Global Mobile Network Aggregators: Taxonomy, Roaming Performance and Optimization

Sergi Alcalá-Marín
IMDEA Networks Institute
Universidad Carlos III de Madrid
Madrid, Spain
sergi.alcala@imdea.org

Aravindh Raman
Telefonica Research
Barcelona, Spain
aravindh.raman@telefonica.com

Weili Wu
Northwestern University
Evanston, Illinois
weiliwu2023@u.northwestern.edu

Andra Lutu
Telefonica Research
Madrid, Spain
andra.lutu@telefonica.com

Marcelo Bagnulo
University Carlos III of Madrid
Leganes, Spain
marcelo@it.uc3m.es

Ozgu Alay
University of Oslo
Karlstad University
Oslo, Norway
Karlstad,Sweden
ozgua@ifi.uio.no

Fabián Bustamante
Northwestern University
Evanston, Illinois
fabianb@cs.northwestern.edu

ABSTRACT

A new model of global virtual Mobile Network Operator (MNO) – the Mobile Network Aggregator (MNA) – has recently been gaining significant traction. MNAs provide mobile communications services to their customers by leveraging multiple MNOs, and connecting through the one that best match their customers' needs at any point in time (and space). MNAs naturally provide optimized global coverage by connecting through local MNOs across the different geographic regions they provide service. In this paper, we dissect the operations of three MNAs, namely, Google Fi, Twilio and Truphone. We perform measurements using the three selected MNAs to assess their performance for three major applications, namely, DNS, web browsing and video streaming. We benchmark their performance comparing it to the one of a traditional MNO. We find that even MNAs provide some delay penalty compared to the service accessed through the local MNOs in the geographic area where the user is roaming, they can significantly improve performance compared to traditional roaming model of the MNOs (e.g. home routed roaming). Finally, in order to fully quantify the potential benefits that can be realized using the MNA model, we perform a set of emulations by deploying both control and user plane functions of open-source 5G implementations in different locations of AWS, and measure the potential gains.

CCS CONCEPTS

• **Networks** → **Network measurement**; **Network performance analysis**; **Network design principles**.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in The 20th Annual International Conference on Mobile Systems, Applications and Services (MobiSys '22), June 27-July 1, 2022, Portland, OR, USA, https://doi.org/10.1145/3498361.3538942.

Mobisys'22, June 2022, Portland, OR, USA

KEYWORDS

Mobile networks, roaming, application performance, network aggregators

ACM Reference Format:
Sergi Alcalá-Marín, Aravindh Raman, Weili Wu, Andra Lutu, Marcelo Bagnulo, Ozgu Alay, and Fabián Bustamante. 2022. Global Mobile Network Aggregators: Taxonomy, Roaming Performance and Optimization. In *Proceedings of ACM Conference (Mobisys'22).* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

In the dawn of commercial cellular service, a handful of telco providers together with MNOs used to offer cellular connectivity services in a monolithic fashion (i.e., a single entity provided all the components of the mobile communications service). As both the market and the technology matured, new models emerged, notably, the onset of Mobile Virtual Network Operators (MVNOs) [24, 27], which leverage the infrastructure of existent base MNO to provide services.

Recently, we have been witnessing a sharp increase in the demand for connecting extremely heterogeneous terminals that operate globally in different environments around the world, with varying performance requirements. This surging demand in global / ubiquitous cellular connectivity comes both from the massive number of connected Internet of Things (IoT) devices as well as from people who are progressively switching (together with all their devices) to a digital nomad lifestyle. The rise of a large, new group of traveling, remote workers is one of the prevailing side-effects of the COVID-19-impacted work world [18]. This growing digital nomad community across the globe is fueling the demand for international seamless connectivity. Support for things roaming globally is now critical for IoT vertical applications, from connected cars to smart meters [17, 21].

The 'Global SIM' is now a product that individuals and IoT companies demand, and it is being satisfied by a a new breed of providers of mobile communications, the MNAs. Similarly to MVNOs, MNAs rely on the infrastructure of MNOs to provide services. However, instead of relying on a single base MNO, MNAs are able to multiplex clients across multiple MNOs in order to ensure optimal service and sustained Quality of Experience (QoE), without the added cost of operating the network.

Roaming is an essential service that can support MNAs operations in multiple countries, without the need of finding a local communication provider in every country where their end-users operate. MNAs benefit from the extensive global network infrastructure that international carriers (e.g., incumbent tier-one operators such as Vodafone, Tata, Telefónica or Orange) have been shaping for the past decades.

The emerging MNA model is appealing to the Internet companies, which now crossed into the telco world, such as Google's Fi Project [28]. Furthermore, cloud communication platforms as a service (CPaaS) such as Twilio provide MNA services by aggregating networks at the international level, thus aiming for global service, and making connectivity available through simple interfaces to application and service developers world-wide.

In this work, we dissect the operational models of MNAs, and characterize them from different angles, including performance aspects and implications on the end-user. Measuring international roaming performance is challenging, as it requires international cooperation and logistics are daunting. We detail our experimental setup in Section 4. We run performance measurements for each MNA both from the "home" location (i.e., the US), and also capture their operations while roaming internationally in Europe (i.e., in Spain and Norway).

Our contribution in this paper is three-fold, as follows.

First, we introduce the operating models of MNAs, and we position them in relation to well-known models for MVNO operations by offering a complete taxonomy (Section 2). We discuss the role of the roaming function in the context of MNAs with global coverage, and map its implementation in the case of the different models of MNA (Section 3).

Second, we uncover the underlying network architecture (Section 5) supporting the operations of three commercial MNAs (e.g., Google Fi, Truphone and Twilio). We focus on the impact of global operators in roaming on end-user performance, and we focus on evaluating several aspects, including DNS resolution delay and application (i.e., web and video) performance (Section 6). We capture the performance of the MNAs from the end-user perspective, and discuss the implication of different solutions for implementing international roaming. Despite their promise and potential, our results show that MNA models are currently in their infancy, and still suffer from the impact of Home-routed Roaming (HR) roaming (Section 6), similar to traditional MNOs [20].

Third, we leverage the 5G Control and User Plane Separation (CUPS) concept, and implement a realistic approach for global cellular operations that tackles some of the shortcoming of current roaming solutions (Section 7). CUPS is essential to 5G networks because it allows operators to separate the packet core into a control plane that can sit in a centralized location (e.g., the "home" country), and for the user plane to migrate closer to the application it is supporting (e.g., in the visited country of the end-user). We present a pilot implementation of a Regional Breakout (RBO) solution that leverages CUPS, using open-source software. We deploy our RBO solution on top of AWS infrastructure, and present our results that demonstrate the performance benefit of this roaming implementation.

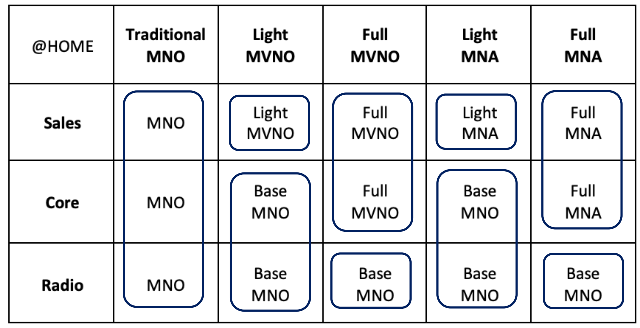

Figure 1: Types of MNOs.

2 TAXONOMY: MNO, MVNO AND MNA

There are several types of mobile operators with different operation models available in the market today; we capture these configurations in Figure 1.

An MNO[1] is an entity that owns (or has the exploitation rights) of a cellular network (i.e., base stations, network core, spectrum, etc). This was the initial operation model deployed to provide mobile communication services. Examples of MNOs include Vodafone, Orange, O2, AT&T, NTT to name a few. Later on, the MVNO operation model emerged [29]. Specifically, the MVNO is an entity that offers mobile network services to end-users, but does not own nor operate fully a cellular network. The MVNO is defined by its lack of ownership of radio spectrum resources.

In order to operate, an MVNO needs to have agreements in place to access the network of a base MNO. The implementation of the MVNO varies, and thus there are many different types on MVNOs. The type of MVNO is determined by how "thick" or "thin" a technological layer the MVNO adds over its access to its base MNO's network [24, 25, 27].

A *light MVNO* is a service provider that has its own customer support, marketing, sales and distribution operations, and may have the ability to set its tariffs independently from the retail prices of the base MNO. One such example is giffgaff in the UK, which uses O2 UK as a base MNO.

A *full MVNO* has a core network implementation operating essentially the same technology as an MNO, only missing their own radio network. They thus run their own core network, and rely on a base MNO who can offer access to radio resources. One example is Sky Mobile, which operates as a full MVNO in the UK, using O2 UK as a base operator.

Much more recently, we have witnessed the emergence of a new type of MVNO, namely the MNA [28]. While "traditional"

[1]This terminology distinguishes mobile operators as a general concept and Mobile Network Operators (MNOs) that is a specific type of mobile operators (i.e., mobile operators = set(MNOs, MVNOs, MNAs)).

MVNOs have agreements with a single base MNO, an MNA is an MVNO that exploits more than one base MNO, either in one single economy, or across different economies. Examples of MNAs include Google Fi, Truphone, Twilio or Lycamobile. Aggregating multiple base MNOs allows the MNA's customers to dynamically change the base MNO to which they attach. This change of base MNOs depends on different factors, including policy, coverage or performance.

In this paper, we extend the currently used MVNO-specific taxonomy [24, 25, 27] to include the MNAs. We further classify them into *full MNAs* or *light MNAs* , depending on whether they operate their own core network or not. We also differentiate the MNAs based on the geographic coverage of the multiple base MNOs they aggregate. In the general case, the base MNOs aggregated can cover the same or different geographic regions. A particular case is when the different base MNOs aggregated provide coverage in different geographic regions that do not overlap, notably different economies. If this is the case, we call this specific type of MNA a multi-country MVNO. These multi-country MVNOs usually have commercial offers in each of the different economies where they operate.

We acknowledge that, as in most taxonomies, there are corner cases that we cannot neatly classify into one of the categories. In our case, there is the case where a full MVNO has a commercial agreements with one or several IP Packet Exchange (IPX) Provider (IPX-P) [16], and does not depend on a specific base MNO (e.g., the MVNO might use global IMSI ranges). In this case, with a single agreement, the MVNO has "direct" access to several base MNOs located in different economies, depending on the footprint of the IPX-P. This configuration lies somewhere between the full MNA and the full MVNO, since it has a single agreement but connects to multiple base MNOs. However, we classify this the full MVNO category, since it is closer to the case where the MVNO has an agreement with a single entity and leverages its roaming agreements.

We highlight the lack of knowledge in our community around how these different models of MNAs satisfy the need for global coverage for their end-users. We further give background on roaming, which is one of the fundamental functions mobile operators ensure to their end-users, and is specifically relevant for these operators that aim for global uninterrupted service.

3 ROAMING FOR VIRTUAL OPERATORS

MNOs have customers, which are the end-users of mobile devices that normally attach to the MNO's network to access mobile communications services. The MNO represents the Home Mobile Network Operator (HMNO) for these end-users. The customers of an MNO can also attach to other radio networks owned by different MNOs. This happens when the HMNO does not offer radio coverage in the geographical region where the end-user wants to connect and access mobile communication services. A typical example for this is when the end-user travels abroad, in a foreign country. In this case, the end-user is *roaming*, and thus can attach to a "visited" cellular network, which the *visited MNO* operates in the foreign country. Mobile roaming is a fundamental characteristic of mobile service, which the cellular ecosystem enables through a tightly interconnected network of carriers and MNOs [16].

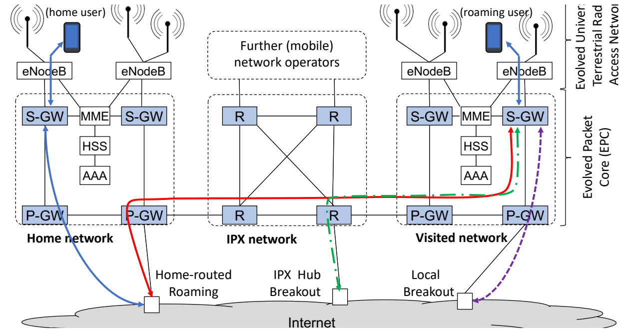

Figure 2: Internet access options when roaming home-routed roaming, local breakout, and IPX hub breakout roaming configurations.

3.1 Roaming Background

MNOs commonly connect with each other through an IPX network. An IPX [16] is a hub service, to which MNOs connect over a private IP backbone network. Usually, telco carriers operate as IPX-Ps, offering the IPX service, and interconnecting in a full mesh with all the other IPX-Ps to for the IPX Network. An IPX-P has connections to multiple MNOs, and thus enables each MNO to connect to other operators via a single point of contact. The interconnections between Mobile Network Operators (MNOs) are accompanied by roaming agreements that enable the operators to apply policies, control network access for roaming subscribers, and manage their roaming services. Figure 2 illustrates the three main schemes that MNOs employ for providing data roaming services, which we further describe below.

With **HR** [9], the IP address of the roaming user is provided by the home network. All traffic to and from the mobile user is routed through the home network, for which a GPRS Tunneling Protocol (GTP) tunnel is set up between the Serving Gateway (SGW) of the visited network and the Packet Data Network Gateway (PGW) of the home network (red path in the figure). With the IP end point in the home network, all services will be available in the same way as in the home network.

When **Local Breakout (LBO)** [9] is used, the IP address of the roaming user is provided by the visited network. The GTP tunnel is terminated at the PGW of the visited network and IP-based services can be accessed directly from there (purple path in the figure). This does not add latency and reduces network resource usage, but may restrict access to private services in the user's home network. Service control and charging also become more complex using LBO.

IPX Hub Breakout (IHBO) [16] provides an alternative to overcome the limitations of home-routed roaming and local breakout. Here, the IP address of the roaming user is provided by the IPX network. The GTP tunnel from the SGW in the visited network terminates at a PGW in the IPX network (green path in the figure). There may be multiple PGWs so that latency and resource usage can be reduced by selecting one geographically close to the visited network. As the IPX network maintains a trusted relationship with the home network, it may assign an IP address recognized by the home network to the roaming user, thereby allowing the user access also to private services in the home network. IHBO can also

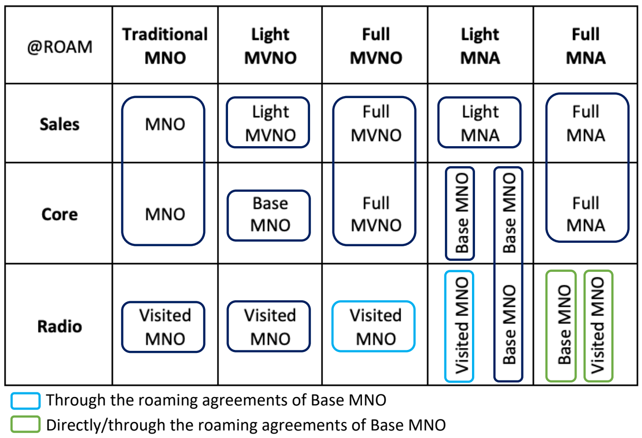

Figure 3: Roaming operations for the different types of mobile operators.

simplify setup and management as a single GTP tunnel, terminated in the IPX network, can be used for roaming users from different home networks.

These configurations might have an impact on the communication performance [20]. For instance, when the node accesses services inside the *visited network*, the performance is likely to be worse in the HR case, because all packets travel twice between the visited and the home country; less so when the communication peer is in a third country and minimally when accessing services in the home country. This may also have implications in the selection of Content Delivery Network (CDN) when roaming abroad, because the mobile user will access a server in the home network rather than one close to their location.

3.2 Mix-and-match: MVNO and MNA Roaming

In this section, we discuss and analyze the roaming operations for the different types of mobile operators we introduced in Section 2 (see Figure 3). Previous work [20] shows that currently, the vast majority (if not all) MNOs use HR for roaming, assuming the associated performance penalties that it implies. As we show in the first column of Figure 3, this means that the MNO relies for radio access on the visited MNO, while using its own core network functions.

In the case of the MVNO, when end-users of MVNOs roam internationally, there are several options for managing their connectivity. Given that, by definition, a light MVNO relies on a single base MNO, it then follows that they also rely on the roaming agreements that the base MNO has with visited networks in the roaming location. The difference between the light MVNO and the full MVNO is that, in the latter case, since the MVNO operates their own core network, they also handle roaming independently from the base MNO (e.g., they rely on a roaming hub service from an IPX-P). Thus, the full MVNO use their own core network for the roaming solution, while the light MVNO relies on the base MNO's core network.

Regarding the MNAs, since they rely on multiple base MNOs (accross different economies), they can obtain connectivity while roaming through a local MNOs with which they have a direct agreement (i.e., the local MNO in the visited country acts as a base MNO for the MNA) or they can connect to a visited MNO that has a roaming agreement with one of the MNA's base MNO elsewhere. Depending on whether it is a full or a light MNA, it will use its own core network, or it will rely on the core network of one of the base MNOs. Notably, multi-country MNA break out in the same country or in the same region where the device is roaming.

4 EXPERIMENTAL SETUP

We perform a number of experiments to understand roaming operations for MNAs and characterize their performance. To this end, we subscribe to several MNAs in the U.S., and we perform the experiments while roaming in different countries in Europe (Spain and Norway). We use three different MNAs, namely Google Fi, Truphone and Twilio.

Fi is Google's MNA service. Fi is only available for U.S. customers. This means that we must activate any new Fi account in the U.S., and, only after that, we can use it internationally.[2] Fi only works with a limited set of mobile phones, and some features (e.g., the capability of dynamically switching between underlying MNOs) are supported only in a subset of the devices that are "designed for Fi",[3] which in most cases only include the U.S. version of the devices. Fi automatically connects to the Virtual Private Network (VPN) provided by Google. This not only provides security for the connection, but it also allows to preserve the IP address used by the mobile device when communicating (even if the MNO used to attach to the network varies). It is possible to disconnect the VPN service manually.

In the US, Fi uses T-mobile, Sprint and U.S Cellular as base MNOs[4]. Fi also connects to a large number (millions, as claimed by Fi) of pre-selected WiFi hotspots. While roaming, Fi claims that it has coverage is over 200 countries, but provides very little information about how they achieve this. In particular, Fi announced an agreement with MNO Three (owned by Hutchison Telecommunications) to improve performance for end-users in roaming [15]. While the selection of the base MNO is automatic, it is possible to force the the MNO used by Fi using dialer codes [3]. These codes allow the selection of the base MNO; however, they do not allow the explicit selection of the visited MNO while roaming. Changing the base MNO may affect the visited MNO (which depends on the roaming agreements of the base MNO with the visited MNOs available in the visited countries).

Based on the publicly available information, we classify Fi as a light MNA, according to our proposed taxonomy.

Truphone is an MNA with headquarters in the UK, and with (MVNO) separate agreements with base MNOs in 8 countries (Australia, Germany, Hong Kong, Poland, Spain, the Netherlands, United Kingdom and the United States) and "bi-lateral roaming agreements in place with a wide range of operators around the world" (from https://en.wikipedia.org/wiki/Truphone). Truphone is a light MNA, as per the taxonomy we propose. The company has built a mobile

[2]Activate Google Fi service: https://support.google.com/fi/answer/6078618?co=GENIE.Platform%3DiOS&hl=en&oco=0
[3]Fi supported phones: https://fi.google.com/about/phones/.
[4]See the answer to "What is unique about Google Fi's network?" in the FAQ: https://fi.google.com/about/faq/.

network with core network technology distributed across four continents. Truphone uses these local points of presence (POPs) to reduce the distance that mobile traffic has to travel, which comes with a promise to reduce latency, and improve the end-user experience.

Twilio's Super SIM is an MNA targeting IoT devices [6]. Twilio's Super SIM can connect to 343 networks in 174 countries [7] and it uses its own mobile core that runs in the AWS cloud. As such, Twilio's Super SIM is a full MNA according to our taxonomy. However, local breakout outside the U.S. is still under development at the time of this writing. We also tested Twilio's Wireless SIM, which is an MVNO operating on top of T-Mobile (US) as base MNO (similar to Google Fi) [5].

To perform the measurements, we subscribed to the three aforementioned MNA services in the US, and ran experiments while roaming in two different countries in Europe (Spain and Norway). For end-user equipment, we use Pixel 4A (U.S. version) mobile phones in all experiments. In addition to the MNA services, we also subscribe to a local MNO, in order to be able to compare the MNA roaming solution with a local breakout roaming configuration. Moreover, we also subscribed to Three (UK/AT based), in order to also consider the case of a regional breakout in Europe (based on the configuration of Fi). In the case of Fi, we can use the dialer codes we described earlier to select the base MNO, so we toggle between the different possibilities available during the measurements. Also, in the case of Fi, even though the default behavior is to connect the VPN, we performed measurements with and without VPN. In Table 1 we include all configurations that we used for our experiments while measuring in Spain.

For each of these configurations, we run the following set of experiments:

- Traceroute: We performed a number of traceroutes to discover and characterize the paths from the end-user to reach different destinations. We selected targets in the US (i.e. home), in Spain (visited country), in Belgium (visited region) as well as content served by a CDN.
- DNS measurements: We measured the resolution time for both cached and non cached domain names.
- Web performance: We measured the Page Load Time (PLT) for web pages hosted by a server located in the US (home), Spain (visited country), Belgium (visited region) as well a served by a CDN.
- Video performance: We measured the average quality, the number of rebuffering events and the bandwidth obtained while streaming.

5 MNA ROAMING CONFIGURATION

In order to learn about the MNAs' underlying infrastructure and roaming configuration, we run traceroute from end-users (roaming) in Spain towards four different targets in Europe and in the US. We repeat the measurements for each experimental end-user SIM configuration that we include in Table 1. We select as targets simple web pages that are not served by CDNs, and operate from servers located in different countries (namely, ucla.edu (US), uclouvain.be (Belgium) and url.edu (Spain)). In addition, we also target a web page served by a CDN (i.e., mit.edu). We performed 20 traceroute measurements for each target, and we only keep the minimum Round-Trip Time (RTT) value observed for each hop, as we are interested in measuring the fixed components of delay at this stage. We also run additional tests to further assess performance in Section 6.

We analyze the data paths, and infer the roaming configuration these operators deploy (e.g., HR, LBO or IHBO). We infer the geolocation of each hop along the data paths that traceroute uncovers. For this, we use reverse DNS information, WHOIS information and the Maxmind geolocation database. We acknowledge the limitations of all these approaches, and mention that given the country-level geo-location we aim to achieve, the approach we use is accurate enough [14].

5.1 Traceroute to US

Figure 4a shows the delays for each hop replying to the traceroute probes from a mobile device (roaming) in Spain towards a server (ucla.edu) located in the US (the home country). We can observe that the overall delay to the target varies between 200 ms and 240 ms (20% variation) across the different SIM configurations (Table 1).

One major difference that we observe is the relative location of the transatlantic link in the path. When we measure Fi/T-Mobile (with and without VPN) or Fi/3/VPN, we find that the transatlantic link is before the first hop. The large delay value we measure in the first hop (≈200ms), and the inferred geo-location of the IP address of the first hop in the US both corroborate our deduction.

However, when using the Orange Spain subscription (i.e., the visited MNO that Fi attaches to in Spain when using Three as base MNO), we find that the first hop geolocates in Spain (the delay is ≈40 ms, and the geolocation of the IP of the first hop maps to Spain). We can easily identify the transatlantic link later in the path due to the large increase in the delay (≈200ms), and due to the fact that the IP geolocation shifts from Spain to the US.[5]

In the case of Fi/3/noVPN we find that the first hop is within Europe (but not in Spain), and that the transatlantic jump also happens later along the path. Similarly, for Truphone, we observe that the first hop is within Europe (but not in Spain) and it behaves similarly to a local MNO (namely, similar to Orange Spain, as we observe in the Fig. 4a). However, Truphone has slightly higher delay values in the first hop compared to Orange Spain, since it breaks out in the Netherlands.

When using Twilio, we observe two large increases in the delay, one in the first hop (≈150ms), and another one later on (an extra of ≈100ms).

For Google Fi, we further explain the different delay values we observe as effects of two specific design factors that the MNA includes in their roaming implementation, namely the HR roaming configuration and the location of the VPN endpoint. When the VPN is enabled, the first hop in the data path is the other VPN tunnel endpoint. From the experiments, we conclude that this is located within Google's infrastructure in the US. So when the VPN is enabled, the traffic is routed to Google in the US, and then towards the Internet through Google.

When we disable the VPN, we observe the effect of HR roaming. When using Fi/T-Mobile, we note that the traffic is first routed

[5]We observed similar behaviour using Vodafone (i.e., the visited MNO that Fi attaches to when using T-Mobile as underlying MNO); we do not include it the graph to improve readability.

Table 1: Experimental configurations: MNAs, base MNOs and visited MNOs we used for measurements in Spain. the Breakout column includes the breakout point identified through our experiments.

Name	MNA	Base MNO	Visited MNO	Comments	Breakout
Fi/TM/VPN	Fi	T-Mobile (US)	Vodafone (ES)	VPN enabled	US
Fi/TM/noVPN	Fi	T-Mobile (US)	Vodafone (ES)	VPN disabled	US
Fi/3/VPN	Fi	Three	Movistar / Orange/ Vodafone/ Yoigo	VPN enabled	US
Fi/3/noVPN	Fi	Three	Movistar / Orange/ Vodafone/ Yoigo	VPN disabled	UK
Truphone	Truphone	Orange (ES)	N/A		Europe
Twilio_WS	N/A	T-Mobile (US)	Movistar	MVNO	US
Twilio_SS	Twilio	N/A	Movistar	MNA	US
Orange	N/A	Orange (ES)	N/A	Baseline (Spain)	Spain
3NET	N/A	Three	Movistar / Orange/ Vodafone/ Yoigo		UK

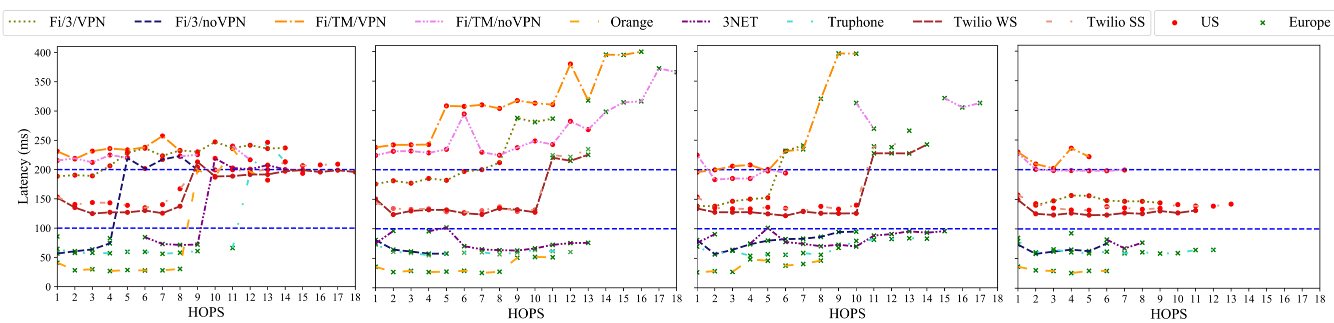

(a) Spain –> US (b) Spain –> Belgium (c) Spain –> Spain (d) Spain -> CDN

Figure 4: Hops encountered with the latencies towards targets in US, Belgium, Spain and within a CDN from a mobile device roaming in Spain with different SIM configurations (see Table 1).

to T-Mobile in the US, and it then exits to the Internet through T-Mobile's US network. When using Fi/Three or directly Three, we note that the traffic is first routed to Three network in Europe and then to the Internet through Three's network.

In the case of Twilio, we can explain the behavior we capture because Twilio breaks out in the East Coast (VA) [6], while our target is located on the West Coast.

For Truphone, since Truphone breaks out in Europe, the behavior we observe is very similar to the local MNO (i.e., Orange Spain).

5.2 Traceroute to Europe

We next analyze the traceroute results we collect from measuring towards a server located in Belgium (uclouvain.be), which we show in Fig. 4b. In the previous section, the measurements target located in the US. (ucla.edu) forced all the data paths to traverse from Europe to the US. Consequently, the home routed roaming configuration of the MNAs did not translate into a significant impact in the overall experienced delay, even when using a local MNO from Spain, such as Orange. In our current setup, we do observe significant difference in the overall delays as a consequence of HR roaming, because we keep both the target and all the end-user devices within Europe.

We observe that for Orange, Truphone and Fi/3/noVPN we measure a significantly lower delay (between 50 ms and 75 ms) than the remaining SIM configurations. For both Twilio setups we measure 220 ms of total delay (200% to 300% increase compared to the above-mentioned operators). Even more, both Fi setups with VPN enabled as well as Fi/T-Mobile/noVPN have a delay ranging between 300 ms and 400 ms, which represent 500% to 700% increase compared to the values we measured for the above-mentioned group of operators. We conclude that these latter SIM configurations are impacted by circuitous routes from Spain to Belgium through the US. This hairpinning effect is either because of the VPN (for Google Fi) or because of the HR setup for roaming traffic flowing from Spain to Belgium. We observe smaller delay values for Orange, Truphone, Fi/3/noVPN because the corresponding traffic never leaves Europe, taking a much shorter path. We find that Twilio (in both configurations) also uses HR roaming, but the breakout geo-locates on the US East Coast. Thus, the penalty in terms of latency is lower than we measure for Fi/T-Mobile, which breaks out in the US West Coast.

To further analyze the impact of the content location within Europe, we further run traceroute experiments between a mobile device in Spain and a server also located in Spain. We show our result in Figure 4c. If we look at the total delay we captured, we find that Orange Spain gives the smallest overall delay. We observe the second-smallest overall delay when using Fi/3/noVPN and Truphone, while the other cases (Fi/T-mobile with and without VPN and Fi/3/VPN) suffer an exceedingly larger delay (with Fi/T-Mobile/VPN being the largest).

Similar to the previous cases, we can explain these results as side-effect of a combination of HR and VPN tunnelling. However, in this case, the situation is more extreme because both the mobile device and the server are located in the same visited country, so the circuitous routing through the US or through the UK injects the extra delay towards the foreign network twice (to go and to come back to Spain). We validate this by the two large leaps in delay in Figure 4c (one for the first hop and another one later in the path). Nevertheless, the breakout in Europe (Three, Fi/Three, Truphone) introduces a significantly lower delay than the breakout in the US (for all other Fi configurations and Twilio).

5.3 Traceroute to a CDN

In this section, we discuss the traceroute results we collected when measuring towards a CDN-hosted web service (i.e., mit.edu). In this case, the CDN replica located close to the breakout point (where the end-user traffic is injected to the public Internet) provides the web content, and thus represents the target for our measurements. We illustrate our results in Figure 4d. As expected, we observe that the end-user experienced the shortest delay when using Orange Spain, followed by Fi/Three without VPN, Truphone and Three – all Europe-based operators. Both Twilio solutions, Fi/3/VPN, and Fi/T-mobile (with and without VPN) resulted in much higher delays. This implies that when the MNA deploys the VPN and/or the HR roaming to the USA, the end-user is accessing a content replica in the US. When using Fi/3/noVPN or Truphone, the end-user accesses a content replica in Europe (but not in Spain), resulting in a shorter overall delay. We measure the smallest overall delay when using Orange, which means that the end-user retrieves the content from a replica in Spain.

5.4 Takeaways

MNOs usually deploy the HR configuration for international roaming (Section 3). This configuration results in an added latency penalty for the end-user, especially when the other end of the communication is located topologically close to the visited location of the roaming device [19, 20]. We observe a similar behavior for the three MNAs we measured. The difference, however, comes from the capability of the MNA to change their base MNO (nationally and internationally), thus implicitly also changing their "home" operator.

For the delay measurements, we used as a baseline the delay we observe when using Orange Spain subscription within Spain. This is the delay the end-user would experience when directly attaching to a local network in the visited country.

Google Fi: When T-Mobile (without VPN) is the base MNO, Fi home-routes the traffic back to the US, resulting in large penalties for end-users roaming in Spain. This is especially noticeable when accessing content available locally in Spain/Europe. Even in the case when we target content served by a CDN, we still observe a large delay penalty. The only case when no significant delay is added is when the end-user targets content only available in the US. However, using T-Mobile as base MNO is *not* the default policy Fi has for end-users roaming in Europe. When in Europe, Fi switches to using Three as a base MNO. When using Fi/Three without VPN, the latency penalty (albeit still existent in some cases) drops. This is because, even with HR roaming, the distance between the visited location and the new "home" location (i.e., U.K.) is smaller compared to relying on T-Mobile. This approach is a middle-ground between LBO and HR, and significantly reduces latency with a small overhead in terms of roaming agreements (only one extra agreement for a whole region). We model this configuration as a "regional breakout" model, similar to the IPX breakout model. However, the benefits the "regional breakout" bring are lost when the Fi end-user enables the VPN service (which is the default behavior for Fi), because the other endpoint of the tunnel is located in the US.

Twilio: In the case of Twilio, both configurations (i.e., Super SIM and Wireless SIM) rely on HR roaming back to the USA with the corresponding penalties in terms of delay.[6] However, because Twilio breaks out in the US East Coast, the penalties are reduced in our case that the device is roaming in Europe. Devices roaming in the Asian Pacific region for instance, should observe the opposite effect.

Truphone: Overall, we observe that Truphone provides a delay experience slightly higher than that of the local MNO in the visited country (i.e., Orange Spain) as it breaks out in Netherlands. To confirm this, we also perform measurements in Norway using a Truphone SIM purchased in Spain. We find that the end-user traffic breaks out in UK, instead of the further "home" location in Spain. We thus confirm that Truphone uses their POPs to implements the "regional breakout" approach to reduce the distance that mobile traffic has to travel. We observe that this indeed reduces the delay.

6 MNA PERFORMANCE

Our traceroute results confirmed that MNOs implement "regional breakout" to reduce the distance that mobile traffic has to travel compared to the HR roaming configuration. However, it is unclear how this approach further improves the end-user experience for popular applications and services. In this section, we evaluate the impact of the MNA roaming configuration on DNS resolution delay, web browsing performance and video streaming quality.

6.1 DNS resolution delay

We measure the DNS resolution delay using the different experimental configurations in Table 1 to compare their performance. We first measure the resolution delay for a non-cached name. In order 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. We then perform queries for unique names under our domain, ensuring that the name queried is never repeated, ensuring that caching is not involved in the resolution. The authoritative server is located in Spain, serving the worst case for roaming devices as we inferred from the traceroute experiments (see Figure 4). We did 20 queries for each experimental configuration for the end-user SIM.

In Figure 5, we show the delays we measure while querying a non-cached and cached domain names. We note that two clusters of operators emerge while querying non-cached domain names.

[6]Some resources within the Twilio (i.e., the Super SIM API and Internet breakouts outside of the United States, such as Frankfurt, Germany) are still in Pilot or Beta at the time we performed our measurements in July-August 2021. Thus, our results are consistent with the Twilio configuration for Super SIM, where traffic breaks out to the Internet via the Twilio Mobile Core in Ashburn, US.

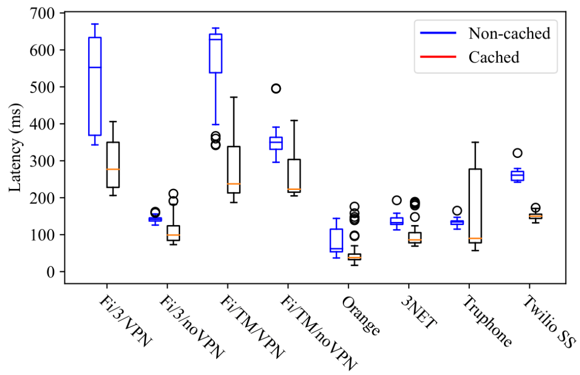

Figure 5: DNS resolution delays we measured for a non-cached and a cached DNS name using different experimental configurations for the end-user SIM (see Table 1).

First one includes both VPN-enabled Fi configurations (namely, Fi/TM/VPN and Fi/3/VPN), which exhibit the largest DNS resolution delay, followed by Fi over T-Mobile without VPN (Fi/TM/noVPN), and the two Twilios SIMs.[7] We observe that a second cluster of operators includes Three, Fi over Three (without VPN), Truphone and, finally, the (native) Orange with the smallest delay. We obtained similar results to Orange Spain with other local MNOs (e.g., Vodafone Spain), which, again, we do not include in Figure 5 to improve readability. These operators show much smaller resolution delay, comparable to the one a use connecting via a local operator in the visited country might experience.

Overall, we find that both VPN-enabled Fi configurations experience a surprisingly long delay (i.e., the mean delay is over 550 ms) and also a high variance. Fi over T-Mobile without VPN follows, with a mean delay of over 350ms. Then, we find Twilio with a mean delay close to 250 ms.

The two remaining configurations (Fi/3/noVPN and Truphone) have a mean delay close to 150 ms, while for the native Orange Spain configuration we obtain a mean close to 50 ms. Based on this, we conclude that the HR Roaming configuration which routes the traffic back to the USA imposes a penalty of 200-500ms compared to LBO roaming configuration. This accounts for a penalty that can vary between 300% and 1,000%. In the same time, we find that the regional breakout MNA leverage (e.g., here in the case of Fi using Three as base MNO and no VPN) only bring an extra 100 ms of delay, which accounts for 200% penalty. In particular, Truphone delivers best on the promise of offering optimal (i.e., close to using a local operator in the visited country) global experience to the end-user. Even when measuring with an US Truphone subscription in Spain, we note that the Mobile Country Code (MCC)/Mobile Network Code (MNC) (MCCMNC) changes to the one of the local operator Truphone registered in Spain. When traveling to a country where Truphone does not operate locally (e.g., Norway), we found that the MNA will indeed route the traffic to their nearest Point of Presence (PoP) (i.e., UK).

[7]Figure 5 depicts only one Twilio configuration (Twilio_WS) to avoid cluttering the plot; however, we found very similar results for both Twilio solutions for connectivity.

We next measure the delay for the resolution for a query that is present in the resolver's cache. To ensure this, 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 again), and query for the same domain name to measure the DNS resolution time. We repeat this 20 times for three popular domains (namely, www.amazon.com, www.facebook.com and www.youtube.com).

In Figure 5, we compare the DNS resolution delay for a cached name with that of non-cached ones. Overall, we find similar relative performance across the different configurations to that we observed in the case of querying a non-cached domain. The absolute values, however, as well as the absolute differences are smaller than in the case of the DNS queries for the non-cached domains. Indeed, for both configurations of Fi over T-Mobile and for Fi/Three with VPN, the delays are in the order of 250 ms. This is followed by Twilio with 150ms mean delay, while for Fi/Three without VPN and Truphone the delays are in order of 100ms. We observe that for Orange the DNS resolution delay drops to 50ms in median. This means that the penalty for using Home Routing back to the US is 300% - 400%, while for the regional breakout the penalty is 100% compared to the potential local breakout.

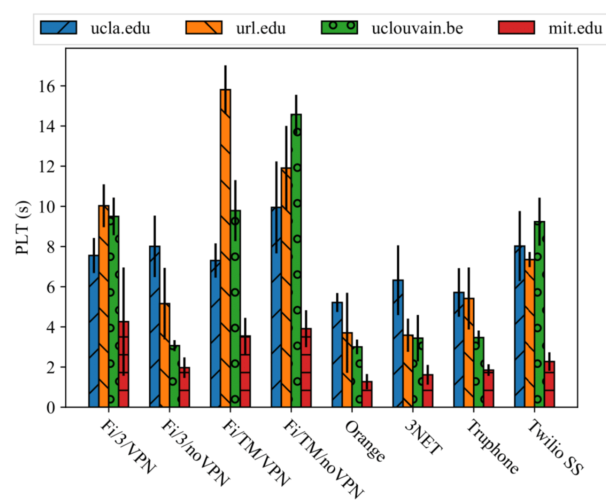

Figure 6: PLT measured towards web pages in US (ucla.edu), Belgium (uclouvain.be), Spain (url.edu) and within a CDN (mit.edu) from a mobile device roaming in Spain with different SIMs (see Table 1).

6.2 Web performance

While DNS resolution plays a crucial part in the overall performance of the service MNAs offer to their end-users, it might not necessarily translate into a significant impact on the end-user QoE. In order to further capture how the different MNA solutions actually deliver on end-user experience, we measure the impact of the delay introduced by the different experimental configurations on web browsing. We use the PLT metric to characterize web performance, which captures the time it takes for a webpage to load.

We extract the PLT from the Navigation Timing API [10], available in native android web browser (Google Chrome). We calculate the PLT from initiation (the `LoadEventStart`) to completion (the `NavigationStart`), 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 actually required to complete the rendering of the visible portion of the page. Though many other metrics focus on different aspects of webpage performance, recent studies [23] showed that PLT is good enough to capture the experience of the users in various radio contexts, showing strong correlation to other QoE metrics such as First Paint or Speed Index. Thus, in this paper we focus on the PLT, which we use as a proxy for end-user QoE.

We measure the PLT using the different configurations for the end-user SIM (see Table 1), while targeting 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 US, www.uclouvain.be in Belgium, www.url.edu in Spain andwww.mit.edu which is served by a CDN. For each end-user SIM configuration and for each web page, we collect 20 different measurements. We show our results in Figure 6.

We observe similar trends to the ones we obtained in the case of the DNS resolution delay (Figure 5) and traceroute (Figure 4). As before, the penalties are much larger when the web server is in Spain, followed (closely) by the case of the web server hosted in Belgium , and, finally, the server in the US, where the overall penalty is significantly smaller. Regarding the content served by the CDN, the overall PLT is smaller in all the cases, but the relative differences remain similar to the previous cases. In all cases, the smallest PLT is achieved with the native service offered by Orange in Spain. We then observe the cluster of MNAs which delivers the closest experience to that of the native operator, namely Fi over Three without VPN (Fi/3/noVPN) and Truphone. Next, we have Twilio and, finally, with the highest PLT, we see both VPN-enabled configuration for Fi (Fi/3/VPN, Fi/TM/VPN), as well as Fi over T-Mobile with no VPN (Fi/TM/noVPN). We would like to highlight that serving content through a CDN – while it reduces the delay for all configurations – does not eliminates the differences in the PLT for the different configuration. This is so because, even when a CDN is used, the content is retrieved from the replica that is closest to Internet access point associated to the end-user (i.e., the breakout point). When HR roaming or VPN are used, the mobile accesses to a replica in the U.S.A. while when using Three, and Fi/Three, the mobile connects through 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.

6.3 Video performance

Though web browsing represents a critical service in the mobile Internet ecosystem, video traffic accounts for the largest proportion. In this section, we thus investigate how the different MNA configurations in roaming impact the video delivery performance. We run active end-user measurements using the YoMoApp [26]. YoMoApp is developed by Wurzburg University that allows to rate the stream quality of YouTube videos, as well as obtain different

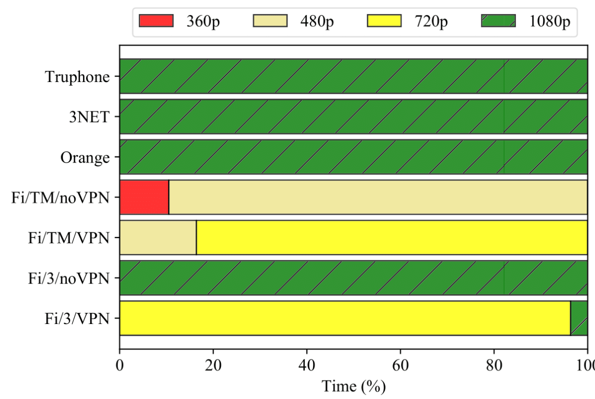

Figure 7: Video quality for the different SIMs in Table 1 while playing a music video.

metrics from the network. Once the measurement completes, it uploads to yomoapp.de/dashboard/, from where it is possible to retrieve the results. The tool allows us to measure the video quality, download throughput, stalling events, and also capture the radio access technology and buffer level.

We measure the video performance using the different configurations in Table 1), except for Twilio. We omitted Twilio from these tests because their products are oriented to IoT. Though they do offer video streaming APIs, the amount of traffic we needed to perform the video tests proved to be prohibitive.

We played short videos (2-3 min) from YouTube. We used three different videos, a sports video[8], a music video[9] and a movie trailer[10], to be representative of the different types of content available. We measured each video/configuration combination 20 times. We next present and analyze the results of the measurements obtained when the MNA end-user is roaming in Spain. In Figure 7 we show the time spent in each quality for the different configurations for the music video. We mention that we obtained similar distribution for the other types of videos.

We can observe that for Three UK, Fi with Three (without VPN), Truphone and Orange, the videos render in the highest resolution (1080p) – while in Fi over Three with VPN and for Fi over T-Mobile with VPN, the most common quality is 720p. Furthermore, for Fi with T-Mobile without VPN the most common quality is 480p, dropping to 360p 10% of the times.

Moreover, we look into the initial delay and the buffer levels in the video traces (these results are not presented here due to space limitations). We observe that Fi over Three with VPN and Fi over T-Mobile with and without VPN have introduced significant initial delay (highest initial delay is experienced by Fi over T-Mobile without VPN) compared to the other setups. We also observe that these three setups experience much lower buffer levels compared to the other setups, indicating inefficient use of buffering mechanism. We conclude that the long link delays experienced with these three

[8]https://www.youtube.com/watch?v=JEoubjE2PBQ
[9]https://www.youtube.com/watch?v=l6N-Yq9Fw4U
[10]https://www.youtube.com/watch?v=0WVDKZJkGlY

setups, impact the ABR mechanisms in a negative way, leading to a much lower video quality experience compared to the other setups.

6.4 Takeaways

We find that the HR roaming configuration all MNAs deploy impacts in a similar way critical services for the end-user, namely DNS resolution, web browsing and video streaming. However, by switching the "home" (i.e., the base MNO) closer to the end-user visited country, some MNAs such as Google Fi and Truphone succeed in reducing the penalty that home routed roaming introduces. We notably find that (based on our limited measurements and set of operators we test) Truphone owns the most mature MNA deployment.

Google Fi: The use of regional breakout in Europe on top of Three's network helps Fi to reduce significantly the delay their users experience in Europe. For DNS resolution, we find that using Three as base MNO reduced the delay penalty to 200% from 300%-1000% compared to using a local native operator. In terms of further impact on application performance, we find that Fi/3/noVPN allowed the end-user to stream videos in the higher resolution (1080p), and also provided the closest web browsing performance to that a local MNO would provide.

Twilio: With the SuperSIM solution still in a very early roll-out phase at the time of our measurements campaign (July-August 2021), we find that the Twilio end-user is still impacted significantly by the HR roaming setup. Given that their breakout point is closer to the visited location (Spain) than the ones for Fi over T-Mobile or Fi with VPN active, we find a slightly better performance for this MNA.

Truphone: Leveraging their mature setup with different PoPs, we find that Truphone is able to deliver the performance closest to the one provided by a local MNO in a visited country in Europe. This is true for all the different services and applications we tested.

7 REGIONAL BREAKOUT IN 5G

The split of user/data plane from the control plane at both the radio front [12, 13] and the core [22] is one of the major upgrades in 5G. This paves the way not just for better management of data and control packets, but also to enable truly global operations of an MNO. We argue that, with this approach, any MNO can potentially convert into a global provider, avoiding HR roaming configuration, and enabling the end-user to achieve a good experience potentially world-wide through regional breakout. We envision a scenario where the MNO provides the local breakout solution to each end-user. In this section, we provide a proof-of-concept implementation and evaluate how the separation at the mobile core can support global operations.

7.1 Experimental setup

We utilise the edge (wavelength), local and regional deployments of Amazon global infrastructure [2] to deploy control and user plane functions of open-source 5G implementations (namely, Open5GS [4] and UERANSIM [8]).

Infrastructure. The setup we build aims to emulate different roaming configurations (see Section 3). For this, we rely on the global infrastructure including both storage and compute services that AWS offers. This includes:

- a *regional* infrastructure with data centers present in a region (e.g., US East/Ohio region). Within this deployment, a cluster of isolated and physically separate data centers are found in a geographical area.
- a *local* infrastructure hosted as an extension of regional infrastructure to run latency sensitive and high bandwidth applications. For example, Netflix uses AWS local zone deployments for their content creation process.
- an *edge* infrastructure hosted within telecommunications providers' data centers and connected to the operator's 5G network. We consider this as first point an user can breakout to the Internet from MNOs network.

Connectivity. For our pilot deployment, we assume an end-user with 5G connectivity who has their network home location in the UK. To emulate the user roaming behavior, we test two different scenarios, where the user roams in two locations: (*i*) in the US (San Francisco) and (*ii*) within Europe (Berlin, Germany). With current 4G/LTE technologies, the default roaming configuration would be HR roaming, where the traffic is routed back to the UK. We argue that, by using 5G Control and User Plane Seperation at the core, we can keep the control plane functions in the trusted, centralized home network location, while dynamically moving the user plane function with the roaming user. We handle this connectivity by deploying control and user planes built using Open5GS [4]. We deploy the control plane, which includes Session Management Function (SMF) and Access and Mobility Management Function (AMF), at the regional infrastructure in the user's home location (London, UK). The user plane enables breakout to the internet, and hence we deploy it across multiple locations in the US and EU (as per Table 2). The selection of UPF to breakout is chosen by the DNN setting in the 5G UE. We use UERANSIM [8] to deploy a simulated environment of 5G RAN and 5G UE in the edge infrastructure.

We leverage the different existent AWS architecture to place the UPF at various locations in relation to the roaming end-user. Namely, we emulate the case of *edge breakout* by placing the function in the AWS wavelength deployment within the carrier infrastructure. We use one such existing Wavelength Zone in the US (Verizon central office in San Franciso), corresponding to breaking out the visited MNO, very closely to the location of the end-user. We then migrate the UPF further from the visited location (but still within the visited country) in each of the roaming scenarios, as we show in Table 2.

The UE and the migrated UPF instance across the Amazon EC2 zones in the US, UK and EU are connected via a private backbone network [11] and through virtual private cloud peering connections [1]. We show ping latency within the instances in the US and to UK (home) in Table 3.

7.2 Pilot Results

In order to capture the end-user performance under the scenarios we include in Table 2, we measure web performance. We focus on the PLT as the representative metrics, similar to our approach in Section 6. We use as target the same four website as previously in Section 6.2 (namely, www.ucla.edu in the US, www.uclouvain.be in

Table 2: Locations of deployed User Plane Function (UPF) and Data Network Name (DNN) names as identified by the UEs to breakout at a visited location (namely, US or Germany).

	UK		US				EU (Germany)	
BO type	Home (London)	Edge (Vodafone)	Edge (Verizon)	Local (Las Vegas)	Regional (Oregon)	National (Ohio)	Edge (Vodafone)	Regional (Frankfurt)
DNN	edge.london	home	edge.sfo	local.las	regional.or	national.oh	edge.ber	regional.fra

Table 3: Ping latency (in ms) to the instances deployed across various zones in the AWS from an instance located at same zone as Edge.

Edge	Local	Regional	National	Home
0.22 ± 0.03	26.2 ± 0.08	49.8 ± 0.1	78.6 ± 1.5	157.3 ± 0.2

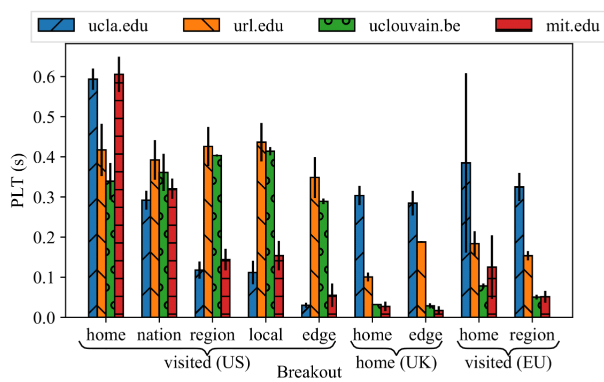

Figure 8: Page load time to different websites in function of the breakout point considered for the roaming configuration.

Belgium, www.url.edu in Spain and www.mit.edu which is served by a CDN). We show our results in Figure 8. We use as a baseline the measurements for the non-roaming scenario (marked "home (UK)" in Figure 8). The goal of our measurements analysis is to establish which roaming configuration can offer the same performance that the end-use enjoys while at home. We find than when content is served by a CDN, the regional breakout configuration offers comparable performance to the no-roaming scenario, regardless of the location where the end-user travels (e.g., US or Germany, in our case). This is a direct consequence of the close location of the content replica to the end user, which is dictated by the location of the breakout point. From the case of end-user roaming in the US, we see that the local breakout configuration yields similar performance to regional breakout, likely as a results of the small distance between the locations of the infrastructures used in this scenario.

When a CDN is not serving the web content, the distance between the location of the end-user breakout point and the content location impacts the web performance. For example, if edge breakout in San Francisco offers the optimal performance for accessing content hosted in California (ucla.edu), we see this is no longer the case when accessing content hosted in Europe (uclouvain.be, url.edu). The PLT we measure in this latter case is, in fact, similar to the one we measure under the HR roaming configuration. The same is true for accessing US-based content from Germany, under the regional breakout configuration. Surprisingly, however, we find that the PLT for Europe-hosted web content is slightly smaller in the US (San Francisco) edge breakout scenario than all of the other configurations (local/regional/national/home breakout). We conjecture that this is a side-effect of relying on the carrier's infrastructure (i.e., Verizon), while for the other configurations the AWS private backbone impacts the delays between the various instances (see Table 3).

The results we present here invite more broader discussions on the role of cloud service providers in supporting cellular connectivity providers. Given the freedom that the 5G allows in terms of dynamically locating the UPF for a roaming user, we envision this as a first step towards an adaptive approach for managing the end-user cellular connectivity.

8 APPLICABILITY

There are a growing number of use cases that heavily rely on roaming, enabled by both the surge in demand for global deployment of IoT devices (e.g., smart meter, connected cars) with pre-provisioned connectivity [17], and in the demand for global seamless connectivity from digital nomads. A new breed of global operators aims to directly respond to the needs of such users (e.g., Twilio for IoT and Fi & Truphone for digital nomads).

Our work shows that light MNAs such as Fi still have to mitigate the impact of latency, and consider regional breakout solutions. Furthermore, MNAs that deploy VPN features should also consider deploying them in the same regional breakout approach, to enable the privacy and security benefits for their end-users without penalizing performance.

Full MNAs (such as Twilio) should carefully consider the underlying infrastructure hosting their core network, as this will impact the performance. Our work highlights that one option might be to explore using CDNs with a wide geographical footprint as underlying infrastructure, and benefit from dynamically migrating different core functions to the optimal location.

For app developers, our work gives tangible insights into what they might expect in terms of performance from this new breed of global operators. Thus, an app developer or app provider (i.e., the entity providing services through the internet) who is aware of the performance limitations of a global connectivity model can accommodate and potentially mitigate them (e.g., for multimedia streaming apps, the developer might consider using larger buffers for video). This is especially relevant for the services aimed at IoT devices that IoT vendors distribute worldwide, and that are permanently roaming.

9 CONCLUSIONS

The mobile communications landscape is in constant (r)evolution. Very recently, we witnessed the emergence of a new type of global virtual operator, the MNA. In this paper, we provided the means to map out the operations of MNAs and reason about their potential evolution. To that end, we extended the currently available taxonomy of MVNOs, and included the MNA concept.

Our measurements allowed us to gain insights about the roaming operations of three MNAs with distinct operational models, namely Twilio, Truphone and Fi. We characterized their performance and quantified the impact of their operational approach on the performance of relevant applications, namely DNS, web browsing and video streaming.

We find that, opposed to MNOs and MVNOs which rely on HR roaming, MNAs are (slowly) moving to some limited forms of LBO roaming. We observe that, while current MNA operations do not enable LBO in the visited country, some implement regional breakout that at least keeps the traffic in the same continent/region and avoid long transoceanic links. This is clearly reflected in the performance of the different applications we tested. The purpose of our measurements is to learn about the MNAs operations models and their impact on the performance of several applications. However, it is not our purpose to benchmark the different MNOs and MNAs tested, as our measurement campaigns are limited to a few vantage points and a few applications.

Finally, we looked into possible evolution of the MNAs model and we explored, through emulations, the potential performance gains that the full exploitation of the LBO can provide. We find that, depending on the application, the benefits of regional breakout may be good enough, and few additional benefits come from implementing a full blown LBO.

ACKNOWLEDGMENTS

We thank the MobiSys anonymous reviewers and our shepherd - Rajesh Balan - for their helpful comments and guidance. This work has been partially funded by the NGI Atlantic MCCA project, and by the European Union Horizon 2020 Research and Innovation Action Grant Agreements 101017109 (DAEMON) and 101016509 (CHARITY). The work of Andra Lutu was supported by the EC H2020 Marie Curie Individual Fellowship 841315 (DICE).

REFERENCES

[1] [n.d.]. Amazon Virtual Private Cloud. https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html. Accessed: 2021-12-10.

[2] [n.d.]. AWS Global Infrastructure. https://aws.amazon.com/about-aws/global-infrastructure/. Accessed: 2021-12-10.

[3] [n.d.]. Google Fi dialer codes. https://arkienet.com/2018/01/google-fi-dialer-codes/. Accessed: 2021-12-10.

[4] [n.d.]. Open5gs v2.3.6. https://github.com/open5gs/open5gs. Accessed: 2021-12-10.

[5] [n.d.]. Programmable Wireless: Gain global cellular IoT connectivity. https://www.twilio.com/iot/wireless. Accessed: 2021-12-10.

[6] [n.d.]. Twilio Super SIM. https://www.twilio.com/docs/iot/supersim. Accessed: 2021-12-10.

[7] [n.d.]. Twilio Super SIM: Available Networks. https://www.twilio.com/docs/iot/supersim/available-networks. Accessed: 2021-12-10.

[8] [n.d.]. UERANSIM v3.2.4. https://github.com/aligungr/UERANSIM. Accessed: 2021-12-10.

[9] GSM Association. 2016. LTE and EPC Roaming Guidelines. https://www.gsma.com/newsroom/wp-content/uploads/IR.88-v15.0.pdf. [Online; accessed 06-March-2018].

[10] MDN contributors. 2020. Navigation Timing API. https://developer.mozilla.org/en-US/docs/Web/Performance/Navigation_and_resource_timings. In *MDN Web Docs*. last accessed in 23/6/2021.

[11] Lorenzo Corneo, Maximilian Eder, Nitinder Mohan, Aleksandr Zavodovski, Suzan Bayhan, Walter Wong, Per Gunningberg, Jussi Kangasharju, and Jörg Ott. 2021. *Surrounded by the Clouds: A Comprehensive Cloud Reachability Study*. Association for Computing Machinery, New York, NY, USA, 295–304. https://doi.org/10.1145/3442381.3449854

[12] Xenofon Foukas, Mahesh K. Marina, and Kimon Kontovasilis. 2017. Orion: RAN Slicing for a Flexible and Cost-Effective Multi-Service Mobile Network Architecture. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking* (Snowbird, Utah, USA) *(MobiCom '17)*. Association for Computing Machinery, New York, NY, USA, 127?140. https://doi.org/10.1145/3117811.3117831

[13] Xenofon Foukas, Navid Nikaein, Mohamed M. Kassem, Mahesh K. Marina, and Kimon Kontovasilis. 2016. FlexRAN: A Flexible and Programmable Platform for Software-Defined Radio Access Networks. In *Proceedings of the 12th International on Conference on Emerging Networking EXperiments and Technologies* (Irvine, California, USA) *(CoNEXT '16)*. Association for Computing Machinery, New York, NY, USA, 427?441. https://doi.org/10.1145/2999572.2999599

[14] Manaf Gharaibeh, Anant Shah, Bradley Huffaker, Han Zhang, Roya Ensafi, and Christos Papadopoulos. 2017. A Look at Router Geolocation in Public and Commercial Databases. In *Proceedings of the 2017 Internet Measurement Conference* (London, United Kingdom) *(IMC '17)*. Association for Computing Machinery, New York, NY, USA, 463–469. https://doi.org/10.1145/3131365.3131380

[15] Tyler Kugler. [n.d.]. Google Fi: Stay connected abroad with high speed data from Project Fi. https://www.blog.google/products/project-fi/stay-connected-abroad-with-high-speed/. Accessed: 2021-12-10.

[16] Andra Lutu, Byungjin Jun, Fabián E. Bustamante, Diego Perino, Marcelo Bagnulo, and Carlos Gamboa Bontje. 2020. A First Look at the IP EXchange Ecosystem. *SIGCOMM Comput. Commun. Rev.* 50, 4 (Oct. 2020), 25?34. https://doi.org/10.1145/3431832.3431836

[17] Andra Lutu, Byungjin Jun, Alessandro Finamore, Fabián E. Bustamante, and Diego Perino. 2020. Where Things Roam: Uncovering Cellular IoT/M2M Connectivity. In *Proceedings of the ACM Internet Measurement Conference* (Virtual Event, USA) *(IMC'20)*. Association for Computing Machinery, New York, NY, USA, 147 – 161. https://doi.org/10.1145/3419394.3423661

[18] Andra Lutu, Diego Perino, Marcelo Bagnulo, Enrique Frias-Martinez, and Javad Khangosstar. 2020. A Characterization of the COVID-19 Pandemic Impact on a Mobile Network Operator Traffic *(IMC'20)*. Association for Computing Machinery, New York, NY, USA, 19 – 33. https://doi.org/10.1145/3419394.3423655

[19] Andra Lutu, Martino Trevisan, Ali Safari Khatouni, Anna Maria Mandalari, Ana Custura, Marco Mellia, Ozgu Alay, Marcelo Bagnulo, Vaibhav Bajpai, Anna Brunstrom, Jorg Ott, and Godred Fairhurst. 2021. Measuring Roaming in Europe: Infrastructure and Implications on Users QoE. *IEEE Transactions on Mobile Computing* (2021), 1–1. https://doi.org/10.1109/TMC.2021.3058787

[20] 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* (New Delhi, India) *(MobiCom '18)*. Association for Computing Machinery, New York, NY, USA, 179–189. https://doi.org/10.1145/3241539.3241577

[21] Emeka Obiodu, Aravindh Raman, Abdullahi Abubakar, Simone Mangiante, Nishanth Sastry, and Hamid Aghvami. 2022. DSM-MoC as Baseline: Reliability Assurance via Redundant Cellular Connectivity in Connected Cars. *IEEE Transactions on Network and Service Management* (2022), 1–1. https://doi.org/10.1109/TNSM.2022.3153452

[22] Idelkys Quintana-Ramirez, Anthony Tsiopoulos, Maria A. Lema, Fragkiskos Sardis, Luis Sequeira, James Arias, Aravindh Raman, Ali Azam, and Mischa Dohler. 2018. The Making of 5G: Building an End-to-End 5G-Enabled System. *IEEE Communications Standards Magazine* 2, 4 (2018), 88–96. https://doi.org/10.1109/MCOMSTD.2018.1800024

[23] Mohammad Rajiullah, Andra Lutu, Ali Safari Khatouni, Mah-Rukh Fida, Marco Mellia, Anna Brunstrom, Ozgu Alay, Stefan Alfredsson, and Vincenzo Mancuso. 2019. Web experience in mobile networks: Lessons from two million page visits. In *The World Wide Web Conference*. 1532–1543.

[24] Paul Schmitt, Morgan Vigil, and Elizabeth Belding. 2016. A study of MVNO data paths and performance. In *International Conference on Passive and Active Network Measurement*. Springer, 83–94.

[25] Narseo Vallina-Rodriguez, Srikanth Sundaresan, Christian Kreibich, Nicholas Weaver, and Vern Paxson. 2015. Beyond the radio: Illuminating the higher layers of mobile networks. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services* (Florence, Italy) *(MobiSys '15)*. 375–387.

[26] Florian Wamser, Michael Seufert, Pedro Casas, Ralf Irmer, Phuoc Tran-Gia, and Raimund Schatz. 2015. YoMoApp: A tool for analyzing QoE of YouTube HTTP adaptive streaming in mobile networks. In *2015 European Conference on Networks and Communications (EuCNC)*. IEEE, 239–243.

[27] Ao Xiao, Yunhao Liu, Yang Li, Feng Qian, Zhenhua Li, Sen Bai, Yao Liu, Tianyin Xu, and Xianlong Xin. 2019. An in-depth study of commercial MVNO: Measurement and optimization. In *Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services* (Seoul, Republic of Korea) *(MobiSys '19)*. 457–468.

[28] Zengwen Yuan, Qianru Li, Yuanjie Li, Songwu Lu, Chunyi Peng, and George Varghese. 2018. Resolving Policy Conflicts in Multi-Carrier Cellular Access. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking* (New Delhi, India) *(MobiCom '18)*. Association for Computing Machinery, New York, NY, USA, 147–162. https://doi.org/10.1145/3241539.3241558

[29] Fatima Zarinni, Ayon Chakraborty, Vyas Sekar, Samir R Das, and Phillipa Gill. 2014. A first look at performance in mobile virtual network operators. In *Proceedings of the 2014 conference on internet measurement conference* (Vancouver, BC, Canada) *(IMC '14)*. 165–172.