Scalable Mobile Backhauling via Information-Centric Networking

Giovanna Carofiglio*, Massimo Gallo†, Luca Muscariello‡ Diego Perino †
* Cisco Systems, † Bell Labs, Alcatel-Lucent, ‡ Orange Labs Networks,
gcarofig@cisco.com, first.last@alcatel-lucent.com, luca.muscariello@orange.com

***Abstract*—The rapid traffic growth fueled by the spread of mobile devices and the availability of high speed network access calls for substantial innovation at network layer. The content-centric nature of current Internet usage highlights the limitations of the host-centric communication model in coping with dynamic content-to-location binding, mobility management, multicast, multi-homing, etc. If transmission capacity speedups in the backhaul may hide such inefficiencies in the short term, the host-centric communication model needs to be revisited to sustain the mobile demand on the long term.**

In this paper, we first identify and quantify the opportunities for backhaul evolution by analyzing a large set of traffic measurements collected between the mobile core and backhaul of Orange France. The analysis reveals that 50% of HTTP requests are cacheable and traffic can be reduced from 60% to 95% during the peak hour by using 350GBs to 1TB of memory overall. Motivated by such significant opportunities for latency reduction and network cost savings, we present a solution based on Information-Centric Networking (ICN). First results of a large scale experimentation using about 100 Linux servers with customized software, in a realistic network setting, provide a glimpse into ICN gains even in presence of naive caching policies: a factor three reduction in content delivery time and almost 40% bandwidth savings, when compared to existing alternatives.

I. Introduction

In the last few years, broadband access penetration and rich web services development have stimulated user demand beyond any previous expectation. Especially in the mobile segment, we witnessed fermenting innovation in mobile terminals and access rate increase. Among other factors, multiple access network convergence and all-IP architecture simplification via LTE pave the way to close the gap between fixed and mobile network user experiences.

Traffic growth is increasingly driven by video-centric and, more broadly, by content services dissemination. The content-centric nature of current Internet usage has turned the network into a medium to connect people with content, highlighting inefficiencies of the current host-centric communication model (e.g. difficulties in dynamic content-to-location binding, mobility management, multicast, multi-homing, etc.).

The observed traffic growth is predicted to significantly increase and already solicits mobile network operators (MNOs) for innovation and significant investments in the infrastructure to guarantee high quality connectivity and better user Quality of Experience (QoE). From a technical point of view, realizing efficient data delivery without capacity over-dimensioning demands for a flexible and adaptive traffic control, reducing congestion phenomena via effective statistical service multiplexing and by natively employing in-network caches.

Today there is a lack of mechanisms for dynamic and flexible traffic control within the backhaul for mostly two reasons:

- the transparent use of backhaul imposed by the 3GPP standard (IP tunneling);
- the connection-oriented nature of current application layer solutions and the use of a connection-oriented transport model (e.g HTTP on TCP).

This results in a backhaul used as a passive network segment, where congestion phenomena are suffered and not dynamically managed. To prevent backhaul saturation different edge caching solutions have been proposed in the last year by startup companies: to cite a few, DatE (Data-at-the-Edge) product by Altobridge, iQstream solution by Sycamore networks, LiquidNet by Nokia. These products apply the concept of micro-caching at the edge of the network and specifically at base stations to reduce backhaul load, while improving consumer experience by delivering content faster. However, they share two major drawbacks: they provide limited performance gains because of their applicability at the network edge only and of the use of HTTP-based, TCP connection-oriented transport model. Alternative attempts, not yet technically specified, are also available at ETSI [10].

In this paper, we show how Information-Centric Networking (ICN) provides a natural answer to such needs by rethinking data delivery in the network via a connectionless name-based transport model over IP, by enabling in-network caching and multi-point to multi-point communications (cfr.[6], [14], [2] for a primer to ICN).

ICN basic principles include name-based routing, pull-based multi-point to multi-point communication controlled by user requests and in-path caching as a built-in network feature. It results in a connectionless data delivery where packet requests are sent to the network with no need for a point-to-point connection setup. As the end user and the content are not bounded to a network addressable interface network mobility is part of the communication model itself. In the mobile backhaul, ICN creates the lacking protocol substrate for managing traffic growth and matching it with the minimum required resources for a given QoS. Also, it brings into the network the lacking content-awareness offering to operators the capability to better manage their resources and to pursue novel business opportunities.

If the shift brought by ICN to networking primitives is truly disruptive, such substantial gains cannot be achieved by application-layer solutions relying on the current transport

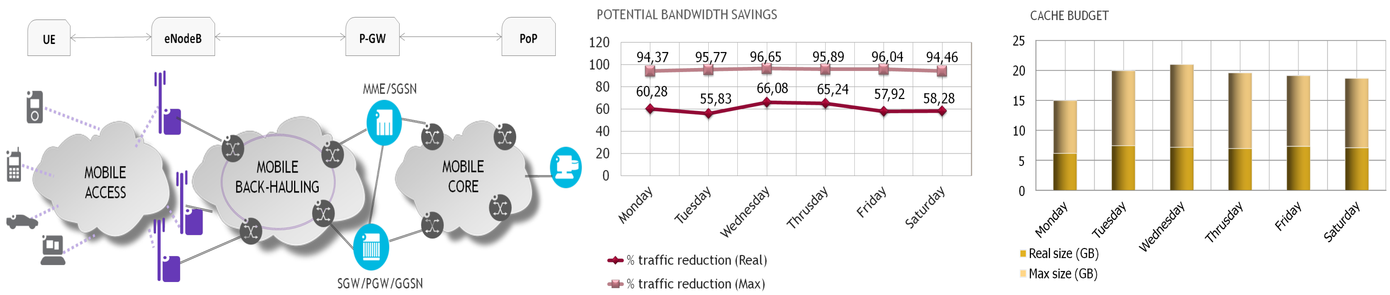

(a) (b) (c)

Fig. 1. (a) Network segments identification; (b) Bandwidth savings from traffic measurements; (c) Cache budget requirements in a regional area.

model. To the best of our knowledge, this is the first quantification of the benefits of a complete ICN design.

In the paper, we first present an analysis of traffic measurements collected over the nation-wide Orange mobile network in France and revealing that 50% of HTTP requests are cacheable and traffic can be reduced from 60% to 95% during the peak hour by using limited memories (Sec.II). Further, by means of extensive experimentation made in a large scale test-bed under realistic settings, we show how our ICN solution enables such gain and overperforms currently envisaged alternatives (Secc.III, IV). Finally, we discuss a feasible integration of ICN in the mobile backhaul that limites upgrades to currently deployed technologies in Sec.V and conclude the paper in Sec.VI.

II. Mobile backhaul opportunities

This section reports some facts about mobile traffic characteristics and insights on the opportunities of caching in terms of transport cost reduction and user performance improvement.

We rely on passive measurement points monitoring traffic between the mobile core and the mobile backhaul, placed at various SGWs (Fig.1(a)). It is worth noticing that the current infrastructure serves different radio access technologies: 2G, 3G (with its different evolutions), 4G via a unified backhaul network deployment. To evaluate the potential of embedding in-network caching capabilities, we quantify the fraction of data that could be cached in the network. To this aim, we start by collecting all IP packets belonging to a portion of total nationwide network coverage, during the peak hours over an entire week. The sample includes packets flowing upstream and downstream in order to rebuild the complete set of HTTP sessions delivered to the active terminals. The average downstream rate in the peak hour is about 250Mbps for an observed user fan-out of tens of thousands of terminals. We focus on HTTP sessions of the following predominant applications: (i) Web browsing; (ii) Audio/Video : all HTTP traffic characterized including You Tube traffic; (iii) You Tube: representing a significant portion of total Audio/Video traffic category.

1) Cacheability: For the three traffic categories, we quantify the cacheability, namely the percentage of requests related to objects requested at least twice in the considered time period (e.g. the peak hour). In average 52% of total requests result to be cacheable according to the definition of cacheability given above. Audio/video applications and You Tube in particular can attain values up to 86%, which constitute a significant potential. We derive hereafter the implications in terms of bandwidth savings and associated amount of useful memory to place into the network. It is also important to notice that such statistics remain stable during the whole week. Similar numbers have been reported in [12] in the USA and [13] in South Korea.

2) Potential bandwidth saving: To accurately quantify amount of bandwidth that can be saved by caching data we analyze in detail the amount of data delivered by each web transaction, which may differ even for the same content, i.e. users may retrieve only partially the same content. Partial content retrieval might be caused by users lack of interest towards some web content or by users impatience caused by poor network performance. In many cases, different components of a web transaction might be locally cached in the mobile terminal. Moreover, the same content item is often delivered to different users with additional customized information: targeted advertisement or personal information. In average, we observe that *about 50% of web transactions downloads the whole content (maximum observed size) while the left 50% downloads only fractions of it*. We take into account such characteristics of data traffic to evaluate our key statistics. In Figure 1(b), we report the percentage of saved bandwidth during the peak hour by using in-network storage for cacheable data of the size reported in Figure 1(c). The real values are compared with the case where every web transaction is assumed to always download the entire content.

Note that the effect of all users downloading the entire objects would be a traffic increase up to six times during the peak hour. Such phenomenon would be impossible to control neither with traffic management systems available in today networks, nor by using the most common levels of link over-provisioning. The interesting finding is that, in such adverse conditions, by using three times larger but still limited memories (as shown in Figure 1(c)), traffic can be reduced up to 95% during the peak hour. The incentives for an ISP to cache traffic within the backhaul become clear by observing the spectacular potential benefits in terms of traffic reduction.

3) Cache budget: The computed gains can be obtained by using very small caches. To be convinced about this claim, it is sufficient to observe that the considered regional area requires from 7 to 20 GB of storage for a single peak hour. If we consider a worst case scenario for a national network composed of 100 such regional areas and a traffic peak period of 5 hours, the total required storage would range from 350GB to 1TB. The computation assumes that in all 100 regions and during each peak hour, the requested web content would change completely. Even in the worst case scenario, the memory required is very small and implementations would be feasible with technologies enabling at high speed networking in hardware [11]. While the presented calculations are just rough estimations, they suggest a rather wide spectrum of design choices and technological options to achieve significant gains, for viable network backhaul evolutions in the long term.

4) Locality of requests: The above presented statistics prove that it is a misguided belief that caching data close to the users, e.g. at regional level, as we propose in the backhaul, is inefficient. Such belief is often based on the opinion that user fan-out is too low that deep in the network, and memory sharing would have little interest. Another misguided belief about caching is that the required storage must be very large to achieve measurable gains. Such belief is based on estimations of total available web content in the Internet which scales as 10^{10} and makes the assumption that content is uniformly spread around the globe. Instead, we have observed popularity is very much affected by spatial and temporal locality leading to a regional segmentation of all Internet content that is also popular within relatively short time windows. Thus, the relevant portion of content should be that measured during the peak hours in view of reducing bandwidth usage and network congestion. Indeed, there is little value in saving more bandwidth when the network is already lightly loaded (no additional latency reduction).

III. EXPERIMENTAL FRAMEWORK

In this section we introduce the experimental scenario used for the quantification of ICN benefits in the mobile backhaul (Sec.IV). We first describe the topology and traffic demand used in our evaluation. We then introduce the evaluated solutions and details of the experimental platform.

A. Network topology and Traffic Demand

Figure 2 reports the network topology that we use in our experiments, which is a down-scaled model of the backhaul topology of a mobile operator. It consists of four PDN GWs connected by a full mesh. SGWs are assumed to be co-located with the PDN-GW, as it is typically the case in todays deployments. In addition, we assume to have two CDN servers external to the backhaul network reached via two of the four PDN-GWs. Each PDN-GW node is the root of a fat tree topology composed of 20 nodes, including PDN-GW and eNodeBs. Each eNodeB aggregates traffic generated by users in three adjacent cells. Cells are dimensioned so that every eNodeB serves the same average traffic demand, i.e. users are uniformly distributed over eNodeBs. Network links capacities are the same for the four fat tree topologies and reported only for one of them in Figure 2. Moreover, network links are drawn in continuous or dotted lines when modeling primary or secondary links. Secondary links are used only in case of failure in today backhaul networks while always available in the ICN mobile solution.

The workload model used to generate mobile users demand is taken from the trace characterization analysis presented in [4], down-scaling scaling the amount of available content items and number of active users for test-bed scalability. Users request content items belonging to two different content catalogs (i.e. two different services), according to a Poisson process of rate $\lambda = 0.4$. Each catalog consists of 10,000 content items divided in 250 data chunks of size 4096 bytes. CDN servers installed outside the backhaul network permanently store a copy of one catalog each. In the experiments, users belonging to trees from 1 to 4 respectively request 95%, 90%, 85% and 80% of Catalog 1 content items, while the rest of the requests is for items of Catalog 2. Finally, content catalogs' popularity is modeled by a Weibull distribution with shape 0.8 and scale 500 which mirrors the trimodal distribution identified in [4].

B. Evaluated solutions

In the experiments presented in this paper, we compare an ICN backhaul deployment with today solution based on single path and without in-network caching capabilities. For the sake of completeness, we add to the set of experiments two additional alternatives, currently under development by different startups (as e.g. iQstream) integrating caching at PDN GWs or at NodeBs. These two GTP tunnel endpoints, as preconized by the 3GPP standard, are the unique locations where caching could be implemented with relatively little efforts. However, caching within the backhaul cannot be implemented, currently, without breaking the mobility management model standardized at 3GPP, based on protocol anchors in the PDN GWs and GTP tunnels across the entire backhaul.

In the experiments, we consider three architectures for the mobile backhaul and report performance gains of each solution normalized with respect to the baseline architecture not implementing caching:

- ICN, introduced in the mobile backhaul, with transport, multi-path forwarding and caching implemented and enabled (see [2] for a detailed description of the used optimized protocols). We assume a naive non-uniform caching budget deployment: more cache memory is installed close to the users, hence in first and second level cache (4 times w.r.t. other network nodes). As an important remark, we consider in this paper naive LRU (Least Recently Used) caching for a fair comparison with other alternatives. However, the space for improvement to approximate optimal caching policies is huge either by using popularity prediction and content placement or smart distributed caching and content eviction policies. Some of these mechanisms are actually already available

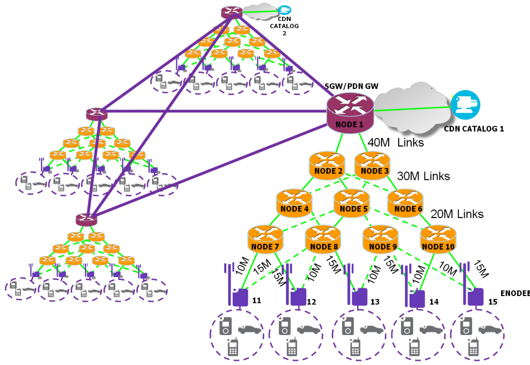

Fig. 2. Network testbed setup description.

and in part implemented in the software prototype described in Sec.III-C.

- PDNCache: caches are deployed at PDN GWs only. As for the baseline case with no caches, all traffic is routed through a single shortest path between eNodeBs and PDN GWs.
- eNodeBCache: caches are deployed at eNodeBs only. As for the baseline and the PDNCache scenarios, traffic follows the shortest path between eNodeBs and PDN GWs.
- ICN+PDNCache - In addition to ICN in the backhaul, we also consider the placement of a large cache at PDN-GW, that would be a feasible solution for a drastic reduction of network cost both in the backhaul and towards the core.

Finally, notice that in all the scenarios involving caching, we assume that total caching budget distributed in network nodes remains constant (i.e. 6GB) for all the solutions being compared, except ICN+PDNCache.

C. Experimental platform

For our evaluation, we use a customized version of the open source CCNx prototype [9], where we implement our ICN design, more specifically caching, forwarding and transport mechanisms (as in [2]). While the original goal of the prototype is to implement ICN communication mechanisms, we used it to evaluate all three solutions described above.

To run our experiments, we used a nationwide data center composed of a few thousands of general purpose servers: the Grid'5000 facility [5]. The infrastructure allows booting custom Linux kernel images, including our customized version of CCNx software. In our testbed, IP is used as a link layer technology and link rates are parameterized by using token buckets rate shapers. Moreover, a CCNx overlay network is built on top of IP and name based routing is configured according to pre-computed name based FIBs.

Experiments configuration, management and monitoring is performed using a centralized application that orchestrates all required operations in a scalable way: from servers' reservation, kernel booting, network configuration, experiment launching and test logs collection at the end of each test. Performance measurements are then reported in the following section. The entire set of softwares and tools are now made available in open source for the new NDN prototype [7], [8].

IV. Experimental results

The section gathers the results of the evaluation carried out using the previously described framework. We quantify either the improvement brought by ICN on the QoE perceived by end users, either the reduction in terms of cost for the network provider. Further, we analyze the reactivity of the ICN solution to the arise of a congestion phenomenon and its ability to adapt to a sudden variation in network conditions, as compared to current solutions.

A. Improving User QoE

To assess the impact of an ICN deployment on user quality of experience, we measure the average delivery time on all content and on a per-catalog base. Results reported in Figures 3(a), 3(b) are normalized with respect to the baseline scenario without caching (i.e. the current deployment). The most striking result is that ICN can reduce up to a factor 3 the average content delivery time, thus largely improving user quality of experience (a factor 2.5 on all items, 3 on the most popular ones). While the introduction of caching at eNodeB and PDNCache also provide respectively 26% and 14% average delivery time reduction, ICN gains are significantly larger (60%). The improved user QoE is due to the introduction of two ICN features (and their interplay):

- in-network caching: a distributed caching approach with larger cache close to the users significantly reduces backhaul traffic, thus increasing per-flow bandwidth share and throughput;
- dynamic multipath transfer: the use of multiple paths in parallel is a clear advantage of ICN in terms of faster download and finer adaptation (at packet-level) to available network resources. This results in a better utilization of network resources and, consequently, in a shorter content delivery time.

If we look at the detail of per-application performance, gains slightly vary over catalog as a consequence of the traffic breakdown, but they remain proportionally of the same order of magnitude for each catalog. In Figure 3(b), we report the dependency of the content delivery time on popularity distribution: we distinguish the average content delivery time experienced by the 6 most popular objects within each catalog and that of less popular items. ICN shows an even larger reduction of content delivery time for popular items, while a slightly larger-than-average delivery time for the less popular ones. Overall, results highlight the effectiveness of the pervasive ICN caching and multipath forwarding in the mobile backhaul and confirm the enormous potential related to traffic cacheability, quantified by the statistics presented in previous sections.

B. Reducing network cost

We now quantify the cost reduction in terms of bandwidth/link load that ICN would provide if deployed within the mobile backhaul. To this purpose, we measure the total

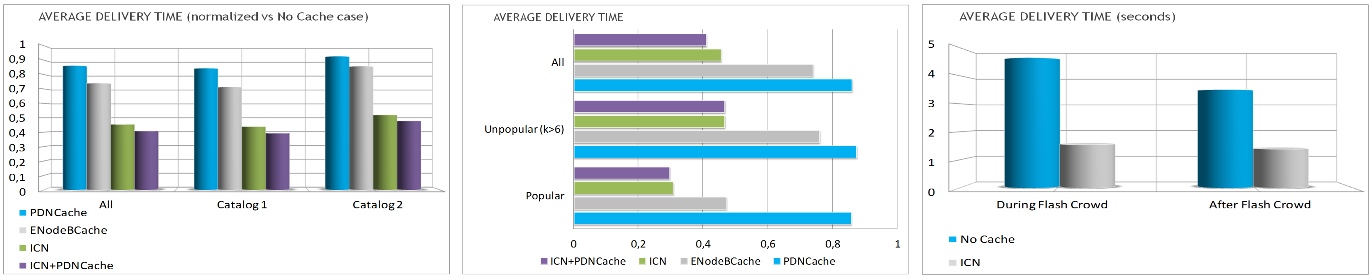

(a) (b) (c)

Fig. 3. User's performance indicators: average delivery time. Statistics reported as global averages (a), and as a function of content-type and content popularity (b). The impact du to flash crow is also reported (c).

bandwidth saving with respect to the baseline scenario without caching. Figures 4(a), 4(b), 4(c) report the saving of different solutions w.r.t. the baseline scenario with no caching. We present the total amount of bandwidth saved within the backhaul and on links connecting the backhaul (PDN GWs) and the core network, in Figure 4(a), 4(b) respectively. ICN allows saving up to 38% total bandwidth within the backhaul (33% with ICN in backhaul only, 38% with ICN+PDNCache). The alternative solutions PDNCache and eNodeBCache eNodeB-Cache, show 10% and 23% bandwidth saving respectively. The bandwidth reduction towards content servers within the core shows good results for ICN only (42%), while it is clearly more important for the PDNCache approach, where the whole cache budget is concentrated in a single point. These results clearly highlight a trade-off: putting caches closer to the users helps reducing the overall bandwidth consumption of the backhaul, at the price of a higher bandwidth usage towards the core. Differently, caching at the PDN GWs is more effective to save bandwidth towards the core, at the price of higher bandwidth consumption in the backhaul. ICN provides a good balance on such trade trade-off, while realizing PDNCache savings when coupled with PDNCache (77% bandwidth savings at server for the ICN+PDNCache approach). Finally, in Figure 4(c) we report the maximum link load reduction w.r.t. the No Cache scenario. Unlike single path alternatives, ICN multipath transport minimizes the maximum link load by load balancing traffic flows at packet granularity. This leads to a reduction of the maximum link load of 71%, largely better than the other solutions.

C. Enhancing network flexibility

The characteristics of ICN transport, namely its connectionless nature, the rate and congestion control placed at the receiver and the dynamic multipath packet forwarding, guarantee high reactivity to network changes a and, hence, enhanced flexibility in facing sudden variations of traffic load with respect to currently deployed solution. To analyze these properties, we emulate a flash-crowd crowd phenomenon, where the request rate at eNodeB 7 swiftly increases. The flash crowd starts at time 0 and ends after 3000 seconds creating a congestion period on the link connecting node 7 to node 4 (i.e., link 7-4). In Figure 3(c) we analyze the impact of the flash crowd on the performance of users in the cell. We observe the delivery time is significantly affected during the flash crowd in the scenario without caching, while user's performance in the ICN case is not impacted as fast request forwarding is able to select a non congested solution. Overall, results highlight that ICN enhances network flexibility with high reactivity to dynamic conditions: this allows preserving user performance and reducing backhaul network congestion during sudden variations of the traffic load. With the current solution links should be largely overprovisioned to preserve user QoE during sudden congestion phenomena.

V. ICN DESIGN FOR MOBILE BACKHAUL INTEGRATION

The integration of ICN to the backhaul infrastructure requires few upgrades to the current protocols as defined by the LTE 3GPP standard [1]. In this section, we provide a sketch of a possible implementation to show the feasibility of such upgrades and the limited impact to already deployed networks. We mostly provide details about an implementation of the solution for the data plane. Control plane is out of scope of the present paper and requires substantially less modifications.

Current backhaul architecture - The mobile backhaul network is composed of IP/Ethernet routers organized in a hierarchical topology. Some of the backhaul devices are equipped with additional elements in charge of specific 3GPP data and control plane functionality, i.e. eNodeB, SGW (Serving Gateway) and PDN GW (Packet Data Network Gateway). To manage network mobility, user data plane traffic is encapsulated in tunnels (GTP-U) between eNodeB, SGW and PDN GW. Two additional functional elements, the MME (Mobility Management Entity) and PCRF (Policy Charging and Rules function) completes the backhaul control plane.

Enhancing backhaul equipments with ICN - We propose to deploy ICN as an overlay over IP. As several devices are based on programmable elements (e.g. Network processors), it is possible to integrate ICN mechanisms via software upgrades, without installing any new ad-hoc hardware. Also, we can assume that the memory required for in-network caching is measured in few tens of GBs (Sec.II): this amount of storage is currently already available in many deployed equipments.

Towards an ICN header - ICN mechanisms are based on named data packets. Therefore,we define an ICN extended header in the GTP-U encapsulation, whose interpretation is not required by all intermediate nodes: if a node does not

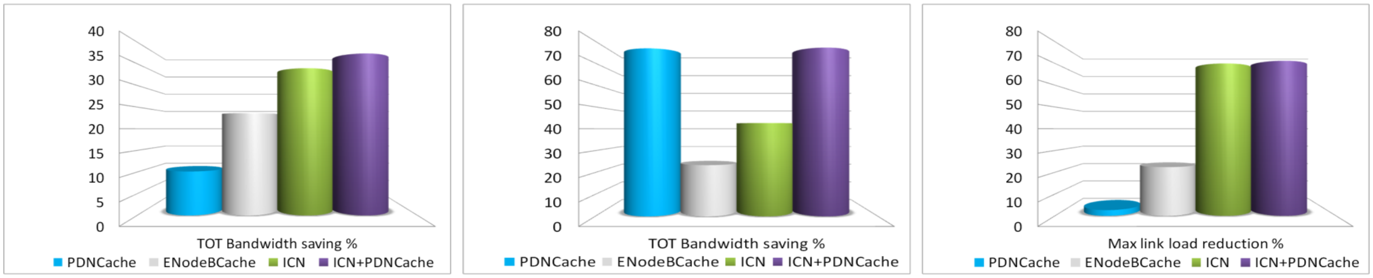

(a) (b) (c)

Fig. 4. Network performance indicators: bandwidth savings in the backhaul network (a), load reduction at the content server (b) and max network load reduction (c).

support ICN mechanisms or if the ICN extended header is not present, the standard forwarding policy is applied.

ICN data delivery - Today, the data retrieval process is driven by the user via application-layer solutions (e.g., HTTP). We propose two solutions to introduce the ICN transport model and to generate packets with an ICN header: i) a proxy in charge of performing ICN data retrieval on behalf of the user and co-located with a backhaul network element (e.g. eNodeB). The proxy generates ICN headers from application requests and insert them in GTP-U packets; ii) ad-hoc applications at the user terminal that are aware of the ICN substrate. A backhaul network element will then be in charge of integrating ICN headers in GTP-U headers.

Application server - A requested data packet can be found in the local cache of any ICN-enabled backhaul node. Every node then periodically reports to the control plane elements statistics about served data packets for policy enforcement and charging purposes. We propose to complete our design with a CDN server installed at PDN GW. This server is programmed to be aware of the ICN substrate, and thus to process packets containing ICN headers directly. In addition the server can use standard CDN communication mechanisms to retrieve content from other CDN servers outside the backhaul network.

VI. Conclusions

In this paper we discuss challenges and opportunities faced by mobile networks in coping with the unprecedented data traffic growth of the last years. We first identify and quantify the opportunities for mobile backhaul evolutions by analyzing a large set of traffic measurements collected over the nation-wide Orange mobile network in France. The analysis reveals that 50% of HTTP requests are cacheable and traffic can be reduced from 60% to 95% during the peak hour by using few GBs of memory in network equipment. Motivated by such significant opportunities for latency reduction and network cost savings, we present a solution based on Information-Centric Networking. By means of large scale experimentation using about 100 Linux servers with customized software, we provide a glimpse into ICN gains, in comparison with the currently deployed solution and potential alternatives for the backhaul evolution. Even in presence of non-optimized naive caching policies, ICN is shown reduce by three times content delivery time and almost 40% bandwidth savings. A fast reactivity to network changes is also proven in facing sudden variations of traffic load. As also reported in [3] ICN does not provide enough gains if analyzed as a caching technology only. In this paper we have shown that ICN brigs substantial gains in terms of user's performance by managing congestion in a much more efficient way than TCP/IP. This is obtained by jointly enabling caching, multipath and adaptive forwarding within the network. The entire experimentation framework and the software router has been release in open source at [7].

Finally, we propose a potential integration of ICN in the mobile backhaul which is 3GPP standard compliant and requiring limited upgrades to existing technologies.

Acknowledgments

This research work has been partially funded by the Technological Research Institute SystemX, within the project "Network Architectures" hosted at LINCS.

References

[1] 3GPP. GTPv2 for evolved GPRS. 3GPP TS 29.274 V8.1.0 (2009-03).
[2] G. Carofiglio, M. Gallo, L. Muscariello, M. Papalini, and S. Wang. Optimal Multipath Congestion Control and Request Forwarding in Information-Centric Networks. In *Proc. of IEEE ICNP*, 2013.
[3] S. K. Fayazbakhsh, Y. Lin, A. Tootoonchian, A. Ghodsi, T. Koponen, B. Maggs, K. Ng, V. Sekar, and S. Shenker. Less pain, most of the gain: Incrementally deployable icn. In *Proc. ACM SIGCOMM*, 2013.
[4] C. Imbrenda, L. Muscariello, and D. Rossi. Analyzing cacheable traffic in ISP access networks for micro CDN applications via content-centric networking. In *Proc. of ACM ICN*, 2014.
[5] INRIA and UPMC. The Grid'5000 Experimentation Test-Bed. http://www.grid5000.fr.
[6] V. Jacobson, D. Smetters, J. Thornton, and al. Networking named content. In *Proc. of ACM CoNEXT*, 2009.
[7] L. Muscariello and M. Gallo. The lurch tool. http://systemx.enst.fr/lurch.
[8] NDN NSF project. The ndn codebase. http://named-data.net/codebase/platform.
[9] Palo Alto Research Center. The CCNx project. http://www.ccnx.org.
[10] M. Patel, B. Naughton, C. Chan, N. Sprecher, S. Abeta, and A. Neal. Mobile-edge computing - introductory technical white paper. Sept.2014 ETSI.
[11] D. Perino and M. Varvello. A reality check for content centric networking. In *Proc.of the ACM SIGCOMM ICN*, 2011.
[12] B. Ramanan, L. Drabeck, M. Haner, N. Nithi, T. Klein, and C. Sawkar. Cacheability analysis of HTTP traffic in an operational LTE network. In *In Proc. of WTS*, 2013.
[13] S. Woo, E. Jeong, S. Park, J. Lee, S. Ihm, and K. Park. Comparison of caching strategies in modern cellular backhaul networks. In *Proc. of ACM MobiSys*, 2013.
[14] L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, k. claffy, P. Crowley, C. Papadopoulos, L. Wang, and B. Zhang. Named data networking. *SIGCOMM Comput. Commun. Rev.*, 44(3):66–73, Jul 2014.