NEXT GENERATION INTERNET

Open Call 1

Measuring Multi-Carrier Cellular Access International Roaming Performance

Deliverable 3: Experiment Results and Final Report

Authors	Marcelo Bagnulo, Ivan Vidal, Jonathan Almodovar, Universidad Carlos III de Madrid, Sergi Alcalá (Universidad Carlos III de Madrid & IMDEA Networks), Aravindh Raman, Diego Perino, Telefónica, Fabián Bustamante, NWU.
Due Date	31st July 2021
Submission Date	23rd July 2021
Keywords	Multi-carrier, performance, roaming

Deliverable 3: Part I

Analysis, results, and wider impact

1 Abstract

Multi-Carrier Cellular Access (MCCA) providers such as Google Fi irrupted recently in the mobile Internet access market claiming to offer enhanced performance at reduced prices. MCCA providers dynamically attach to different carriers based on real time measurements of performance. The goal of the project is to do an experimental comparative study of the performance of these emerging MCCA providers and the incumbent cellular providers while roaming. The proposed methodology is to acquire MCCA services and traditional mobile

services in the US and experimentally compare their performance while roaming (EU). Due to its very nature, measuring international roaming poses significant technical and logistical challenges, as tight cooperation between researchers in different countries is required. This was successfully achieved in this project thanks to a close collaboration between EU and USA partners.

2 Project Vision

Multi-Carrier Cellular Access (MCCA) allows cellular users to dynamically connect to different cellular networks without switching their SIM card. There are a number of advantages that can be realized through MCCA. MCCA users can improve coverage and performance, combining measurement-based dynamic carrier selection and predefined preferences on radio access technology (RAT) use (3G, 4G, 5G, Wi-Fi), irrespectively of which carrier is offering it a particular location.

Google Fi is one provider offering commercial MCCA services. Google Fi carrier selection policy relies on static preferences regarding available carriers and RATs, while also leveraging local or crowdsourced monitors to track various metrics to inform carrier selection [Yuan18].

While MCCA services are becoming increasingly popular, little is known about their technical intricacies beyond what is described in their commercial offers. To date, there is no quantitative-based analysis of the relative performance benefits of MCCA services compared to traditional services. In particular, there is no information about MCCA performance while roaming, which is one of the key selling points of several of these offers.

In a previous study, we analysed the European roaming ecosystem for traditional carriers and its performance [Mandalari18]. We observed that international roaming imposes performance penalties because traditional carriers systematically use Home Routing (HR) while roaming. HR implies that traffic from and to the roaming device is routed back to the home carrier before reaching the Internet, instead of exiting directly through the visited-network Internet access. This extra leg adds latency, penalizing roaming performance. While it would be technically feasible to provide local Internet access to roamers (called Local Break-out), this is avoided by traditional carriers in order to enable visibility of the roaming-client traffic by the home operator. Visibility of roaming-client traffic is deemed necessary to avoid disputes and abuses involving the computation of the traffic carried by the visited network on behalf of the roaming client. Seizing the performance gains resulting from local breakout is yet another potential benefit that MCCA can achieve for their customers in the international roaming context. Measuring international roaming performance is challenging, as it requires international cooperation and logistics are daunting, especially when inter-continental roaming is involved. In this project, we were able to overcome these challenges through a tight collaboration between the EU and USA partners. This involved, acquiring

Co-funded by the Horizon 2020 Framework Programme of the European Union

services and devices in the USA, shipping and distributing them in different countries in the EU in order to perform the measurements.

The objective of the project is characterising the performance of MCCA services while roaming. We will collect quantitative information to help us understand the benefits and downsides of MCCA service when used abroad. We will focus our analysis to clients from MCCA operators in the US that are roaming in the EU and investigate how their performance compares to both traditional US carriers while roaming and local EU carriers while providing local services in the visited network.

3 Details on participants (both EU and US)

Universidad Carlos III de Madrid: Marcelo Bagnulo (male, h-index=24, total citations=3817) is an associate professor UC3M. His research interests include Internet architecture and protocols, transport protocols and measurements. Marcelo has published more than 70 papers in the field of communications in journals and congresses (including Infocom, Mobicom and IEEE/ACM Transactions on Networking) and he is the author of 19 RFCs in the Internet Engineering Task Force (IETF) including the NAT64/DNS64 tools suite for IPv6 transition, and the LMAP framework for large measurement platforms, which have been widely adopted by industry. Marcelo was a member of the Internet Architecture Board (2009-2011). He is the coordinator of the project and will be PI for UC3M. Marcelo will work on the measurements of mobile access performance while roaming in the EU. Ivan Vidal works as Visiting Professor at Universidad Carlos III de Madrid, in the Department of Telematics Engineering. He received the Telecommunication Engineering degree in 2001 from the University of Vigo (Spain), and the Ph.D. in Telematics Engineering in 2008 from Universidad Carlos III de Madrid (Spain). Dr. Vidal has been involved in several international, national and regional research projects, and he has participated in different contracts with public administrations and other organizations, carrying out diverse activities related with consulting, technical assistance and research. These activities include his work as network architect in the regional research network of the Community of Madrid (REDIMadrid). Dr. Vidal has published over 60 research papers in conferences and international journals and magazines, and has served in the technical program committee of different conferences and workshops. His research interests include 5G networks, Unmanned aerial vehicles (UAVs), network security, and multimedia networking. He is currently member of the editorial board of the Internet Technology Letters (Wiley). Jonathan Almodóvar studied at Universidad Carlos III de Madrid, obtaining a BSc. degree in Computer Systems. He worked at Telefonica data centre, as part of systems and infrastructure group of Telyco, giving L1/L2 IT support for Linux servers and as shared storage/backup administrator. His main areas of interests are virtualization and cloud computing. Sergi Alcala is a PhD student hired to contribute on the execution of the project on the UC3M side.

Co-funded by the Horizon 2020 Framework Programme of the European Union

Telefonica: Diego Perino is the Director of Telefonica Research and he leads a team of researchers in Machine Learning, HCI, Privacy/Security and Networks and Systems. Currently, he also leads ARIANA, a Telefonica Innovation start-up project leveraging Artificial Intelligence for Automatic Network Actions. He strongly focusses on partnership with business units to ensure the successful exploitation of research outputs and to help them to solve difficult technical challenges with high relevance. His research work is strengthened via the creation and leading of strong collaborative associations with academic institution, industrial partners, customers, and government agencies. He also focusses on dissemination of research results via publications and demonstrations in major venues worldwide, and generation of intellectual property. Aravindh Raman (aravindh.raman@telefonica.com) is a researcher at Telefonica. He holds a bachelor's degree from Anna University, Chennai, India and MPhil/PhD degree from King's College London, UK, both in computer science. Previously, he was a researcher at academic institutions including King's College London, Queen Mary University of London, Cambridge computer lab and Indian Institute of Technology Delhi. He also worked as a member of technical staff at Gram Vaani Community Media. Aravindh's research centres around contributing to a robust future Internet infrastructure through network measurements and building next generation content delivery architectures. His recent achievements include a Best Paper Award at SIGCOMM Mobile Edge Computing in 2017 and a Best Paper Honourable Mention at WWW 2018. More details about him can be found here: https://aravindhr.am

Fabián E. Bustamante (male, h-index=33, total citations=4761) is a professor at Northwestern University. His does research on computer networks and distributed systems following an experimental approach. Fabian regularly publishes in top conferences, including IMC, Mobicom, Hotnets, WWW. Fabian received the NSF CAREER award and the Science Foundation of Ireland E.T.S. Walton Visitor Award. He has released tens of open-source systems that have gained ~2 million users worldwide. Fabian has co-chaired multiple events, including the first US NSF/FCC Workshop on QoE, the first SIGCOMM Workshop on QoE and the SIGCOMM 2014 conference; he is currently co-chairing the IMC'2020. He will be PI for NWU. Fabian will work on the measurements of mobile access performance while roaming in the EU.

In addition, Ozgu Olay, from University of Oslo is contributing to the project. Ozgu was the leader of the team who built the MONROE platform and among many other contributions, she is helping doing measurements on Norway, to add an additional visited country.

4 Results

4.1 Experiment Description and Implementation Details

MCCA provider selection.

It is not trivial to identify MCCA providers solely from the commercial offer. In particular, it is not easy to distinguish them from virtual mobile network operator (VMNO). The essential difference is that while the VMNO uses a single carrier, MCCA uses multiple of them, but that it is not apparent in all commercial offers. We looked into a number of providers, including Google Fi, Truphone, Apple SIM, Keepgo, Twilio, Gigsky, Flexiroam, Airalo, Surfroam, Skytone. Out of these, we acquired a Google Fi subscription, which is far from trivial, since the Google Fi service is only available in the US. So, in order to obtain the Google Fi subscription, we relied on our US partner. Once purchased, the Google Fi subscription MUST be activated in the US BEFORE it can be used in other countries. We relied also in our US partner for the activation and then started the experiments in Spain. In addition to Google Fi, we are acquired a Truphone subscription as an additional MCCA service provider.

MNOS selection

In order to do the experiments with Google Fi and Truphone, we need to acquire SIMs from the MNOs that Google Fi/Truphone uses as underlying carrier.

Google Fi is a service designed for US users. So, in the USA, Google Fi relies on 3 carriers, namely T-Mobile, Sprint and US cellular. Different sources state that the main carrier is T-Mobile. We are in the process of obtaining also a T-Mobile subscription to compare their performances. While roaming in Spain, Google Fi camps in a local network. According to retrieved information, Google Fi in Spain uses Vodafone and Movistar. We are also in the process of obtaining a SIM card from each provider, in order to compare the performance as well.

Truphone on the other hand, uses a completely different model and has agreements with local MNOs in the visited country.

Hardware selection

Regarding the hardware, Google Fi is only fully functional in a limited number of terminals and ONLY in their US version. We selected a Pixel 4a mobile phone and we acquired 4 of them in the US and then shipped 2 of them back, one to Spain and the other one to Norway.

Software, Monitoring and Data Collection

We use three apps to perform the measurements, namely Roam_Monitor, Android_Tracepath and Yomo.

Roam_Monitor is a native app developed in Android Studio specifically for this project. It allows to make measurements of Radio, end to end latency (ping), DNS Query delay and

Co-funded by the Horizon 2020 Framework Programme of the European Union

Web Performance (several metrics). The app allows us to select up to 8 different IP/hosts for IP and web Measurements, and up to 4 different hosts for DNS.

For Ping, DNS Query, Web Performance and Traceroute, we do 20 measurements for each IP and for each Carrier. Before each measurement, we assure that the cache has been cleared, and all measurements are independent. In Ping measurements, we get : DestIP, RTT, IsVPN, endTime and TTL. In DNS Measurements, we get: Address: Domain IP, Servers: Carrier DNS Servers, Domain, isVPN, QueryTime and endTime.

For web measurements, we use Javascript Timing API [https://developer.mozilla.org/en-US/docs/Web/API/Navigation_timing_API], and Javascript Paint API [https://w3c.github.io/paint-timing/#sec-PerformancePaintTiming]implemented in Webview, to measure the impact of the used network to the web performance on the user side. We used Page Load Time metric to determine how good is each carrier for an end user using a web app.

We developed Roam Monitor as an asynchronous service and configure to run every 2 seconds to capture radio and network metrics. Each data point gets exported as a JSON object and uploaded to the data server which is later used for analysis.

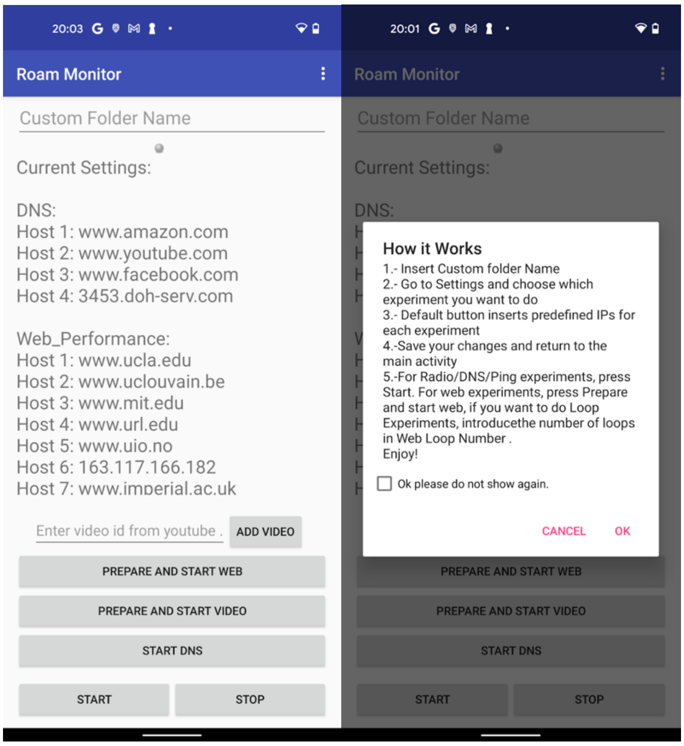

Figure 1: Look and feel of the Roam Monitor App interface

Co-funded by the Horizon 2020 Framework Programme of the European Union

AndroidTracepath is a modification of https://github.com/yanbober/android-tracepath to automatize measurement-taking. It essentially performs a defined set of traceroute tests. All measurements are stored under roaming/roam (Measurement_Type).

YoMo App (http://yomoapp.de), is an app developed by Wurzburg University that allows to rate the stream quality of youtube videos, as well as obtain different metrics of the used network. Once the measure is taken, it is uploaded to http://yomoapp.de/dashboard/ where it is possible to retrieve the lasts measurements results. It allows to measure the video quality, download throughput, stalling events, Radio access Technology, buffer level.

4.2 Traceroute measurement results in Spain.

To assess the impact on performance of the different configurations, we performed a number of traceroutes to different targets both in the EU and in the US using different providers. Specifically, we performed the tests when Fi is using T-Mobile/Vodafone and Three/Orange as MNO/visited network with and without the VPN service. Also we measured using native Vodafone, Orange and Three SIMs.

Regarding the targets, we selected simple web pages that are not served by \ac{CDN}s located in different countries. We tested ucla.edu (US), inl.info.ac.ucl.be (Belgium), uio.no (Norway), imperial.ac.uk (UK), https://www.univie.ac.at (Austria). We also tested against our web server located in Madrid (Spain). In addition, to understand the impact of CDNs, we tested 1.1.1.1 and amazon.com.

Regarding the traceroutes, we performed 20 for each target and we only kept the minimum value observed for each hop, as we are interested in measuring the fixed components of delay at this stage. We do other tests to further assess performance later on. We next present and analyze the results of the measurements obtained when the device is roaming in Spain.

In the following graphs, we depict the delays observed using traceroute with different Fi configurations (T-Mobile and Three, with and without VPN) and also directly using Three, Orange and Vodafone. For each hop, we try to determine its geographical location using the reverse DNS information, the WHOIS information and the Maxmind geolocation database.

We start with the graph plotting the delays for each hop discovered by the traceroute between a mobile device roaming in Spain and a server located in the U.S.A (ucla.edu). We can observe that that overall delay to the target is similar in all tested configurations. One major difference that we observe is the relative location of the transatlantic link in the path. When Fi/T-Mobile (with and without VPN) and when Fi/Three with VPN are used, the transatlantic link is part of the first hop. This can be observed because the delay experience in the first hop is very large (about 250ms) and also because the IP address of the first hop is located in the U.S.A.. However, when directly attaching to Orange (i.e. the visited MNO that

Co-funded by the Horizon 2020
Framework Programme of the European Union

Fi attaches to when using Three as underlying MNO) or to Vodafone (i.e. the visited MNO that Fi attaches to when using T-Mobile as underlying MNO), the first hop is located in Spain (the delay is about 40 ms and the IP of the first hop is geolocated in Spain) and the transatlantic link is observed later on in the path (it can be easily identified in the graph due to the large increase in the delay (200ms) and that the IP location moves from Spain to U.S.A.). Finally, if we look at the case of Fi/Three without VPN, we can see that the first hop is located in Europe but not in Spain and that the transatlantic jump also happens later along the path.

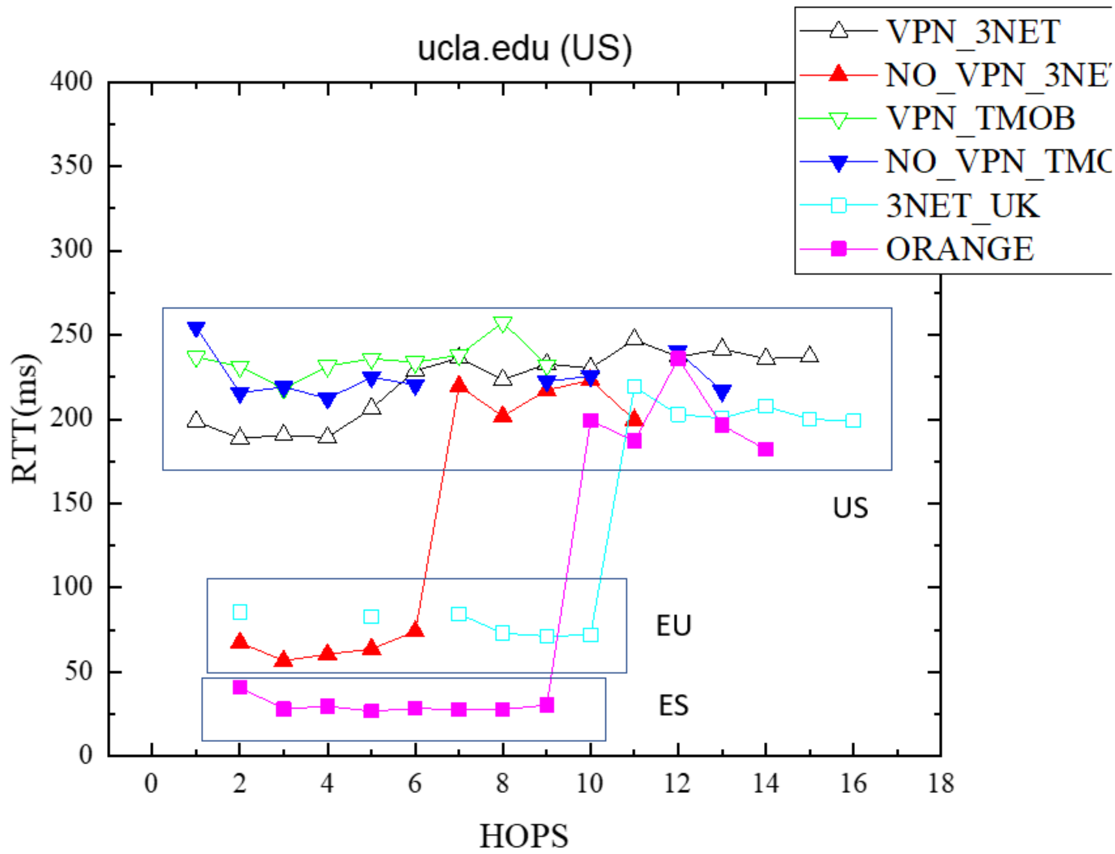

Figure 2: Traceroute from Spain to a server in the USA with different configurations.

The observed behaviour can be explained by two factors, namely Home Routing and the location of the VPN endpoint. When the VPN is enabled, the first hop is the other VPN tunnel endpoint. From the experiments, we conclude that it is located in Google in the U.S.A. So when the VPN is enabled, the traffic is routed to Google in the U.S.A. and then towards the Internet through Google.

When the VPN is not enabled, we observe the effect of HR. When Fi/T-Mobile is used, the traffic is first routed to T-Mobile in the U.S.A and it the exits to the Internet through T-Mobile's network in the U.S.A. When Fi/Three is used and when Three is directly used, the traffic is first routed to Three network (either U.K. or Austria) and then to the Internet

Co-funded by the Horizon 2020 Framework Programme of the European Union

through Three's network. In this case, because the server the Fi device is communicating with is located in the U.S.A. (ucla.edu), none of this has a significant impact in the experienced overall delay.

We next look into the traceroute results towards a server located in Belgium (inl.info.ac.ucl.be). In this case, we do observe significant difference in the overall delays. We observe a first group that includes both Fi/Three and Fi/T-Mobile with VPN plus T-Mobile without VPN that has a significantly larger overall delay (200 ms extra) than a second group that includes Orange, Fi/Three, Three (U.K. and Austria), The reason for this increase can also be observed in the graph. While the first group uses panoramic routing back to the U.S.A., either due to VPN or to HR for traffic flowing from Spain to Belgium, in the second group the traffic never leaves Europe, taking a much shorter path.

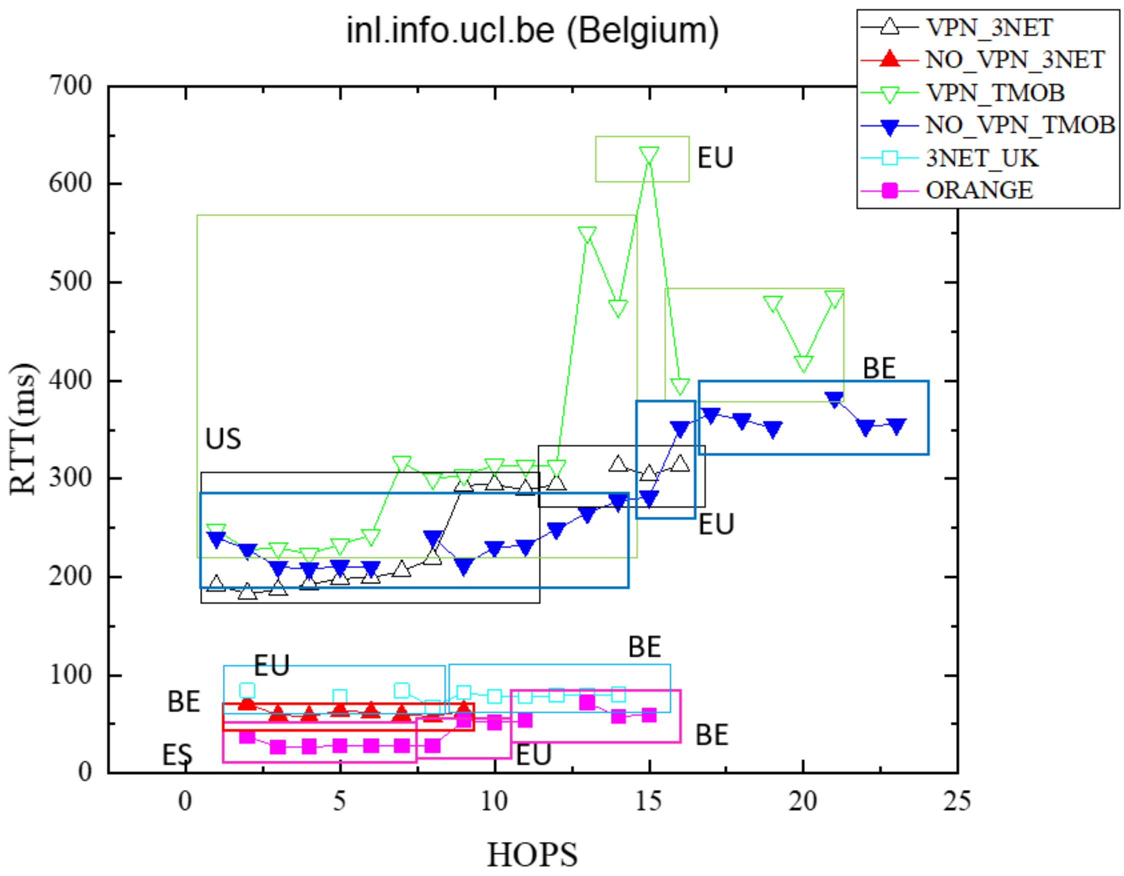

Figure 3: Traceroute from Spain to a server in Belgium with different configurations.

We next plot the delays observed in the traceroute experiments between a mobile in Spain and a server also located in Spain. If we look at the total delay experienced, we can distinguish that the smaller overall delay is observed when using Orange. The second smaller overall delay is observed when using Fi/Three without VPN and Three, while the other cases

Co-funded by the Horizon 2020 Framework Programme of the European Union

(Fi/T-Mobile with and without VPN and Fi/Three with VPN) suffer an exceedingly larger delay (with Fi/T-Mobile with VPN being the largest with a great difference).

Similarly to the previous case this can be explained by a combination of HR and VPN tunnelling, only that in this case, the situation is exacerbated because both the mobile device and the server are located in the same visited country, so the panoramic routing either to the U.S.A. or the U.K. injects the extra delay towards the foreign network twice (to go and to come back to Spain), as it can be observed by the two large leaps in the picture (one for the first hop and another one later in the path).

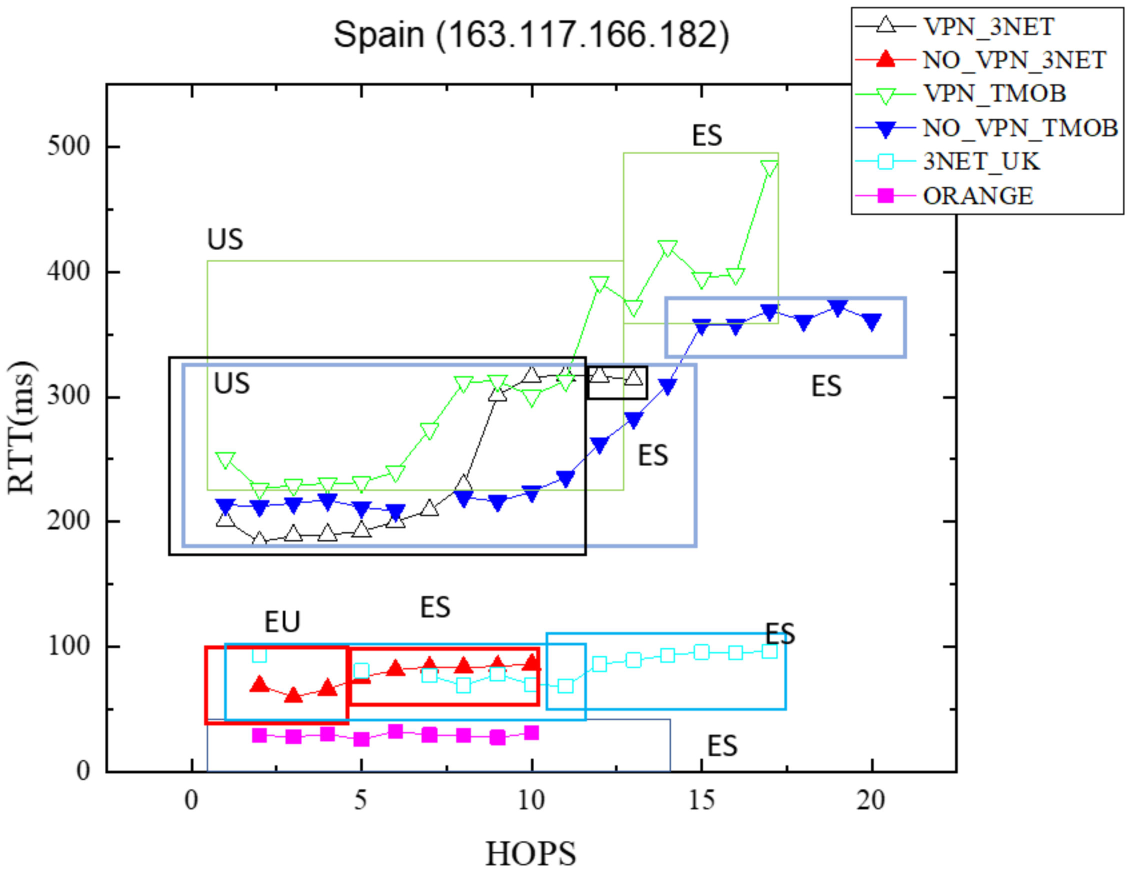

Figure 4: Traceroute from Spain to a server in Spain with different configurations.

We next take a look to the case where the mobile device is communicating with a CDN. As expected, we observe that the shortest delay is experienced by Orange, followed by Fi/Three without VPN and Three and then followed by Fi/Three with VPN and Fi/T-Mobile with and without VPN. In this case, the replica located close to the point where the traffic is injected to the Internet is used to retrieve the content. This implies that when the VPN and when Fi/T-Mobile is used, the mobile is accessing a CDN replica located in the U.S.A. while when Fi/Three without VPN and Three is used the accessed replica is located in Europe (but not in

Co-funded by the Horizon 2020
Framework Programme of the European Union

Spain), while when Orange is used a replica in Spain is used, resulting in a shorter overall delay.

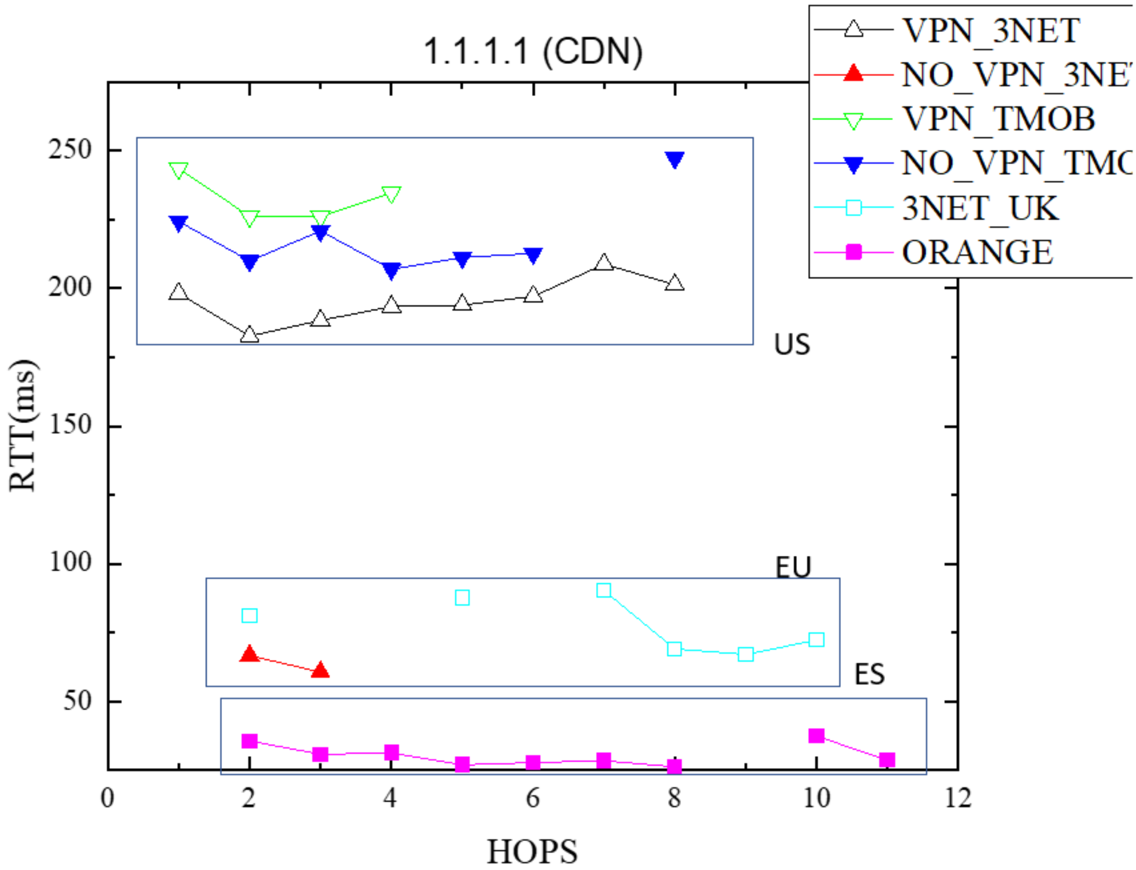

Figure 5: Traceroute from Spain to a server in a CDN with different configurations.

4.3 DNS resolution delay in Spain

We measure the DNS resolution delay using the different configurations, namely, Fi using Three with and without VPN, Fi using T-Mobile with and without VPN, Three and Orange. We next present and analyze the results of the measurements obtained when the device is roaming in Spain.

We first measure the resolution delay for a non-cached name. To do that, we set our own authoritative domain and we configure it with a wildcard DNS record, so it responds to all names under that domain. The we perform queries for unique names under our domain, ensuring that the name queried is never repeated, ensuring that caching is not involved in

Co-funded by the Horizon 2020
Framework Programme of the European Union

the resolution. The authoritative server is located in Spain, serving a worst case for roaming devices as we inferred from the traceroute experiments. We did 20 queries with each configuration.

In the next figure we depict the obtained delays while querying a non-cached name. As expected, both VPN enabled configurations and Fi over T-Mobile exhibit an overall larger delay, followed by Three and Fi over Three (without VPN) and the (native) Orange with the smallest delay. We observe that Fi over Three with the VPN enabled experiences a surprisingly long delay (the mean delay is over 550 ms) and also a large variation. Both T-Mobile configurations (with and without VPN) follow, with a mean delay above 350ms. All the Three configurations (without VPN) have a mean delay close to 150 ms and the native Orange a mean close to 50 ms. So, the Home Routing back to the U.S.A. imposes a penalty of 300-500ms compared to a local breakout (600\% - 1,000\% penalty) while the regional breakout achieved by using Three imposes extra 100 ms (200\% penalty).

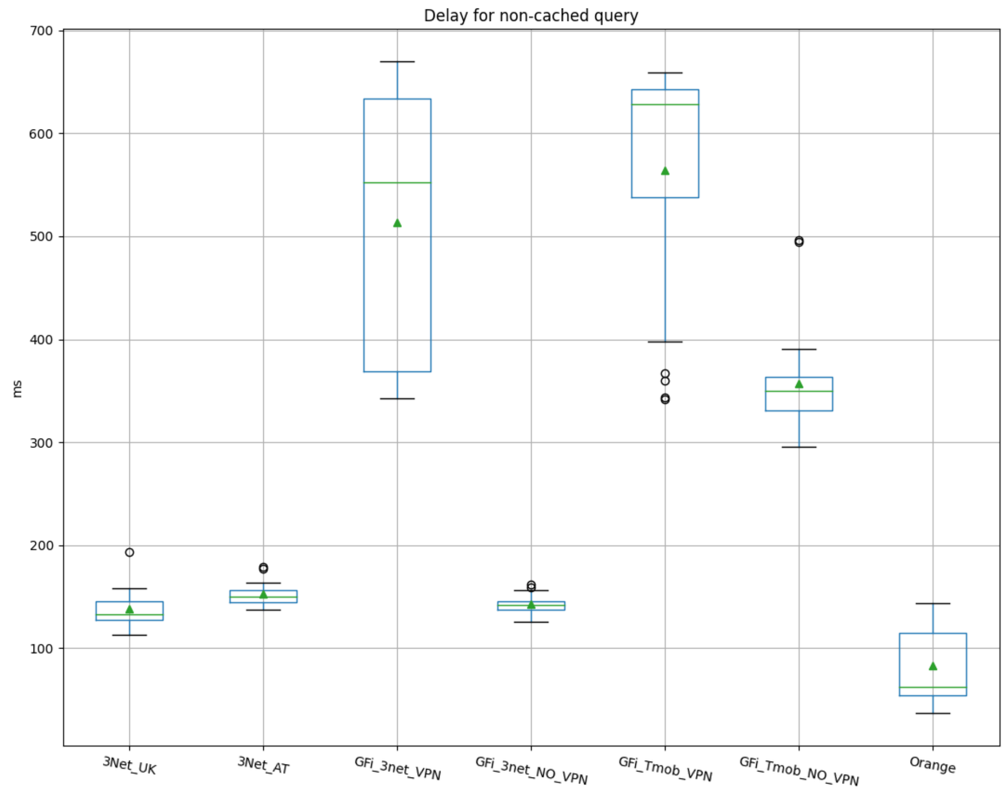

Figure 6: DNS resolution delay for a non-cached name using the different configurations.

We next measure the delay for the resolution for a query that is present in the resolver's cache. To do that, we make a first query for a domain name in order to populate the resolver's cache, we next clear the DNS client cache (to force the client to query the resolver

Co-funded by the Horizon 2020 Framework Programme of the European Union

again), we query for the same domain name and measure the resolution time. We do this 20 times for three popular names (www.amazon.com, www.facebook.com and www.youtube.com).

In the next figure we present the box plots for the resolution delay for a cached name using the different configurations. Overall, we see the same trends as in the case of a non-cached name, with three groups i.e., a first group with the greatest delay formed by both T-Mobile configurations and Three with VPN, a group in the middle consisting on Three and Fi\& Three without VPN and (native) Orange with the smallest delay. The absolute values though as well as the absolute differences are smaller than in the case of the non-cached delay. Indeed, for both configuration of Fi with T-Mobile and for Fi/Three with VPN, the delays are in the order of 250 ms, while with Fi/Three without VPN as well as both Three native configurations, the delays are in order of 100ms, while for Orange it drops to 50ms. This means that the penalty for using Home Routing back to the U.S.A. is 400\% while that for the regional breakout the penalty is 100\% compared to the potential local breakout.

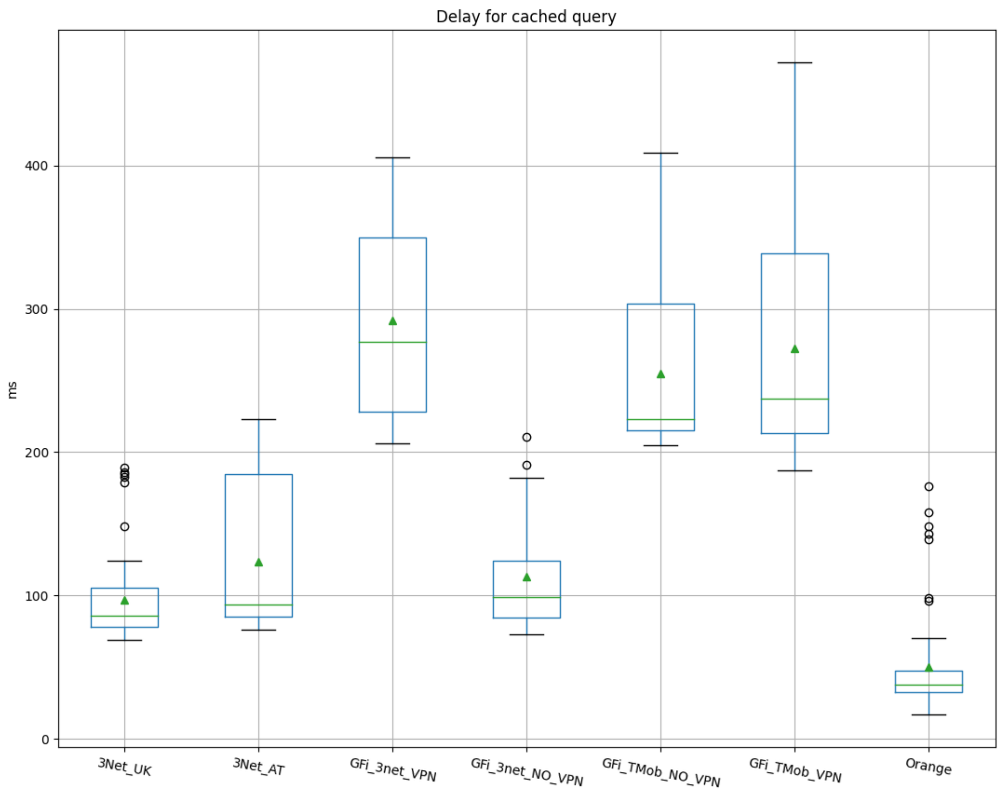

Figure 7: DNS resolution delay for a cached name using the different configurations.

Co-funded by the Horizon 2020
Framework Programme of the European Union

4.4 Web performance measurements in Spain

We next measure the impact on web performance of the delay introduced by the different configurations.

To characterize web performance, we use the Page Load Time (PLT) metric available in the Navigation Timing API in the browser. The PLT is the time it takes for a page to load. We calculate the PLT from initiation (when click on a page link or type in a Web address) to completion (when the page is fully loaded in the browser). essentially, this is the time it takes for the last object in the page to download.

It occurs when all the HTML files and any sub-resources (images, fonts, css, videos, etc.) are loaded. Note that not all these elements are needed to complete the rendering of the visible portion of the page.

Though many other metrics focus on different aspect of webpage performance, recent studies showed that the PLT is good enough to capture the experience of the users in various radio contexts, showing similar patterns to First Paint or Speed Index. Thus, in this paper we focus on the PLT.

We measure the PLT using the different configurations (Fi using Three with and without VPN, Fi using T-Mobile with and without VPN, Three and Orange) and with different web pages. We selected web pages of roughly the same size (with a 5% range) in different geographic locations, namely, www.ucla.edu in the U.S.A, www.uclouvain.be in Belgium, www.url.edu in Spain and www.mit.edu which is served by a CDN. For each configuration and for each web page, we do 20 experiments. Results are depicted in the next figures. We next present and analyze the results of the measurements obtained when the device is roaming in Spain.

Co-funded by the Horizon 2020
Framework Programme of the European Union

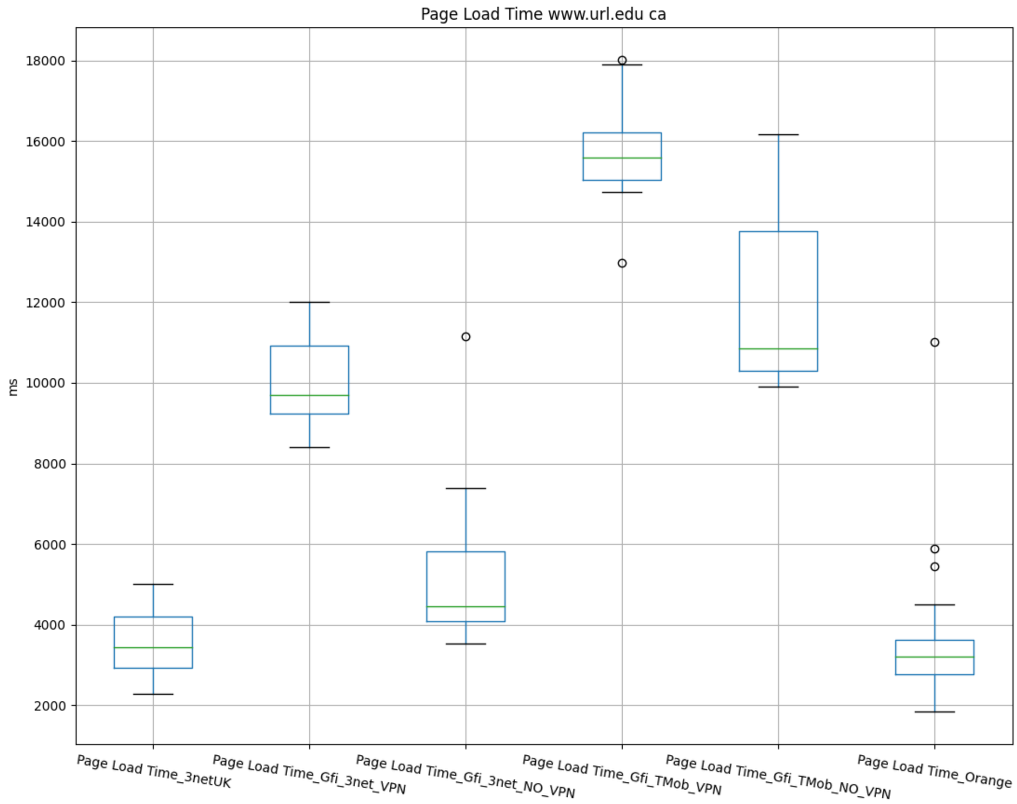

Figure 8: Web performance for a server in Spain using the different configurations.

Co-funded by the Horizon 2020
Framework Programme of the European Union

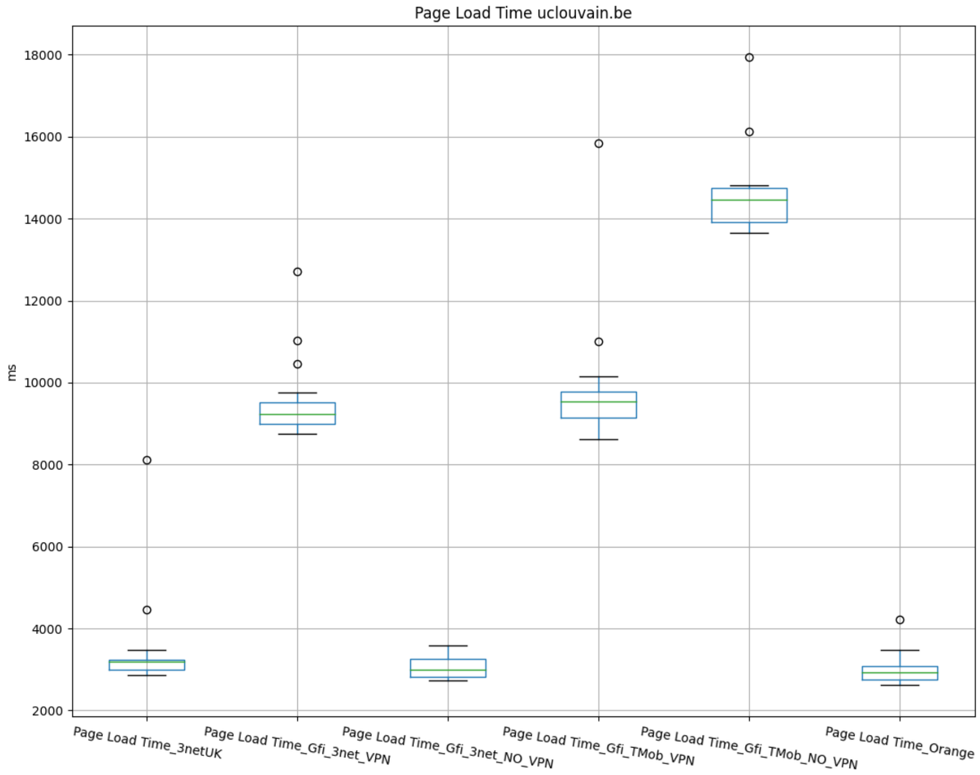

Figure 9: Web performance for a server in Belgium using the different configurations.

Co-funded by the Horizon 2020
Framework Programme of the European Union

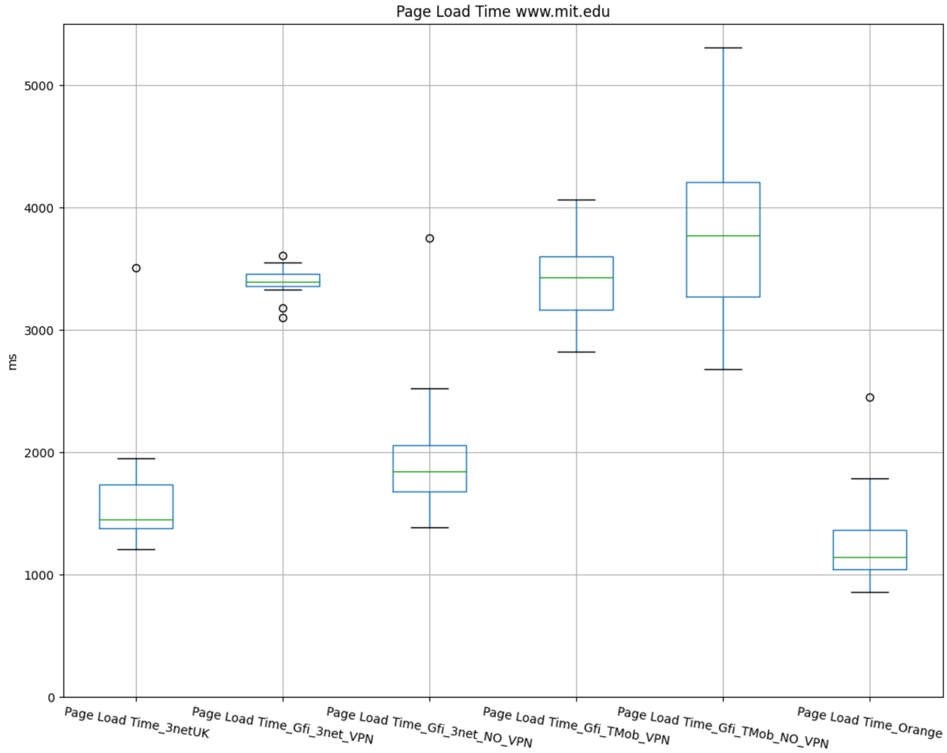

Figure 10: Web performance for CDN hosted content using the different configurations.

Co-funded by the Horizon 2020 Framework Programme of the European Union

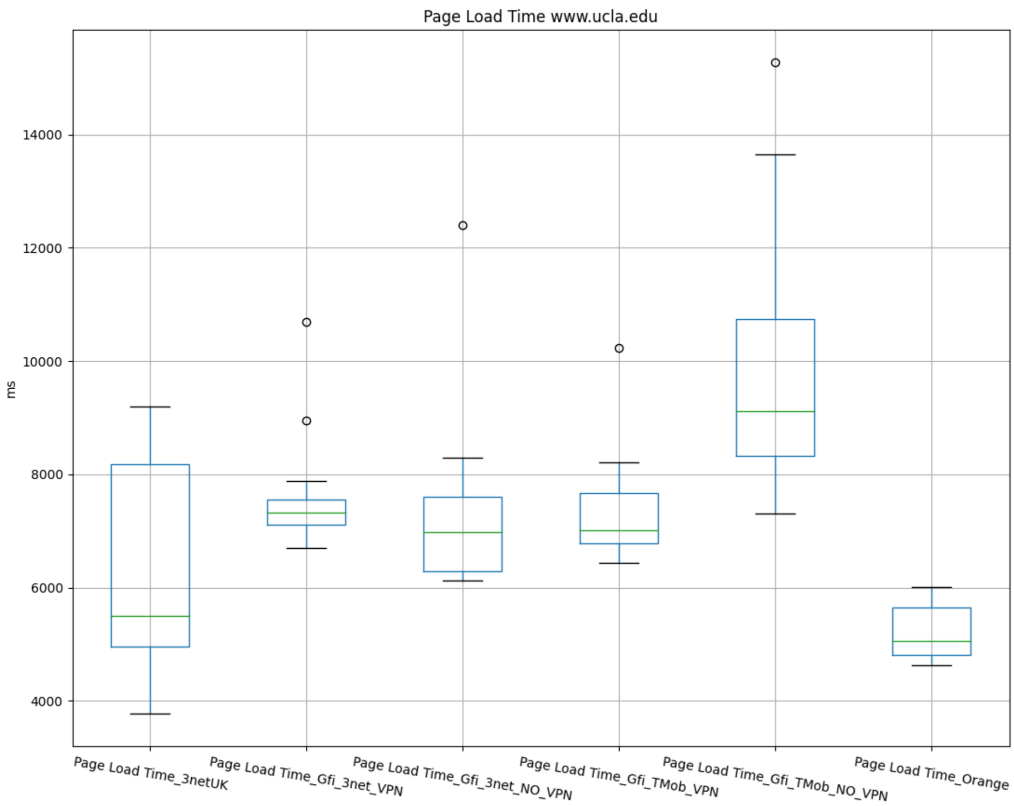

Figure 11: Web performance for a server in the USA using the different configurations.

We can observe similar trends as in the cases of DNS latency and traceroute. The differences are larger when the server is in Spain, followed (closely) by the case of the server in Belgium, then the server in the U.S.A. where differences are significantly smaller. Regarding the content served by the CDN, the overall PLT is smaller in all the cases, but the differences remain. In all cases, the smallest delay is achieved with the native service offered by Orange, then a second tier that includes Fi over Three without VPN (which is very similar to Three) and then finally, a third tier with the longest delay with both VPN enabled configuration plus T-Mobile. We would like to highlight that serving content through a CDN, while reducing the delay for all configuration, it does not eliminate the differences in the PLT for the different configuration. This is so because when a CDN is used, the content is retrieved from the replica that is closest to Internet access point of the mobile. When HR or a VPN is used, the mobile accesses to a replica in the U.S.A. while when using Three, and Fi/Three, the mobile connects through Belgium/U.K. and access to a replica in Europe while connective native, the mobile connects to a replica that is likely to be in Spain.

Co-funded by the Horizon 2020 Framework Programme of the European Union

4.5 Video performance measurements in Spain.

We next present the results of the video measurements performed using the Yomo App. We measure the video performance using the different configurations, namely, Fi using Three with and without VPN, Fi using T-Mobile with and without VPN, Three and Orange. We played a short video (1 min) from youtube using the different configurations. We next present and analyze the results of the measurements obtained when the device is roaming in Spain.

In the figure we show the time spent in each quality for the different configurations.

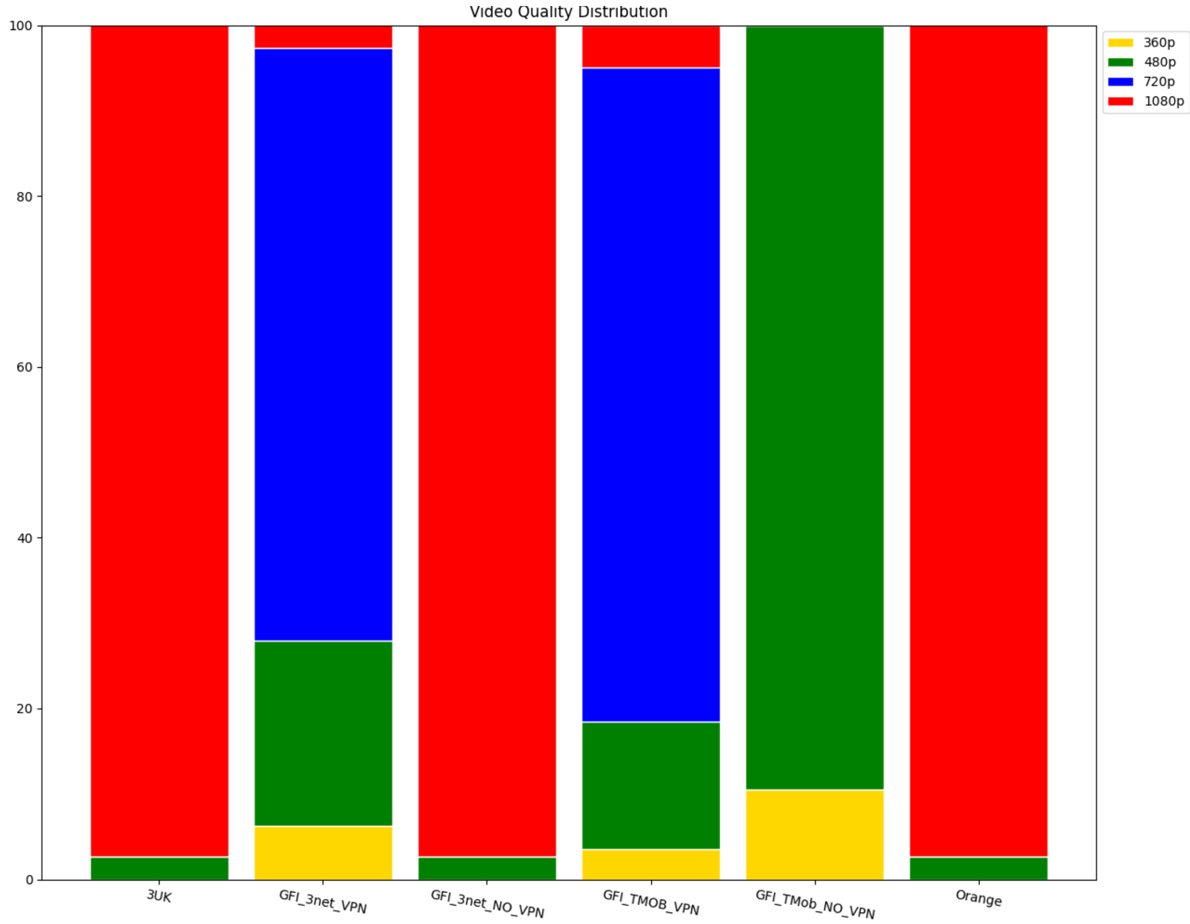

Figure 12: Video quality for the different configurations.

We can observe that for Three UK, Fi with Three (without VPN) and Orange, most of the video is played in the higher resolution (1080p) while in Fi with Three with VPN and for Fi with T-Mobile with VPN, the most common quality is 720p while for Fi with T-Mobile without VPN the most common quality is 480p.

We also measured the bandwidth achieved in each configuration while downloading the video. We depict it next.

Co-funded by the Horizon 2020
Framework Programme of the European Union

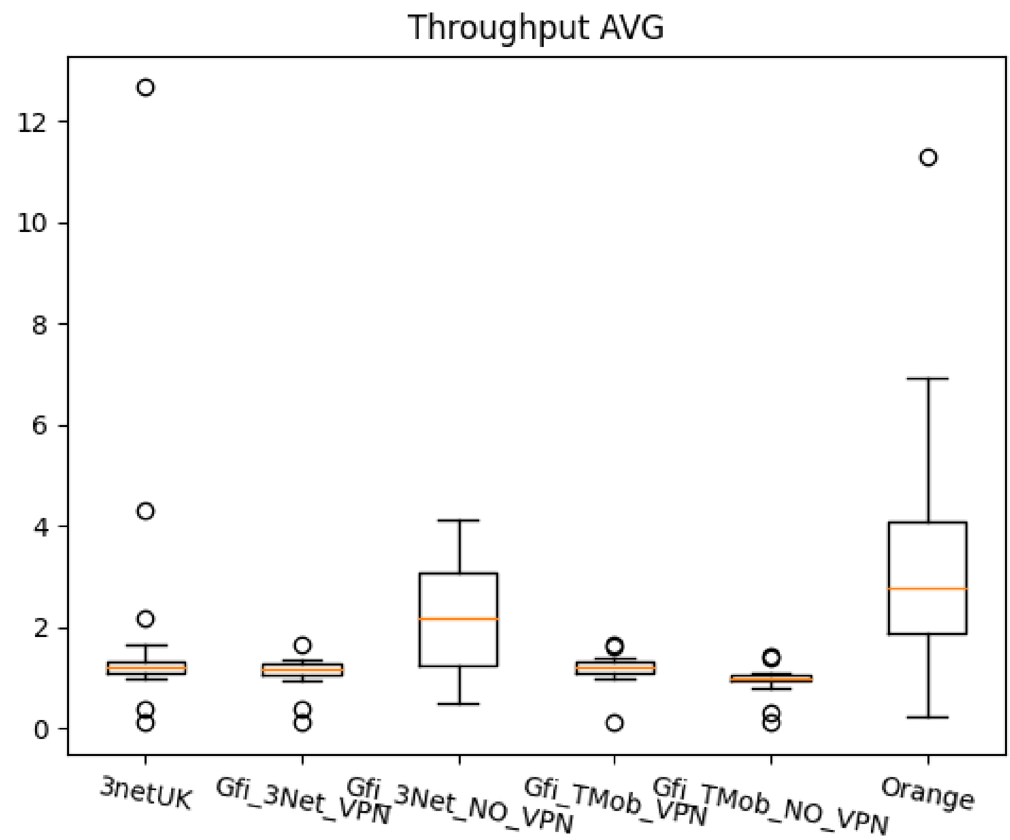

Figure 13: Average throughput while downloading video using the different configurations.

4.6 Traceroute measurement results for Norway.

We next present the traceroute measurements when the mobile device is roaming in Norway. We measured using the different configurations, including all four Fi configurations (using Three and T-Mobile, with and without VPN), as well as Telenor and Truphone. We depict them next.

Co-funded by the Horizon 2020 Framework Programme of the European Union

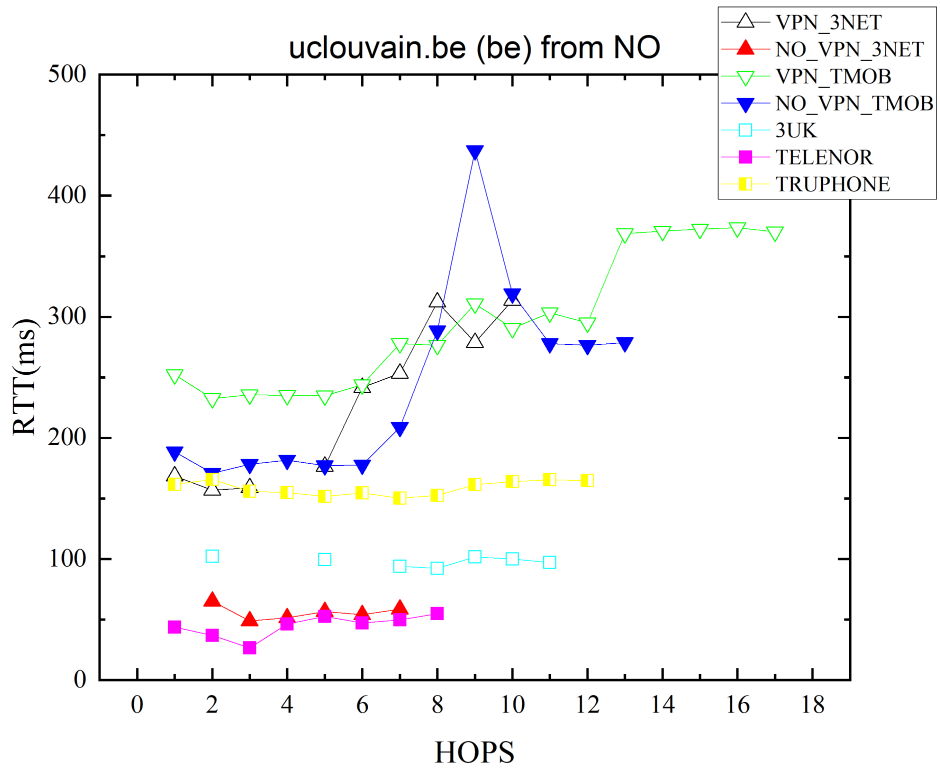

Figure 14: Traceroute from Norway to a server in the Belgium with different configurations.

Co-funded by the Horizon 2020 Framework Programme of the European Union

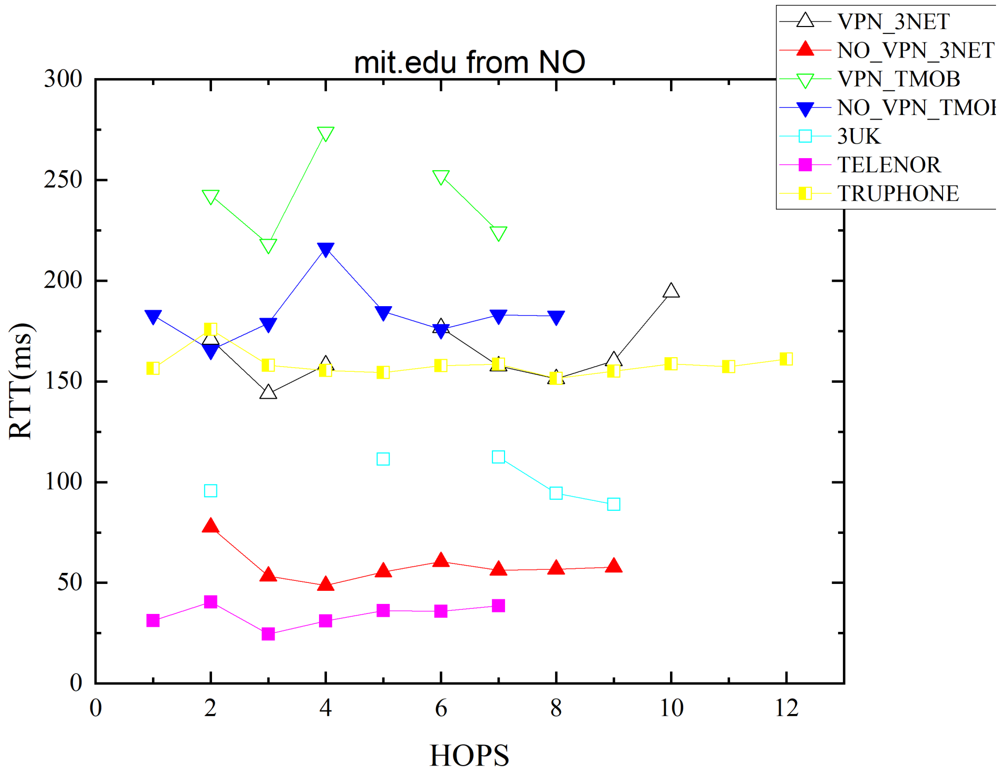

Figure 15: Traceroute from Norway to a server in a CDN with different configurations.

Co-funded by the Horizon 2020
Framework Programme of the European Union

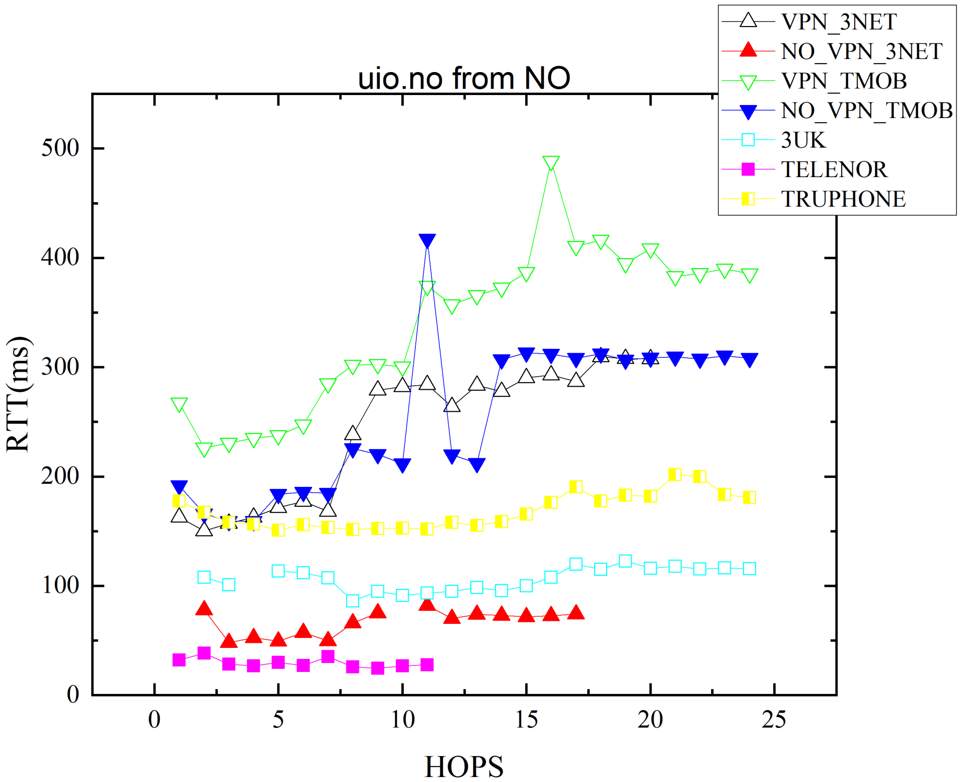

Figure 16: Traceroute from Norway to a server in the Norway with different configurations.

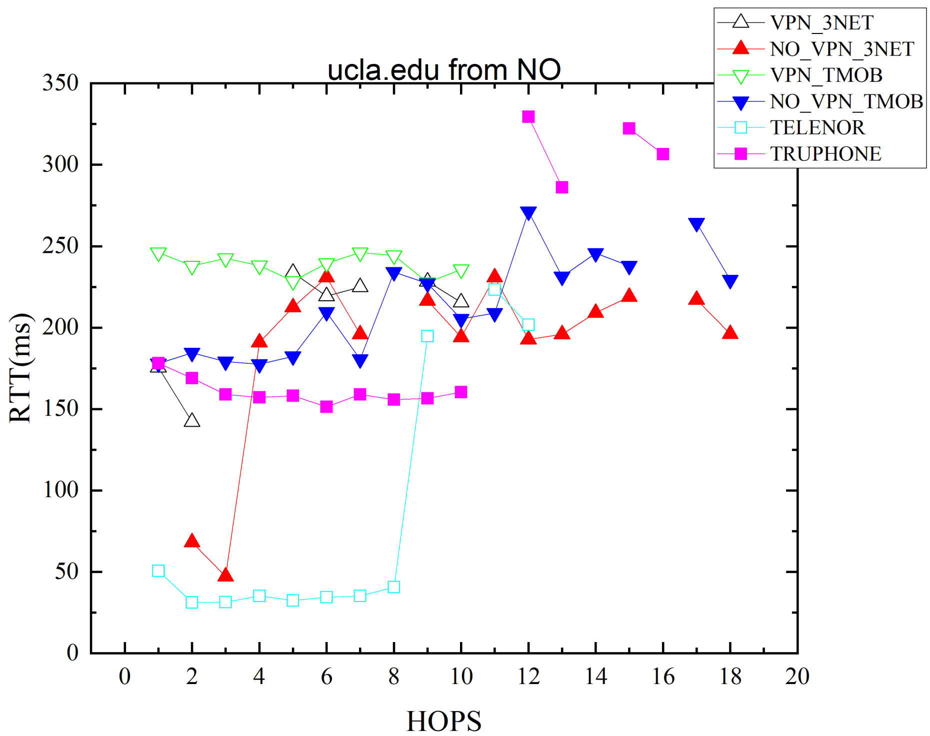

Figure 17: Traceroute from Norway to a server in the USA with different configurations.

Co-funded by the Horizon 2020 Framework Programme of the European Union

Similarly as in Spain, we observe notable differences when the accessed server is either in Norway, in the EU or in a CDN, while the performance is similar when the server is located in the USA. The Home routed configurations and the VPN configurations perform significantly poorer than the regional breakout or local breakout.

4.7 DNS resolution delay in Norway.

We next present the results for the DNS resolution delay when the device is roaming in Norway. We measure with the same configuration used for the traceroute measurements.

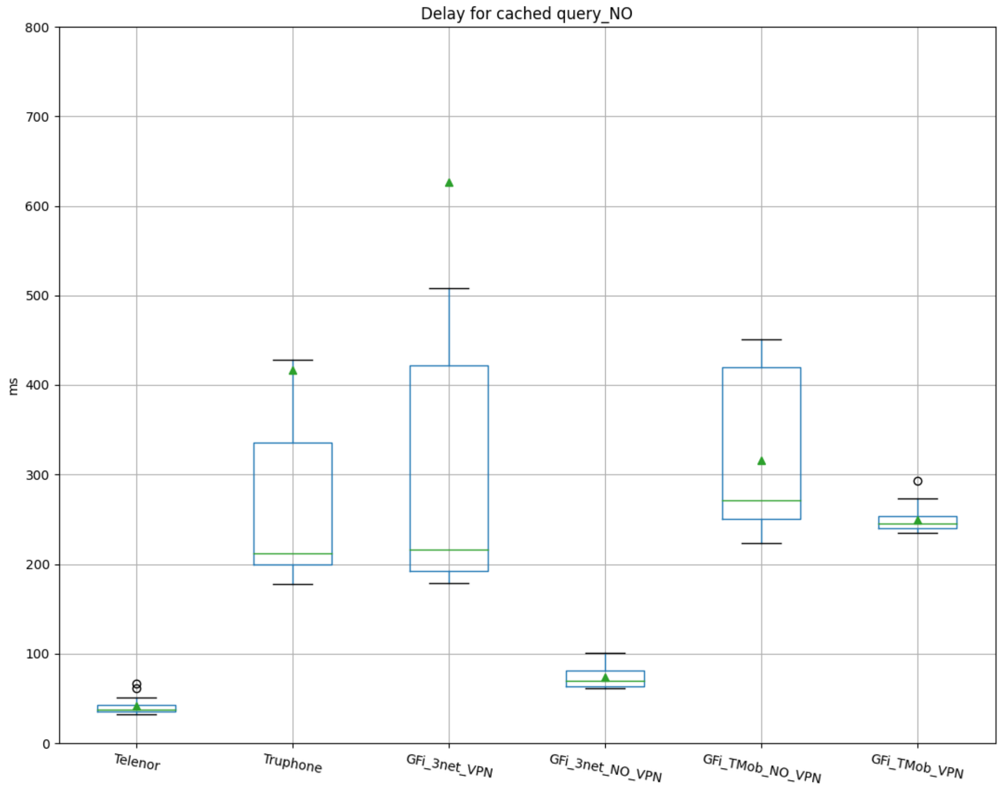

Figure 18: DNS resolution delay from Norway for a cached name using the different configurations.

Similarly as in Spain, we observe notable differences between configurations. The Home routed configurations and the VPN configurations perform significantly poorer than the regional breakout or local breakout.

Co-funded by the Horizon 2020 Framework Programme of the European Union

4.8 Web performance measurements from Norway.

We next present the results for the web performance when the device is roaming in Norway using the different configurations for Fi and Truphone.

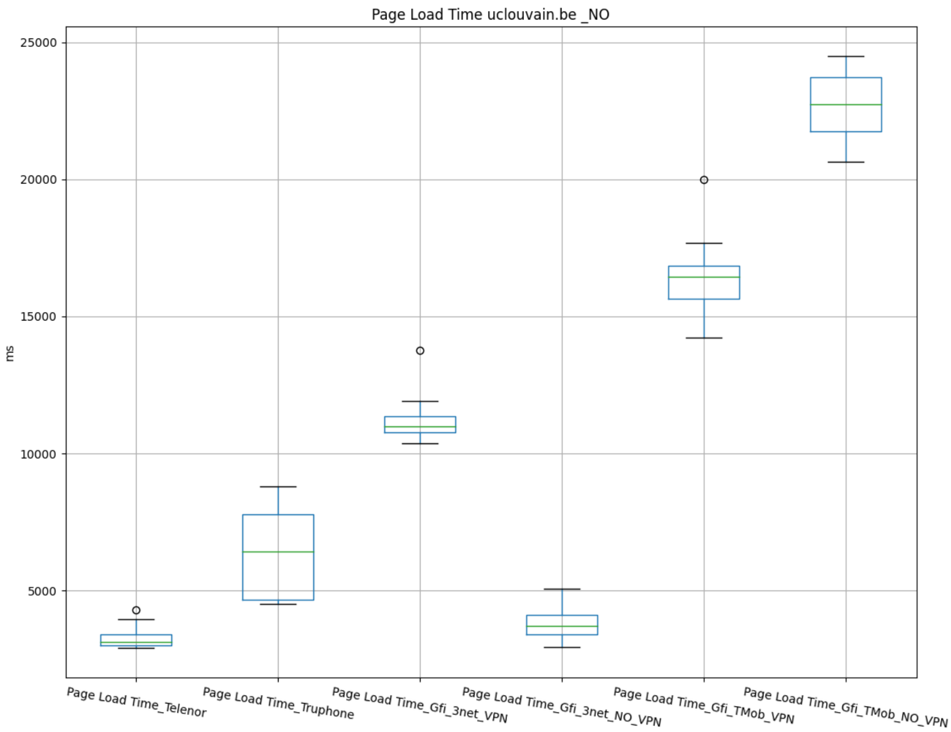

Figure 19: Web performance for a server in Belgium using the different configurations.

Co-funded by the Horizon 2020 Framework Programme of the European Union

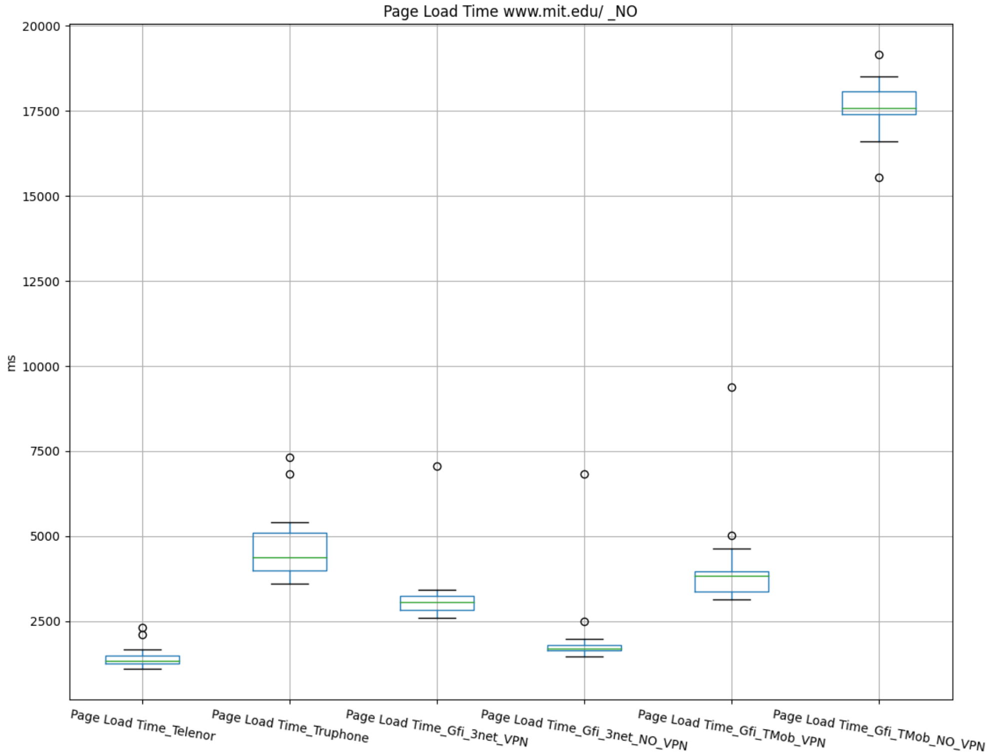

Figure 20: Web performance for a server in a CDN using the different configurations.

Co-funded by the Horizon 2020
Framework Programme of the European Union

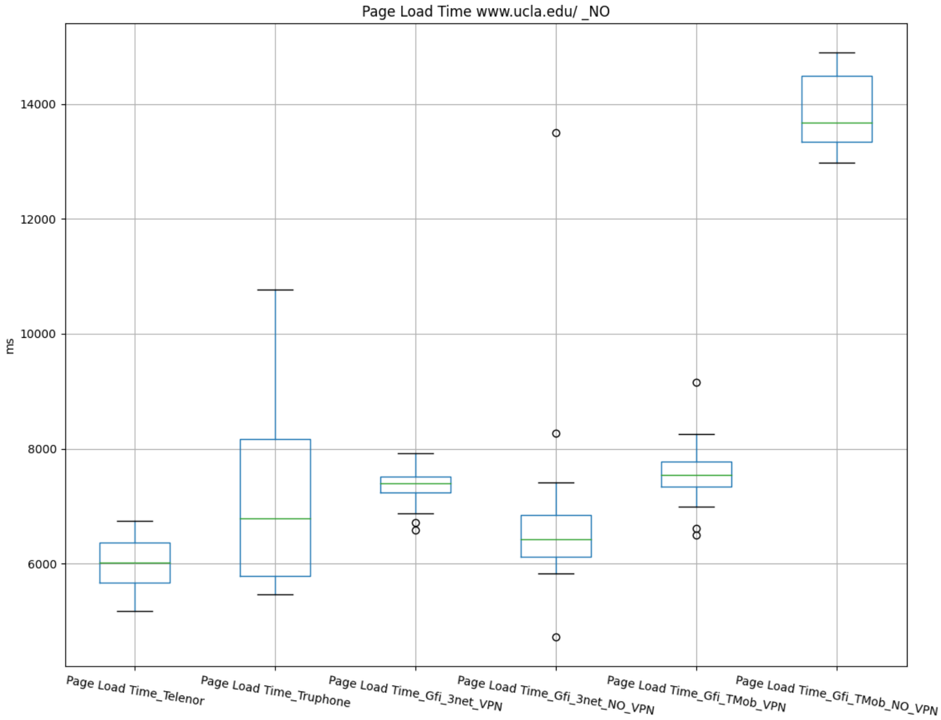

Figure 21: Web performance for a server in USA using the different configurations.

Co-funded by the Horizon 2020
Framework Programme of the European Union

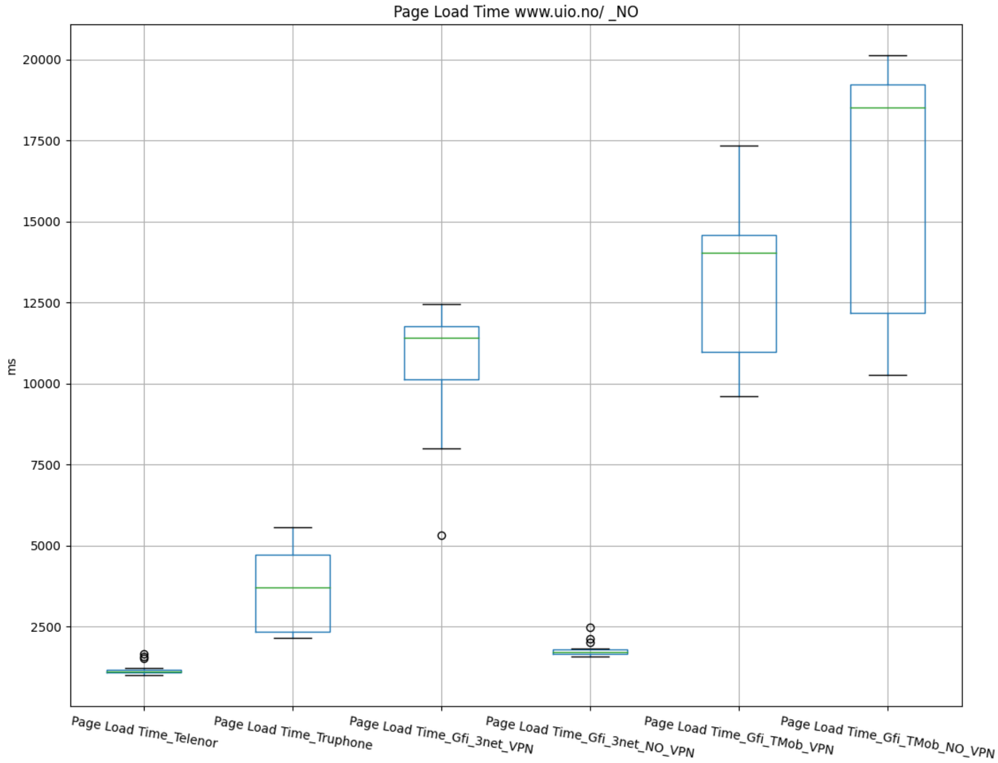

Figure 22: Web performance for a server in Norway using the different configurations.

Similarly to the results in Spain, we observe notable differences when the accessed server is either in Norway, in the EU or in a CDN, while the performance is similar when the server is located in the USA. The Home routed configurations and the VPN configurations perform significantly poorer than the regional breakout or local breakout.

5 Discussion and Analysis on Results

Analysis and discussion of the traceroute measurements results: MNOs rely on HR for roaming devices and this configuration results in an added latency penalty for the roaming devices, especially when the other end of the communication is located topologically close to the current location of the roaming device. For the delay comparisons, we use as a baseline the delay observed by Orange/Vodafone, which is the delay obtained is directly attaching to the local network.

In the case of Fi, when T-Mobile (without VPN) is used, the traffic is home routed back to the U.S.A. resulting in large penalties while roaming in Spain, notably when content locally available in Spain/Europe is accessed. Even in the case when content stored in a CDN is

Co-funded by the Horizon 2020
Framework Programme of the European Union

accessed, a large latency penalty is also observed. The only case when no delay is added is when content in the U.S.A. is accessed. However, using T-Mobile is not the default policy of Fi while roaming in Europe.

When in Europe, Fi relies on Three. When using Fi/Three without VPN the latency penalty, while still exists in some cases, is largely reduced. This is because even if Three uses Home routing the distance to "home" (i.e. U.K. or Austria) is largely reduced compared to the case of home routing back to the U.S.A. This approach is a middle-ground between Local break-out and Home Routing and significantly reduces latency with a small overhead in terms of roaming agreements (only one extra agreement for a whole region). this can be seen as a "regional breakout model" similar to the IPX breakout model.

However, the benefits obtained by the "regional breakout" are lost when the VPN is enabled (which is the default behaviour for Fi), because the other endpoint of the tunnel is located in the U.S.A. A simple improvement that would result in large latency reductions would be to deploy regional VPN tunnel endpoints that terminate the VPN within the region.

Google Fi in Norway performs similarly to Google Fi in Spain.

Truphone in Spain performs local breakout and achieves the better performance (represented by the measurements of the local provider Orange). Truphone in Norway on the other hand, does not perform local breakout, since it exits through Netherlands, performing a regional breakout, similar to the one of Fi with Three without VPN (and exhibiting similar performance).

DNS measurements: Overall, we see the same trends as in the case of the traceroute results, with three groups i.e., a first group with the greatest delay formed by both T-Mobile configurations and Three with VPN, a group in the middle consisting on Three and Fi/Three without VPN and (native) Orange with the smallest delay. In the case of Norway, we find similar results as in Spain. The main difference is regarding Truphone, that while in Spain it performs local breakout and can be assimilated to the local MNO, in the case of Norway it performs regional breakout and performs similarly to Three.

Analysis and discussion of Web performance results: We analyze the correlation of the web performance and the latency metrics. In order to do that, we first measure the RTT to each of the servers using the different configurations. Then, for each server we compute the difference between the PLT/RTT for Orange (the smallest value) and the rest of the configurations and we normalize the obtained difference with the value for Orange. We plot the results below. In Figure 23 following, we can see for instance, that in the case of the web page located in Spain, when using Fi/Three with the VPN enabled, an increase of 5,25 times in the RTT results in an increase of 3 time in the PLT. Overall, we observe that both metrics are directly related and that an increase in the RTT lineally affects the PLT. We observe that

Co-funded by the Horizon 2020
Framework Programme of the European Union

the HR and the VPN enabled configurations impose a higher penalty both in terms of RTT and PLT. We also observe that the relative penalty is smaller for the content in the U.S.A. than for the other locations (which seems natural, since the target is close to the Internet exit point).

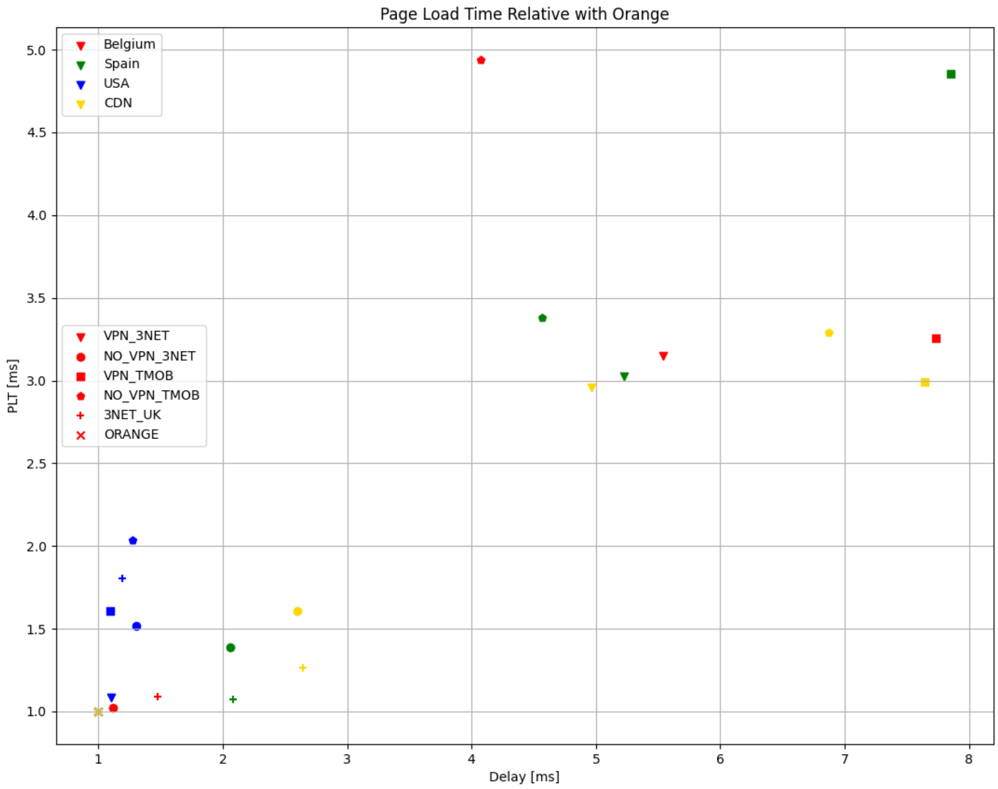

Figure 23: Correlation between PLT and latency for the different configurations.

Regarding Norway, we observe similar results as in Spain for Fi. In the case of Truphone, it exhibits similar results to Fi using Three without VPN rather than Telenor, as in the other metrics.

Video performance considerations: We have observed that the different configurations have a significant impact on video performance. This is somehow unexpected since an additional latency may affect the time to start playing but not necessarily the quality of the video. We measured the available bandwidth during video streaming and we observe that indeed the home routed/VPN configurations exhibit a lower bandwidth, explaining the reduced quality of the video.

Co-funded by the Horizon 2020 Framework Programme of the European Union

6 Present and Foreseen TRL

The project only produced measurements and analysis of the results. The only piece of code produced is the Roam Monitor App, which purpose is to perform the different DNS, web and traceroute measurements. The TRL for the Roam Monitor App is TRL7.

7 Exploitation, Dissemination and Communication Status

The planned dissemination plan includes the publication of the obtained results in a top conference or journal. At this moment, we are considering ACM Mobicom or WWW as candidate venues.
Also, we plan to make all the project results public, including the implementation of the tests, the Roam Monitor App and the dataset obtained from the experiments (in an anonymized form, to preserve privacy).

8 Impacts

The project achieved the following impacts in relation to the NGI initiative, as detailed next:

Impact 1: Enhanced EU – US cooperation in Next Generation Internet, including policy cooperation.: The project nurtured the existent relationships between the US and the EU partners of the project consortium. The experiments also involve Ozgu Olay from Oslo University. We are also in contact with researchers form UK and Italy on this matter. On the US side, we reached out to people from Google Fi (with not so much success) and also to the people who developed Mobile Insight, building additional collaborations.

Impact 2: Reinforced collaboration and increased synergies between the Next Generation Internet and the Tomorrow's Internet programmes.

Impact 3: Developing interoperable solutions and joint demonstrators, contributions to standards. We have developed the Roam Monitor app, which is available in GitHub and was be used by all partners in the project. The Roam Monitor App will become publicly available once the paper presenting the project results is published.

Impact 4: An EU - US ecosystem of top researchers, hi-tech start-ups / SMEs and Internet-related communities collaborating on the evolution of the

Internet: We are working in a joint publication between the US and the EU partners of the project and also including University of Oslo. We expect that we will continue with further collaborations in the future.

Co-funded by the Horizon 2020
Framework Programme of the European Union

9 Conclusion and Future Work

The mobile communication landscape is rapidly evolving with the emergence of new models of global operators. Such emerging models are challenging the current operation models. This is particularly true in the case of roaming. These new operators have a global vision and aim to provide service irrespectively of the user location.

In this work we compare the performance of two of these new global operators, Google Fi and Truphone while roaming and compare it with the incumbent operators. We find that these two operators have different operation while roaming. While Google Fi offers a regional breakout, allowing traffic to exit within the region, Truphone offers local breakout in some countries, allowing a closer exit point. Both cases contrast with the operation model of traditional operators that rely on Home routing to exit traffic.

This results in better performance for all the metrics we tested, including latency, DNS query resolution delay, web performance and video performance. The differences are significant, and we conclude that they have real impact on user experience. We believe that this will have a significant impact on the market for roaming services, and it foreseeable that the new models are adopted by incumbent operators to remain competitive.

As future work, we plan to extend our analysis to new global operators that are targeting the IoT market such as Twilio.

Co-funded by the Horizon 2020
Framework Programme of the European Union

10 References (optional)

[Mandalari18] Anna Maria Mandalari, Andra Lutu, Ana Custura, Ali Safari Khatouni, Özgü Alay, Marcelo Bagnulo, Vaibhav Bajpai, Anna Brunstrom, Jörg Ott, Marco Mellia, and Gorry Fairhurst. 2018. Experience: Implications of Roaming in Europe. In Proceedings of the 24th Annual International Conference on Mobile Computing and Networking (MobiCom '18). Association for Computing Machinery, New York, NY, USA, 179–189. DOI:https://doi.org/10.1145/3241539.3241577

11 Glossary

HR	Home Routing
MCCA	Multi-Carrier Cellular Aggregator
MNO	Mobile Network Operator
MVNO	Mobile Virtual Network Operator
NGI	Next Generation Internet
NWU	North Western University
PLT	Page Load Time
R&D	Research and Development
RTT	Round Trip Time
UC3M	Universidad Carlos III de Madrid (coordinating partner)
VPN	Virtual Private Network

Co-funded by the Horizon 2020
Framework Programme of the European Union