating the IPv6 prefix to the MN, a MAR maintains the reachability for that prefix even after the MN attaches to another MAR.

A MAR routes packets carrying the delegated prefix as a plain IPv6 router when the MN is connected to it. If a handover occurs, i.e., the MN moves to an access network managed by a different MAR, the MN is assigned another prefix from the new MAR's pool, and an IPv6-in-IPv6 tunnel is established between the new and old MARs. The prefix used previously by the MN is thence routed via the tunnel, whereas the new prefix is forwarded without any special handling by the new MAR. In this way, new IP flows started with the latest configured prefix benefit from the optimum route to destination (as the green IP flow in Fig. 2). IP packets bearing an old prefix are redirected through the inter-MARs tunnel, as depicted by the red flow in the picture, so that ongoing sessions are not lost.

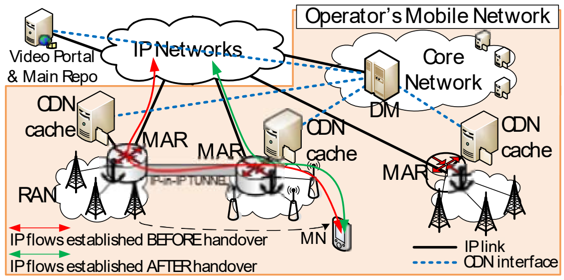

Fig. 2. Global system architecture.

The control plane follows the principles of the fully distributed network based solution that appears in [10]. This solution is a combination of an IP-based DMM protocol derived from PMIPv6 and the IEEE 802.21 Media Independent Handover Services (MIHS) standard, wherein both the data and control planes are managed by the MARs, avoiding the involvement of any centralized entity. Our DMM architecture is intended to be independent from the underneath radio access network, thus allowing for heterogeneous access technologies.

The use of our flat architecture enables the mobile user to be topologically closer to the connection point to external IP services, as compared to the current 3GPP architecture. Therefore, the user experience can be enhanced by placing content and services close to the MARs, which at the same time reduces significantly the load in the network core. Following this key concept, we propose the deployment of a mobile CDN in which the content caches are located in the MARs. We next develop a relevant use case in which we consider video traffic as the target of our study, as it is foreseen to play a dominant role in the years to come.

In our framework, the CDN system comprises a set of CDN caches co-located with the MARs and a controller, called Decision Manager (DM), that monitors and drives the activity of the caches. The video files are stored at a main repository, providing a web-based catalogue to choose them. Note that, in general, the video repository might belong to the mobile operator, or be a third-party service. The caches at the MARs store the most popular videos and host an HTTP proxy server to perform object-oriented caching based on the content's URL. Therefore, if the requested content is locally available at the cache, the video is delivered directly by the MAR. Otherwise, the MAR delegates the DM to decide where the content should be streamed from, which can be another MAR or the main central server. Thereby, the CDN infrastructure delivers the service from the closest node to the MN where

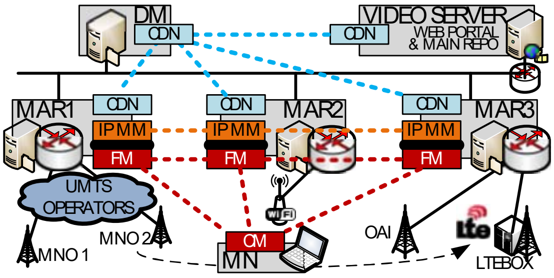

Fig. 3. System architecture test-bed deployment.

the content is available. The DM and the CDN nodes share a signaling interface used to process the video requests and to periodically synchronize the content cached by the MAR, based on the intended popularity metric.

This CDN+DMM architecture is capable of granting the committed service to the final users without forwarding data traffic through the P-GW and its links. This aspect limits the core congestion and reduces the communication latency.

III. Validation and Performance Assessment

The CDN+DMM architecture is implemented in a network prototype and validated through real field experiments to carry out a proof of concept of our design.

A. Test-bed and implementation details

The network prototype consists in a test-bed of GNU/Linux machines running Ubuntu 10.04 deployed as illustrated in Fig. 3, where we highlight the interconnection among elements along with the logical interactions between the nodes' functions. The prototype exhibits three control systems to coordinate the data delivery to moving users: i) the Layer-2 handover management, represented in the picture by the red boxes and dashed lines, ii) the IP mobility management, indicated by orange elements, and iii) the CDN system, whose boxes and interfaces are colored in light blue.

The Layer-2 handover management complies with the IEEE 802.21 protocol suite, Media Independent Handover Services. The IEEE 802.21 implementation used is the ODTONE[1] open platform, installed in all MARs and in the MN. In the MIHS framework, the MIHS-enabled functions export APIs to be used by MIHS users to execute commands and control the handover operations. We have deployed two MIHS users: the Flow Manager (FM), installed in the MARs, and the Connection Manager (CM), running in the MN.

The IP mobility management provides the IP address assignment to the MNs and their reachability within the domain, and is based on the Proxy Mobile IPv6 protocol. Its implementation has been realized modifying the code from the MAD-PMIPv6[2] project and it is installed in the MARs as the IP Mobility Module (IP MM).

The third control mechanism is the CDN system. Such system handles, for instance, the Dynamic Adaptive Streaming over HTTP (DASH) video delivery to the users. We deployed the video service's portal and global repository in the Video Server node, and the MARs are provided with video caches. With DASH, a video file is divided into many segments, which are downloaded using HTTP (after requesting the segments' URL from the Video Server), one by one. The Squid[3] HTTP proxy tool, installed in all MARs, is able to intercept HTTP requests and to accommodate them if the content pointed by the URL (i.e., the video chunk) is available in the local cache. Otherwise, the request is sent to the processing engine at the DM, that computes the best source, e.g., according to the content distribution and the MN location or based on other policies. Then, the DM instructs the MAR to retrieve the segment from that source, which might be the main central server or another CDN node (MAR). In both cases the traffic redirection is transparent to the MN.

Beyond the modules described above, each MAR provides a specific radio access technology. MAR1 supplies the 3G access through real UMTS networks from two commercial operators. The terminal connects alternatively to one or the other using an HSDPA USB dongle. An IPv6-in-IPv4 VPN tunnel is set up to pass through the UMTS network, and transparently join MN to MAR1 as if they were both on the same IPv6 local link. MAR2 offers the WiFi connectivity using a simple IEEE 802.11b/g card. Finally, MAR3 provides the 4G access, through two access modes that can be toggled. The first mode employs the Open Air Interface[4] (OAI): it is a software tool installed in both the MAR and MN that emulates over cable the LTE interface between an eNodeB and a User Equipment (UE). The second mode adopts an Alcatel-Lucent Bell Labs' LTEBOX: it consists of an embedded system hosting a commercial pico-cell eNodeB (2.6 GHz band) connected to a vEPC, necessary for the LTE bearer setup and user credentials. The LTEBOX is connected to MAR3 via an Ethernet link. This method requires an LTE USB dongle plugged in the MN to establish the radio link with the eNodeB.

On the MN's side, the CM groups the three network interfaces (WiFi, 3G and 4G) into a single one. This allows hiding the use of different links to the upper network layers, by showing only one "logical" interface. The CM also performs the MIHS-related signaling with the network upon handover.

B. Video streaming to mobile users scenario

We have examined the system's behavior when a mobile user changes access technology within the domain while playing a DASH video.

The user starts accessing the video portal's web page to pick a video from the on-line catalogue using the 3G connection. Once the video is chosen, the server first sends the list of all video chunks and their corresponding URLs. A typical DASH client[5] then proceeds to download and playback the listed chunks one by one, with a buffer that allows to download the next chunk shortly before the current playback is over. This mechanism prevents the client to consume all available

[1] http://hng.av.it.pt/projects/odtone/
[2] http://www.odmm.net/mad-pmipv6/
[3] http://www.squid-cache.org/
[4] http://www.openairinterface.org/
[5] We used VLC for Ubuntu, http://www.videolan.org

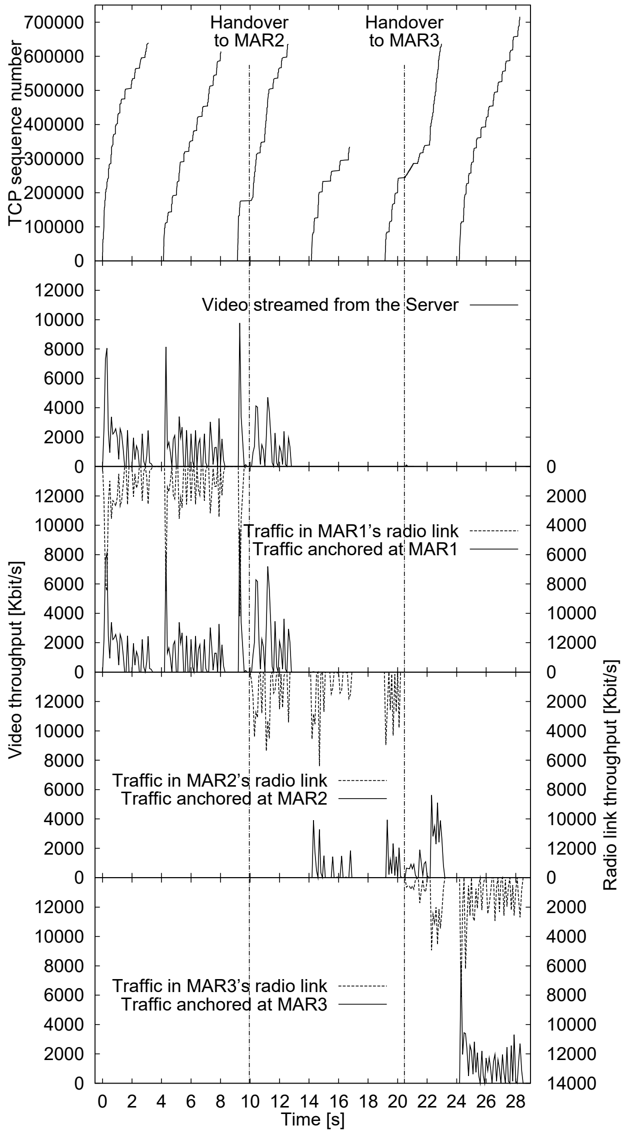

Fig. 4. Traffic captured in the system. From top to bottom: in the MN, in the Video Server, in MAR1, in MAR2 and in MAR3

bandwidth and to download unnecessary chunks if the user is not willing to watch the whole video.

The video for the experiments is formed by 6 chunks, each 5 seconds long, and we have pre-filled the MARs' caches such that the content requested by the user is available in MAR2 and MAR3, but not in MAR1. The MN is programmed to switch to the WiFi access around the video's playback second 10, and then to move to the LTE link around the 20th second, so that the handovers occur in the middle of the download of chunks #3 and #5. We have illustrated in Fig. 4 the throughput captured in the network's nodes. This figure presents 5 graphs, one for each network element: from top to bottom we have the MN, the Video Server, MAR1, MAR2 and MAR3. The top diagram is obtained capturing at the MN's logical interface: we can observe 6 ramp-like plots; each of them is the TCP sequence number evolution for the download of the video chunks[6]. In the second subplot, we have measured the throughput at the server's outbound interface. In the remaining graphs, the solid line represents the traffic anchored at that MAR, and the dashed line is the throughput in the access link.

As long as the MN is connected to MAR1 through the HSDPA link, the content source is the Video Server, because MAR1's local cache cannot fulfil the requests, and the DM instructs the MAR to retrieve the chunks from the main repository. In this period, the MN requests the first three chunks, therefore we can observe in the second and third graph the traffic at the Video Server and at MAR1 accounting for the packets belonging to such chunks. The activity of the MAR1's radio link is interrupted by the handover, but the MAR is still sending traffic for a short interval thereafter. In this moment the MN is handing off from MAR1's access link and attaching to MAR2's WiFi network. Right after the handover, MAR1 redirects the remaining packets of chunk #3 to MAR2, that can deliver them to the MN on the wireless link. We observe this phenomenon in the 4th subplot: MAR2's radio link is active but no traffic is coming from the video server or generated internally. The packets indeed flow from MAR1 to MAR2 through the IPv6-in-IPv6 tunnel established by the IP MMs.

When the MN starts the new request for chunk #4, it uses the address from MAR2, hence the request is not forwarded through the tunnel. In this way, the MAR2's CDN module can intercept the request and, since the chunk is locally available, it immediately transmits the packets to the MN. Thus, the fourth graph shows activity for the locally anchored traffic (solid line), as well as for the radio (dotted line). On the contrary, the second and third graphs, respectively for the server and MAR1, show no traffic for the fourth video segment. The same behavior occurs for chunk #5, up to the second handover takes place: the radio link activity is interrupted, but still MAR2 is sending traffic to the MN. Such packets are collected by MAR3 from the tunnel and delivered, replicating the same reaction observed after the first handover. Finally, the sixth chunk is delivered directly by MAR3.

The architecture proposed in this article makes it possible to deploy a CDN close to the access routers of a mobile network using a simple URL-based caching method. It is obvious that the communication latency can be reduced when the content is downloaded from such CDN caches, as well as the traffic in the core links.

C. Inter-technology handover evaluation

We next evaluate the performance of the system by measuring the latency experienced during a handover.

In our experiments, an MN performs a heterogeneous technology handover. When doing so, the terminal establishes the new radio link before releasing the old one, so there is always an active link to the network. This task is accomplished in three stages, i) the "Handover Preparation", ii) the "Handover

[6]Note that, even if all chunks are a 5-second portion of the video, they are not equally large in KBytes.

Execution" and iii) the "Handover Completion". Operations i) and iii) are realized through the IEEE 802.21 protocol, involving the CM and the FMs, while ii) is performed by the IP MM modules.

During the first stage, the MN's CM indicates to the network that the handover is imminent. Then, the FM residing in the current MAR negotiates the radio resources with the peer MARs available in the MN's surroundings. When the target MAR is ready, the MN establishes a radio link with it, concluding the preparation phase.

The Handover Execution stage is necessary to set up the IP configuration on the new radio link, to import the old link's IP setting from the previous MAR and to build the IPv6-in-IPv6 tunnels. Therefore, once the new radio link is up, the CM switches off the old one and triggers the usual IPv6 Router Discovery procedures in order to obtain IPv6 connectivity on the new link. Then, the new MAR invokes the IP MM functions to perform the IP mobility operations with the old MARs. This phase terminates when the MN is advertised the new prefix to start new sessions, along with the old prefixes to extend their usability for the ongoing flows.

Next, the Handover Completion stage groups the IEEE 802.21 transaction to release the resources on the old MAR and conclude the handover.

The Handover Execution is the most critical for the user's ongoing communications, because, the MN is neither able to send nor receive IP packets. Indeed, in this phase, even if the terminal has already established the new radio link, the IP setup (address and routing) is still associated to the old one.. Therefore we consider this interval as the handover latency, and it is the object of the measurements described below. Table I summarizes the mean and standard deviation values obtained from more than 800 handovers. It should be noted that such interval is lower bounded by the Round Trip Time (RTT) between the MN and the MAR in the target technology, because of the Router Discovery procedure in the access link. For instance, when the MN hands off to the HSDPA technology, such signaling has to traverse the whole operator's UMTS infrastructure, with a high global delay. In our experiments the two operators performed very differently. We also observed a significant difference when comparing the LTEBOX and the OAI platform: the former introduced a larger RTT because of the vEPC embedded in the system, while the latter was tuned to offer the lowest eNodeB-UE latency. WiFi is the technology that offers the best performance.

The results in Table I should be regarded as a proof-of-concept of our design, rather than a real system performance evaluation. Indeed, a real deployment is affected by many factors that degrade the performance. For example, the distance between the MARs involved in the handover adds delay to the handover latency, due to the inter-MARs mobility signaling, as well as the presence of multiple users in the same cell, due to the bandwidth share with other terminals.

TABLE I
HANDOVER LATENCY EXPERIMENTAL RESULTS

IP Handover to	Mean [s]	Std. Dev. [s]
3G (operator 1)	0.085	0.031
3G (operator 2)	0.163	0.033
WiFi	0.003	0.001
4G (OAI)	0.010	0.003
4G (LTEBOX)	0.110	0.023

IV. CONCLUSIONS

Mobile data traffic is foreseen to go through a dramatic growth in future years, and video is expected to be one of the major contributors. Mobile operators are currently investigating solutions to effectively tackle the problem, like flattening the mobile architecture or employing CDNs.

In this article we have introduced a mobile architecture design combining CDN with a mobility protocol based on the DMM paradigm. We have presented how its characteristics can be exploited in order to simplify video caching, thereby reducing the load in the core network. The flat architecture of DMM distributes the handling of the user's mobile context, thus offering the possibility to efficiently push content close to the edge of the mobile network. Therefore, the CDN and DMM concepts nicely inter-work and complement each other.

We have experimentally validated our integrated solution with a network prototype including different radio access technologies. Results show the feasibility of a combined CDN+DMM system, offering the possibility to deliver video contents without congesting the operator's core network and gaining shorter communication latency.

REFERENCES

[1] CISCO, "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2013-2018," in White Paper, Feb. 2014.

[2] 3GPP, "Local IP Access and Selected IP Traffic Offload (LIPA-SIPTO)," 3rd Generation Partnership Project (3GPP), TR 23.829, October 2011.

[3] ——, "LIPA Mobility and SIPTO at the Local Network," 3rd Generation Partnership Project (3GPP), TR 23.859, April 2013.

[4] J. Erman, A. Gerber, M. Hajiaghayi, D. Pei, S. Sen, and O. Spatscheck, "To Cache or not to Cache: The 3G case," Internet Computing, IEEE, vol. 15, no. 2, pp. 27–34, 2011.

[5] B. A. Ramanan, L. M. Drabeck, M. Haner, N. Nithi, T. E. Klein, and C. Sawkar, "Cacheability Analysis of HTTP traffic in an Operational LTE Network," in Wireless Telecommunications Symposium (WTS), 2013. IEEE, 2013, pp. 1–8.

[6] X. Wang, M. Chen, T. Taleb, A. Ksentini, and V. Leung, "Cache in the air: exploiting content caching and delivery techniques for 5G systems," IEEE Communications Magazine, vol. 52, no. 2, pp. 131–139, 2014.

[7] F. Z. Yousaf, M. Liebsch, A. Maeder, and S. Schmid, "Mobile CDN Enhancements for QoE-Improved Content Delivery in Mobile Operator Networks," Network, IEEE, vol. 27, no. 2, pp. 14–21, 2013.

[8] S. Woo, E. Jeong, S. Park, J. Lee, S. Ihm, and K. Park, "Comparison of caching strategies in modern cellular backhaul networks," in Proceeding of the 11th annual international conference on Mobile systems, applications, and services. ACM, 2013, pp. 319–332.

[9] D.-H. Shin, D. Moses, M. Venkatachalam, and S. Bagchi, "Distributed Mobility Management for efficient video delivery over all-IP mobile networks: Competing approaches," Network, IEEE, vol. 27, no. 2, pp. 28–33, 2013.

[10] F. Giust, A. Oliva, C. J. Bernardos, and R. Costa, "A Network-based Localized Mobility Solution for Distributed Mobility Management," Oct. 2011, IEEE Workshop on Mobility Management for Flat Networks (WMMFN).