This is the authors' accepted version of the article. The final version published by IEEE is S. Alcalá-Marín, W. Wu, A. Raman, M. Bagnulo, O. Alay, F. Bustamante, M. Fiore and A. Lutu. "A Comparative Analysis of Global Mobile Network Aggregators," in IEEE Transactions on Network and Service Management, vol. 22, no. 5, pp. 4651-4667, Oct. 2025, doi: 10.1109/TNSM.2025.3582089.

A Comparative Analysis of Global Mobile Network Aggregators

Sergi Alcalá-Marín, Weili Wu, Aravindh Raman, Marcelo Bagnulo, Ozgu Alay, Fabián E. Bustamante, Marco Fiore, *Senior Member, IEEE*, Andra Lutu

***Abstract*—The mobile telecommunication industry is undergoing continuous evolution to cope with ever increasing service requirements and expectations of end users. This has recently led to the rise of Mobile Network Aggregators (MNAs), a new type of global virtual operators that deliver mobile communication services by utilizing multiple Mobile Network Operators (MNOs), dynamically connecting to the one that best meets their customers' needs based on location and time. MNAs can then offer optimized global coverage by connecting to local MNOs that have limited (e.g., national) geographic service. In this paper, we provide a first in-depth analysis of the operations of three major MNAs: Google Fi, Twilio, and Truphone. We conduct performance measurements across these MNAs for critical applications spanning DNS, web browsing, and video streaming, and compare their performance against that of a traditional MNO from two very diverse geographical locations, US and Spain. We find that MNAs may introduce some delay compared to local MNOs in the region where the user is roaming, yet they offer significant performance improvements over the traditional MNOs roaming model, such as home-routed roaming. To fully assess the potential benefits of the MNA model, we also carry out emulation studies assessing the potential performance gains that MNAs could achieve by deploying both control and user plane functions of open-source 5G implementations across different Amazon Web Services locations.**

***Index Terms*—Mobile networks, roaming, application performance, network aggregators, 5G**

I. Introduction

During the early days of commercial cellular service, a few telecommunications providers, in collaboration with Mobile Network Operators (MNOs), used to provide cellular connectivity services in a unified manner, where a single entity offered all aspects of the mobile communications service. Examples of MNOs include T-Mobile, Vodafone, AT&T, or Telefónica, just to cite a few examples. As the market and technology progressed, new models surfaced, including the advent of Mobile Virtual Network Operators (MVNOs) [1], [2], which utilize the existing infrastructure of established MNOs to deliver services. Examples of MVNOs include Lebara and Lycamobile, companies that tend to offer competitive rates with respect to traditional MNOs.

More recently, a new type of service provider, known as Mobile Network Aggregators (MNAs), has emerged. Much like MVNOs, MNAs depend on the infrastructure of MNOs to deliver services. However, in contrast to relying on a single base MNO, MNAs have the capability to multiplex their clients across multiple MNOs. The driver for MNAs is then very different from that of MVNOs: instead of targeting low-cost connectivity, they are oriented at comprehensive international coverage, higher reliability and better performance. Indeed, MNAs establish agreements with MNOs located in different countries and, in case of redundant coverage by multiple MNOs, have the flexibility of potentially selecting the best performing MNO available at each time, hence ensuring optimal service and sustained Quality of Experience (QoE) without incurring additional costs associated with network operation. Representative examples of MNAs are Truphone and Google Fi.

The operations of MNAs heavily rely on roaming, a crucial service that facilitates the operations of MNAs in multiple countries, eliminating the necessity of seeking a local communication provider in each country where their end-users are active. MNAs have the possibility to either directly connect to a local MNO or to leverage on the extensive global network infrastructure established by international carriers, such as incumbent tier-one operators like T-Mobile, Telefónica or Tata.

In this paper, we provide an in-depth investigation on the operation and performance of MNAs, with a number of contributions as follows. First, we present the functional frameworks of MNAs, aligning them with established operational models for MNOs through a comprehensive classification system (see Section III). Then, we examine the significance of the roaming feature within the scope of MNAs that offer worldwide service, detailing how it is applied across various MNA models (see Section IV). We also delve into the specifics of our test procedures, offering a detailed overview of our systematic approach to examining the performance of various MNAs and MNOs. This includes a comprehensive explanation of our three-step comparison strategy, where we contrast the performances of different mobile service providers under various scenarios. We also shed light on the impact of the mobile device used, the chosen MNAs, and the local MNOs that are considered in the study (see Section V).

Next, we delve into the network infrastructure that facilitates the activities of two commercial MNAs, such as Google Fi and Truphone, (see Section VI). Our analysis concentrates on the influence of international operators on roaming, specifically regarding its effect on the performance experienced by end-users. This evaluation covers various factors, including the latency in Domain Name System (DNS) resolution and the

Sergi Alcalá-Marín and Marcelo Bagnulo are with University Carlos III of Madrid, Madrid, Spain. Weili Wu is with Georgia Institute of Technology, Georgia, US. Aravindh Raman is with Sky, London, UK. Ozgu Alay is with University of Oslo, Norway; Karslatd University, Karlstad, Sweden.Fabián E. Bustamante is with Northwestern University, Evanston, IL, US Marco Fiore is with IMDEA Networks Institute, Madrid, Spain. Andra Lutu is with Telefonica Research, Madrid, Spain.

performance of heterogeneous applications like web browsing and video streaming and juxtaposing findings relative to Europe and US (see Section VII). Our results show that, despite their promise and potential, MNA models are currently in their infancy, and still suffer from the impact of Home-routed Roaming (HR), similar to traditional MNOs [3] (see Section VII).

We finally leverage the 5G Control and User Plane Separation (CUPS) concept, and implement a realistic approach for global cellular operations that tackles some of the shortcomings of current roaming solutions. 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 Amazon Web Services (AWS) infrastructure, and present our results that demonstrate the performance benefit of this roaming implementation (see Section VIII).

The analysis conducted in this study reveals several important insights. First, we observe that MNOs with different operational models all experience similar performance on critical user-facing services—DNS resolution, web browsing, and video streaming—when using HR configurations. Second, MNAs like Google Fi and Truphone can reduce performance penalties by selecting a home network geographically closer to the user's visited country. Third, while local MNOs generally outperform MNAs in web performance, the differences in Page Load Time (PLT) are minor and do not significantly affect user experience. Fourth, in some video streaming scenarios, Truphone and Google Fi achieve better performance than local MNOs, possibly due to Internet Service Provider (ISP)-level throttling practices. These findings demonstrate that MNAs can offer robust and competitive services, sometimes even surpassing traditional MNOs, and suggest their potential to disrupt the mobile network market through flexible and efficient service delivery.

Table I: List of Acronyms

Acronym	Definition
AMF	Access and Mobility Management Function
AWS	Amazon Web Services
CDN	Content Delivery Network
CUPS	Control and User Plane Separation
DNS	Domain Name System
GTP	GPRS Tunneling Protocol
HMNO	Home Mobile Network Operator
HR	Home-routed Roaming
IHBO	IPX Hub Breakout
IMSI	International Mobile Subscriber Identity
IP	Internet Protocol
IPX	IP Packet Exchange
IPX-P	IPX Provider
IoT	Internet of Things
ISP	Internet Service Provider
LBO	Local Breakout
M2M	Machine to Machine
MCC	Mobile Country Code
MNA	Mobile Network Aggregator
MNC	Mobile Network Code
MNO	Mobile Network Operator
MVNO	Mobile Virtual Network Operator
PGW	Packet Data Network Gateway
PLT	Page Load Time
PoP	Point of Presence
QoE	Quality of Experience
QoS	Quality of Service
RBO	Regional Breakout
RTT	Round-Trip Time
SEPP	Security Edge Protection Proxy Network Function
SLA	Service-Level Agreement
SMF	Session Management Function
UE	User Equipment
UPF	User Plane Function
VPN	Virtual Private Network

II. Related Work

The research on MNAs and their impact on roaming performance and optimization builds upon a rich foundation of studies in mobile networks, global roaming, and emerging technologies to enhance mobile communication services.

Caushaj et al. [4] extensively evaluated throughput and delay in 3G and 4G mobile architectures for native devices. Their findings illustrate notable improvements in performance metrics as networks evolve from 3G to 4G technologies.

Geissler et al. [5] detailed the signaling traffic between the visited and home networks during roaming scenarios. Their research offers valuable insights into the complexities and demands of signaling in IoT mobile networks, which are critical for optimizing roaming performance and managing network resources efficiently.

Lutu et al. [6] characterized the global roaming support for an operational M2M platform, analyzing the impact of roaming IoT devices on the visited MNO. Their work highlights the unique challenges and requirements of supporting IoT devices in roaming situations, which differ significantly from traditional mobile devices. Ballal et al. [7]. further compare the performance of roaming IoT devices with those deployed in the home network, providing an in-depth analysis of how different cellular IoT technologies affect roaming performance

Schmitt et al. [8] analyzed the Round-Trip Times (RTTs) and route paths of four MNOs and MVNOs while accessing content from popular Content Delivery Networks (CDNs). Their study reveals significant differences in performance and behavior, offering insights that can inform the optimization strategies of global mobile network aggregators.

A significant contribution in this area is that by Vomhoff et al. [9], who explored the latency in global mobile roaming with a focus on regional breakouts within the IP Packet Exchange (IPX) framework (a hub service, to which MNOs connect over a private IP backbone network [10]). Their study investigates how these breakouts can potentially reduce latency and improve performance in mobile roaming scenarios.

Different from the literature above, our work aims to address a thorough analysis from the end-user perspective of the QoE of various MNAs. In an earlier version of this study [11], we carried out an initial exploration based in Spain and aimed at comparing the performance of a few target MNAs with that of local MNOs that have traditionally been the primary providers of mobile services as well as with a conventional MNO in roaming. In this extended version of the work, we present comparisons that cover different scenarios, including

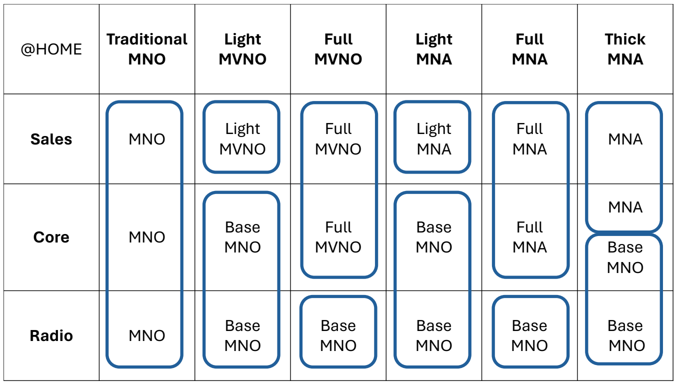

Figure 1: Types of MNOs.

both roaming and domestic traffic observed from multiple vantage points in both Europe and the US. As measuring international roaming performance is a challenging task that requires international cooperation and daunting logistics, the research presented in this paper provides a comprehensive picture of the performance landscape and reveals performance differences across MNAs that operate globally.

III. 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. Note that the terminology we adopt distinguishes mobile operators as a general concept and MNOs, i.e., one particular category of mobile operators along with MVNOs and MNAs. Specifically, an MNO is an entity that owns (or has the exploitation rights) of a cellular network infrastructure composed of spectrum licences, base stations, transport network, network core, etc. This was the initial operation model deployed to provide mobile communication services. Examples of MNOs include Vodafone, Orange, O2, Movistar, AT&T, NTT to name a few.

Later on, the MVNO operation model emerged [12]. The MVNO is an entity that offers mobile network services to end-users, but does not own or operate a full 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 of 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 [1], [13], [14], as follows.

- A *branded reseller* (also known as a ”skinny” MVNO) completely relies on the base MNO facilities to operate. They do not own any network elements, but they may operate their own commercial department.
- A *light MVNO* is a service provider that does not run its own core network, though it 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 *thick MVNO* partially manages its own core network deployment with its own infrastructure, which allows the MVNO more control; however, it still depends on a base MNO for some core network functions. These MVNOs have a heavier focus on branding, customer-ownership, and differentiation through added services like data and SIM applications.
- 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 [15]. While “traditional” 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 (now rebranded to 1Global), 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 [1], [13], [14] to include the MNAs. We further classify them into *full/thick/light* MNAs, depending on whether they operate their own complete/partial 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 IPX Provider (IPX-P) [10], and does not depend on a specific base MNO (e.g., the MVNO might use global International Mobile Subscriber Identity (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 in 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 about how these different MNA models satisfy the need for global coverage for their end-users. We further provide background on roaming, one of the fundamental functions mobile operators ensure to their end-users.

IV. 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 [10].

A. Roaming Background

MNOs commonly connect with each other through an IPX network. As mentioned in Section II, an IPX [10] 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 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** [16], 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 User Plane Function (UPF) of the visited network and the UPF of the home network (solid red path in the figure) for the Data plane connection and Security Edge Protection Proxy Network Function (SEPP) of the visited network and the home network for the Control plane connection. With the IP endpoint in the home network, all services will be available in the same way as in the home network.

When **Local Breakout (LBO)** [16] is used, the IP address of the roaming user is provided by the visited network. The GTP tunnel is terminated at the UPF 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) [10], or regional breakout, 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 UPF in the visited network terminates at a UPF in the IPX network (green path in the figure). There may be multiple UPFs 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 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 [3]. 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 CDN when roaming abroad, because the mobile user will access a server in the home network rather than one close to their location.

B. Mix-and-match: MVNO and MNA Roaming

In this section, we discuss and analyze the global operations for the different types of (virtual) operators we introduced in Section III (see Figure 3). Previous work [3] 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.

Thick MNAs, however, bring more complexity to the landscape: they only operate specific network functions (e.g., they operate only the Packet Data Network Gateway (PGW)/UPF over third-party infrastructure) and they are able to leverage

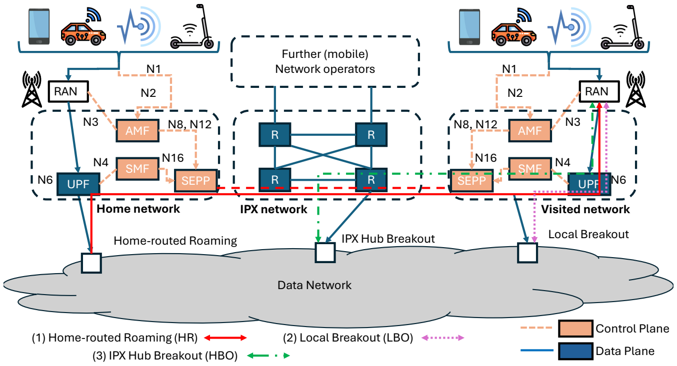

Figure 2: Internet access options when roaming home-routed roaming, local breakout, and IPX hub breakout roaming configurations.

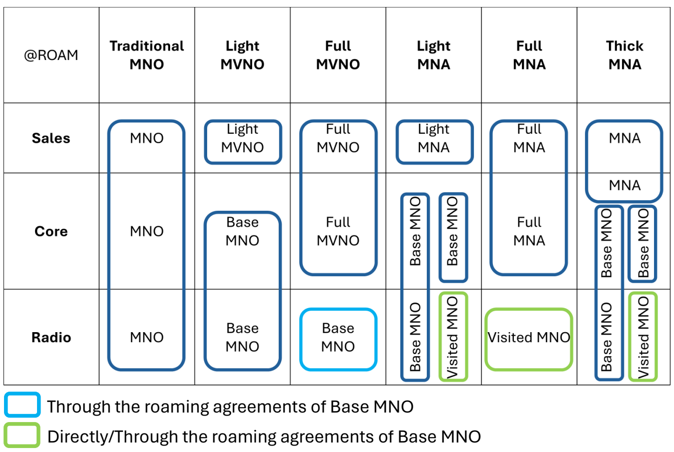

Figure 3: Roaming operations for the different types of mobile operators.

the eSIM technology as well. Providers like Airalo or HolaFly function as thick MNAs, and they partner with few base MNOs to sponsor the eSIMs that travelers can use world-wide, either in roaming or native configuration.

V. 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 US, and we perform the experiments while roaming in different countries in Europe (Spain) and the US (Atlanta, GA). For Europe, we use three different MNAs, namely Google Fi, Truphone and Twilio. For the US, we also characterize the performance of Google Fi and Truphone. Also, we will run the experiments using the local MNOs that have partnership with the aforementioned MNAs. Later in the document, we will provide a detailed explanation for our choice of also using KPN (Netherlands), elucidating the reasons behind this selection within the context of our study. Fi is Google's MNA service.

Fi is only available for US customers. This means that we must activate any new Fi account in the US, and, only after that, we can use it internationally.[1] 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",[2] which in most cases only include the US 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[3]. 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 [17]. While the selection of the base MNO is automatic, it is possible to force the the MNO used by Fi using dialer codes [18]. 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 network with core network technology distributed across four continents. Truphone uses these local Point of Presences (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 [19]. Twilio's Super SIM can connect to 343 networks in 174 countries [20] 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 US 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) [21].

To perform the measurements in Europe, we subscribed to the three aforementioned MNAs services in the US, and ran experiments while roaming in Europe (Spain). For the US measurements, we also subscribe the two most impactful MNAs while roaming in US. For end-user equipment, we use Pixel 4A (US version) mobile phones in all experiments. In

[1] Activate Google Fi service: https://support.google.com/fi/answer/6078618?co=GENIE.Platform\%3DiOS&hl=en&oco=0

[2] Fi supported phones: https://fi.google.com/about/phones/.

[3] See the answer to "What is unique about Google Fi's network?" in the FAQ: https://fi.google.com/about/faq/.

addition to the MNA services, we also subscribe to a local MNO, namely, Orange (Spain), T-Mobile (US) and AT&T (US) , 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). Finally, we subscribed to KPN (Netherlands), since our subsequent experiments in the US reveal that Truphone's e-SIM configurations are based in the Netherlands, and they have chosen KPN as base MNO. 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 II 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 (EU target), and 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 PLT for web pages hosted by a server located in the US (home), Spain (visited country), Belgium (EU target), and one hosted by a CDN.
- Video performance: We measured the average quality, the number of rebuffering events and the bandwidth obtained while streaming.

We selected Belgium as regional EU target because our experiments in Spain revealed that Google Fi's e-SIMs breakout in Belgium. We selected the CDN-hosted website (mit.edu) to represent globally distributed content, helping us to evaluate the interplay of roaming configurations and the CDN.

VI. 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 II. 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 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 VII.

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 [22].

By extending the scope of our analysis to include traceroutes from the US, we aim to not only corroborate our previous findings but also to deepen our understanding of the infrastructure and roaming configurations specific to the US context. To ensure consistency and comparability with the previous study, we employ the same set of target destinations consisting of web pages hosted on servers in Europe and the US. This set includes umich.edu[4] as a representative target in the US, along with uclouvain.be (Belgium) and url.edu (Spain) as targets in Europe. Additionally, we consider a web page, mit.edu served by a CDN. We repeat the measurements for each experimental end-user SIM configuration described in Table III. It is worth noting that the difference in Google Fi between both scenarios (Europe and US) is that in Europe we were able to connect to a regional breakout (e.g. 3 Network) and the partner T-Mobile whereas in US we only connect to T-Mobile. Similar to the previous study, we perform 20 traceroute measurements for each target from the US end-users, compute all available routing paths and record the minimum RTT values observed for each hop. The geolocation of each hop is inferred using reverse DNS information, WHOIS and MaxMind geolocation database, following the methodology established in the previous study. The limitation of these approaches will be discussed later. By conducting this comparative analysis, we aim to gain insights into the potential variations in the infrastructure and performance characteristics of MNAs when serving domestic users in the US compared to international roaming users in Europe. In the subsequent subsections, we will conduct a detailed analysis and comparison of parallel experiments. These include tracing routes from Europe to the United States and vice versa (Sec.VI-A), within Europe, and within the U.S. (Sec.VI-B), as well as to CDNs (Sec.VI-C).

A. Traceroute EU-US and US-EU

Figure 4e 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 II).

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 ($\approx$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

[4]In our previous paper, we used ucla.edu as an US domain, but at the day of performing experimentation in the US, ucla.edu runs in a CDN. Thus, we found another website with the same features that ucla.edu had in 2022

Table II: 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

Table III: Experimental configurations: MNAs, base MNOs and visited MNOs we used for measurements in the US. 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)	N/A	VPN enabled	US
Fi/TM/noVPN	Fi	T-Mobile (US)	N/A	VPN disabled	US
Truphone	Truphone	KPN (NL)	AT&T (US)		Europe
KPN	KPN	KPN(NL)	AT&T (US)		Europe
T-Mobile	T-Mobile	T-Mobile (US)	N/A	Baseline (US)	US
AT&T	AT&T	AT&T (US)	N/A	Baseline (US)	US

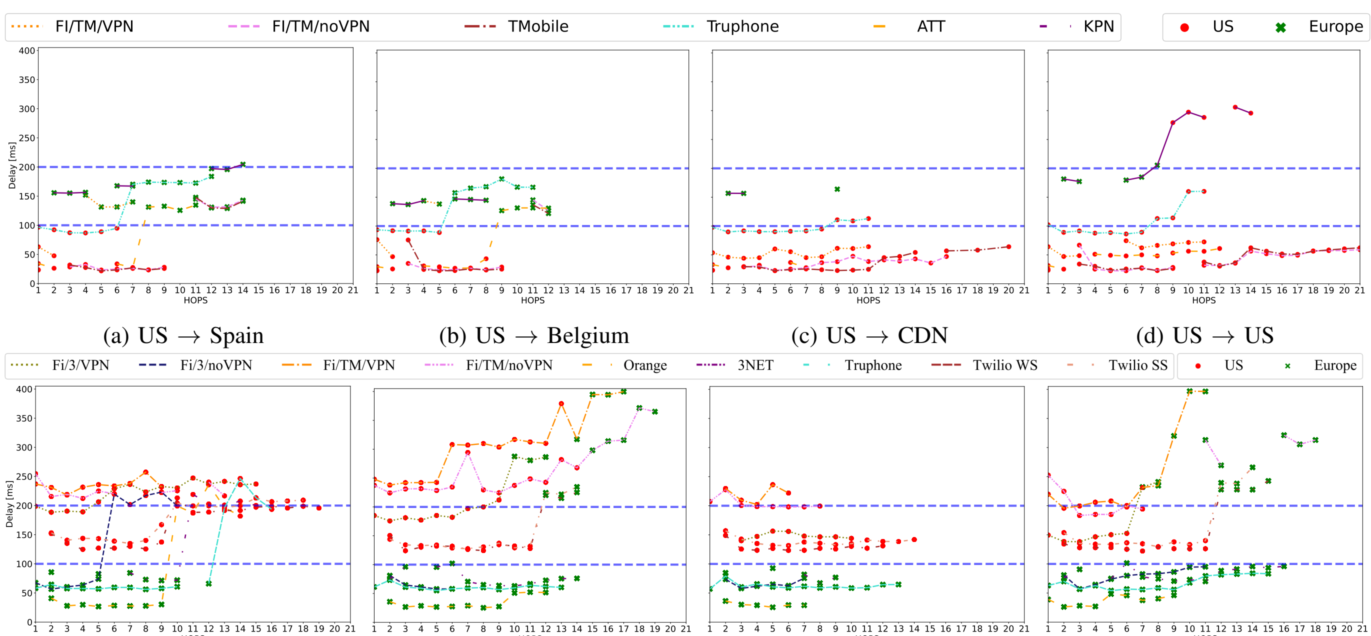

(a) US → Spain (b) US → Belgium (c) US → CDN (d) US → US

(e) Spain → US (f) Spain → Belgium (g) Spain → CDN (h) Spain → Spain

Figure 4: Hops encountered with the latency towards targets in US, Belgium, Spain and within a CDN

(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.

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 Figure 4e). 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. When the VPN is enabled, the traffic is routed to Google in the US and then to 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 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) [19], 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).

Figure 4a and 4b show the delays for each hop replying to the traceroute probes from a mobile device (roaming) in the US towards a server (url.edu) located in Spain and (uclouvain.be) in Belgium. We can observe that the overall delay to the target varies between 150 ms and 200 ms (20% variation) across the different SIM configurations (Table III) for the Spanish server and between 120ms and 180 for the Belgium Server. We observe again the transatlantic link, in this case for all SIM configurations but not KPN. This is because KPN already has the first hop in Europe, inducing a high penalty already in the first hop. The large delay value we measure in the first hop (150ms) , and the inferred geo-location of the IP address of the first hop in the Europe both corroborate our deduction. For the server located in Spain, KPN has a delay on the last hop of $\approx 200ms$ and for the belgian server, 150ms , this is due to having the server located very close to the breakout point. However, when using the T-Mobile or AT&T subscription (i.e., the visited MNOs that Fi and Truphone attaches to in the U.S,respectively), we find that the first hop geolocates in the US (the delay is $\approx$ 30 ms for T-Mobile and $\approx$ 40ms for ATT, and the geolocation of the IP of the first hop maps to US). We can easily identify the transatlantic link later in the path due to the large increase in the delay ($\approx 100ms$), and the fact that the IP geolocation shifts from the US to Spain.

When using GF/TM/noVPN, we observe that it has the same hops and almost same latencies for both servers than TMobile, indicating that has the same routing for the packets than TMobile. In the case of GF/TM/VPN, the VPN tunneling induces to a complete diferent path of hops for the packets, having a higher first hop delay (aprox 70ms) but having an approximate delay in the last hop, but having way less hops compared to GFI/TM/noVPN and T-Mobile (7 vs 14 hops for spain and 5 vs 12 hops for belgium). For Truphone, we observe that the eSIM configuration belongs to its base MNO, KPN, but it has the breakout point in the US, since it first hop is in the US. It is worth nothing that its first hop, even being in the US, is the one with higher latency compared with GFIs and the benchmarks ($\approx 100ms$). We can also identify the transatlantic link later in the path due to the increase in the delay ($\approx 150ms$).

B. Traceroute EU-EU and US-US

We next analyze the traceroute results we collect from measuring towards a server located in Belgium (uclouvain.be), which we show in Figure 4f. 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 4h. 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 4h (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). In the following, we study the traceroute results we collect from measuring towards a server located in the US, which we show in Figure 4d. We observe that for GFI/TM/VPN , GFI/TM/noVPN, TMobile and

ATT the delay is significantly lower (between 50 ms and 75 ms) than for Truphone and KPN. It is the same behaviour that we observed in the EU-EU Traceroute, where Orange, Truphonem and Fi/3/noVPN, having the breakout in Europe, had the lowest latency compared with the other SIM configurations. For Truphone, we observe that it already starts with 100ms of delay in the first hop, and it increases until aprox 180ms. we conclude that this is happening because Truphone's breakout is in the south west coast and the measurements were taken in middle-east of US and we are acessing to umich.edu (Michigan). Finally, for KPN, we observe the same behaviour as FI/TM/noVPN, FI/TM/VPN, and Twilio setups, having a notably big first hop delay (aprox 180ms) and achieving 300ms. Again, this behaviour is observed due to HR setup for roaming traffic.

C. 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 4g. 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.

In the Measurements done in the US, we observe that ATT is the one with the shortest delay (30ms), followed by GF/TM/noVPN, Tmobile and GF/TM/VPN, accessing all of them so the closest CDN from where the measurements were taken. Next, we find Truphone that has a smiliar delay in the first and in the last hop (90-110ms), since it is accessing to the closest CDN in the west coast. Regarding KPN, we also observe the same performance than FI/TM/VPN FI/TM/noVPN and Fi/3/VPN in the prior measurements, the end-user is accessing to a replica in the Netherlands, resulting in a high overall delay.

D. Takeaways

MNOs usually deploy the HR configuration for international roaming (Section IV). 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 [23], [3]. We observe a similar behavior for the three target MNAs. The difference, however, comes from the capability of the MNA to change their base MNO (nationally and internationally), thus implicitly adapting 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 experiences when 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.

When measuring from the US, we observe a very similar performance with its base MNO, namely, T-Mobile, in both configurations (i.e. with and without VPN). As it is a service that you can only subscribe within the US, it makes sense that they have a better performance in the US in all configurations than in Europe if the packets should be routed to the HR. We observe that having the aforementioned Regional Breakout, the first hop has 4 times longer delay than doing HR (20 ms $GFI/TM/noVPN_{US}$ vs 80 ms delay $GFI/Three/noVPN_{Europe}$) for all the accessed domains. Instead, accessing a domain within Europe uses a regional breakout for the last hop, cutting delays dramatically (127 ms $GFI/TM/noVPN_{US}$ vs 58 ms delay $GFI/Three/noVPN_{Europe}$) and demonstrating the huge advantage that Regional Breakout has for users in roaming.

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. From the measurements performed in the US, we find that the breakout point is in the west-coast of the US, having a worse performance than GFI but still improving its base MNO, namely, KPN, that is using HR. Also, Truphone was worse performance than ATT, its visited MNO. We thus confirm that Truphone uses their POPs to implement the Regional Breakout approach to reduce the distance that mobile traffic has to travel. We observe that this indeed reduces the delay.

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.[5] 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.

VII. 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.

A. DNS resolution delay

We measure the DNS resolution delay using the different experimental configurations in Table II for Europe and Table III for the US 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. One of the authoritative servers is located in Spain, serving the worst case for roaming devices as we inferred from the traceroute experiments (see Figure 4). For the US measurements, we also configured a wildcard DNS resolver using Amazon Web Services (AWS) on the Amazon S3 platform located in the US. We run 20 queries for each experimental configuration of the end-user SIM.

In Figure 5a, 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. 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.[6] 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

[5]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.

[6]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.

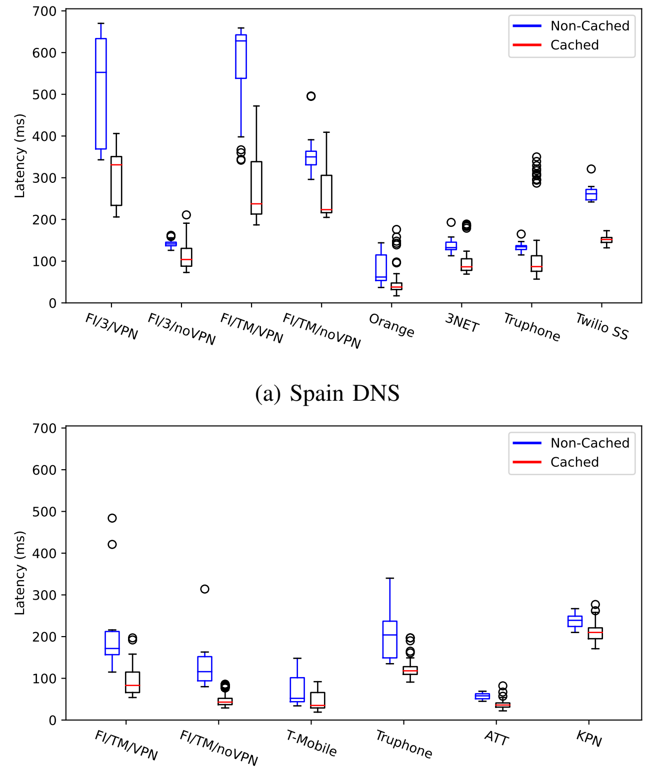

(a) Spain DNS

(b) US DNS

Figure 5: DNS resolution delays we measured for a non-cached and a cached DNS name in (a) Spain and (b) US, using different experimental configurations for the end-user SIM (see Table II) and Table III.

over 550 ms) and also a high variance. Fi over T-Mobile without VPN follows, with a mean delay of over 350ms. 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 US 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) changes to the one of the local operator Truphone registered in Spain.

In Figure 5b we show the delays for the US measurements, for the resolution delay for a non-cached name, we observe that ATT and T-Mobile, as expected, have the mean lowest resolution delay, i.e. 50 ms. Then, we find that FI/TM/VPN and FI/TM/noVPN have a similar behaviour, with some variance, but with a mean delay of 170 ms and 120 ms, respectively. As observed in Figure 4, GFI/TM/VPN has some delay incurred by the VPN tunneling, producing higher resolution delays. Next, we find Truphone, that has a higher resolution delay mean with some variance, i.e. 200 ± 50 ms. This is due to have the breakout point in the west coast as commented in (see Sec. VI). Finally, the remaining configuration, namely, KPN, has a mean value of 240 ms. We see again how the HR roaming imposes a penalty of 100-200 ms compared to LBO roaming configuration. For this exoeriment, we can compare the performance of KPN (EU MNO accessing to US domain from the US) with FI/3/VPN, FI/TM/VPN, FI/TM/noVPN and Twilio's (US services accessing to Spanish domain from the US), and we observe that KPN has a dramatically smaller DNS resolution delay for the Non-cached domains.

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: www.amazon.com, www.facebook.com and www.youtube.com.

In Figure 5a, we compare the DNS resolution delay for a cached name with that of non-cached ones for Spain measurements. 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.

In Figure 5b, we compare the DNS resolution delay for a cached name with that of non-cached ones for Spain measurements. Again, we follow the same procedure as in the measurements taken in Spain, doing a first query for a domain name to populate the resolver's cache, cleaning the DNS client cache and query for the same domain name to measure the DNS resolution time. Overall, the performance we observe is relatively similar to the both (cached- non cached measurements in Spain) and non-cached in the US, that in absolute values are smaller. For T-Mobile and ATT, the mean resolution delays are in the order of 40 ms. These are followed by the both FI Configurations (noVPN / VPN) with around 60 ms and 80 ms, respectively. For Truphone has a delay of 110 ms and KPN of 200ms. Again, in the measurements in the US the overall resolution delay was lower, because the Regional Breakout was in the same country. What we observe is that, for Truphone, the geolocation of the POP plays a role in the latency of the packets and resolution delay, even being in different countries in the case of Europe, it has a lower resolution delay compared to the one in the US, having twice the distance to the breakout point (ES-NL is 1400 km and US middle-west is 3000 km).

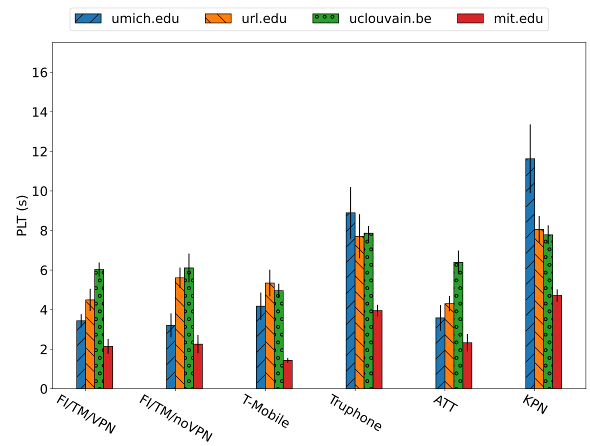

(a) Page Load Time for each configuration US

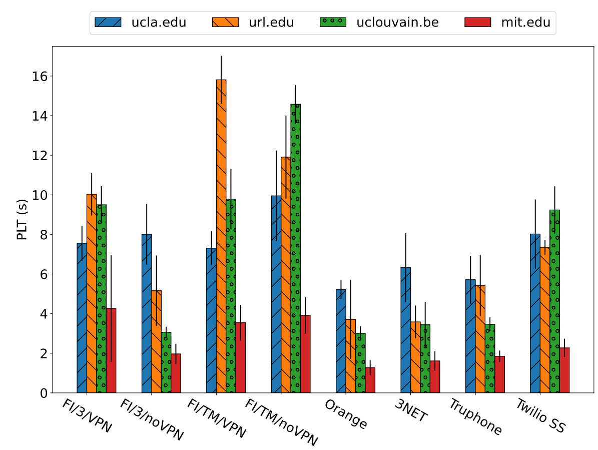

(b) Page Load Time for each configuration ES

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 (a) Spain and (b) US with different SIMs (see Table II and Table III. For the web pages located in US, in (a), we use ucla.edu and in (b), we use umich.edu.

B. 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 [24], available in the 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 [25] 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 II and Table III), while targeting different web pages. We selected web pages of roughly the same size (with a 5% range) in different geographic locations For the measurements done in Spain, we select www.ucla.edu in the US, www.uclouvain.be in Belgium, www.url.edu in Spain and www.mit.edu which is served by a CDN. For the measurements done in the US, we use the same web pages, but changing the one in the US because at the day of the extension, www.ucla.edu has moved to a CDN. Thus, we select www.umich.edu as web page in the US for the experiments in the US. 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.

For the US measurements, we observe analogous trends from the ones obtained in Spain. In the context of this experiment, the US functioned as both the home and visited country for Google Fi and the local MNOs. However Truphone had a different home and visited country, which might have contributed to its comparatively larger latency. In this case, the penalties when accessing to a web server in Belgium are slightly large than the ones for Spain and followed by the server in the US, where the global penalty observed is substantially lower. Again, for the CDN web server, the overall PLT is the smallest in all cases, and the configuration with the least one is T-Mobile. We observe a similar experience when using GFI wtih and without VPN to T-Mobile, even having a marginally better experience when accessing to the web server located in the US (i.e. 0.5s faster). This is something that will not change the QoE of an end-user. When we use Truphone, we observe the largest PLT of the MNA configurations, having an even larger PLT for accessing the web server located in the US than the ones located in Europe. When roaming in the US (KPN configuration) we observe the clear effect of HR roaming: When accessing to the US server, the penalty obtained is much larger than the one in Europe, having a slightly difference between Belgium and Spain (being Belgium smaller because it is closer to the Breakout point).

C. 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 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 II, except for Twilio. We omitted Twilio from these tests because their products are oriented to Internet of Things (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[7], a music video[8] and a movie trailer[9], 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 7b 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.

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

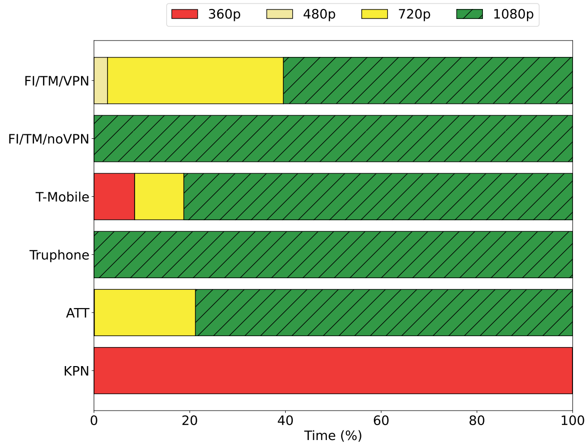

(a) Video performance for each configuration US

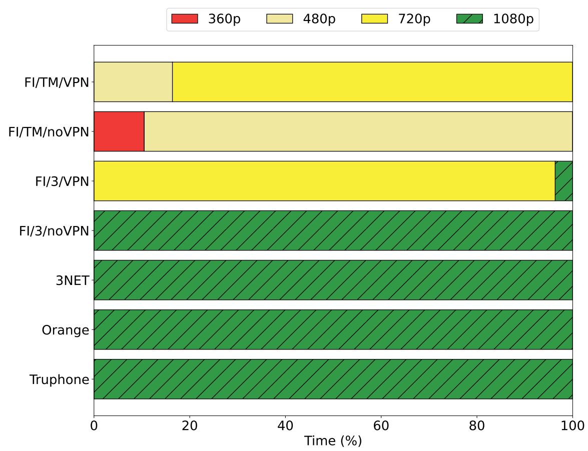

(b) Video performance for each configuration Spain

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

We again confirmed the result of the previous experiment. We observe that for Truphone, and Google Fi without VPN, the videos are mostly rendered in the highest resolution, and for Google Fi with VPN, the most common resolution is 720p. The good performance of Truhpone counter the previous results from DNS and web experiments, where Truphone was the worse configuration. It proves that depending on the use case, the QoE can change. TMobile and ATT are outperformed by Truphone and Google Fi without VPN. These findings contradict the expectation in the sense that MNOs should have better performance than MNAs who leverage their infrastructures to provide services, as is confirmed in the web experiments. We attribute this anomaly to ISP-level throttling. Previous research has uncovered instances of US ISPs throttling multimedia streaming services. To investigate this, we utilized WeHe [27], a VPN-based traffic differentiation detection tool. Our test results revealed that T-Mobile and AT&T were indeed throttling YouTube video playback, while Truphone and Google Fi showed no signs of such behavior.

In conclusion, the exceptional performance of Truphone and Google Fi suggests that US MNOs can indeed provide superior services to subscribers. Given that Google Fi mainly relies on T-Mobile, and Truphone primarily depends on AT&T in the area where we conducted our experiments, the performance discrepancy between these MNAs and their partner MNOs indicates that the MNO's does not classify the traffic generated by all the MNAs in our study.

D. 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.

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. When operating in the US, we observe a very similar performance to its MNO, namely T-Mobile, having the big difference when using video application, where Google Fi was able to outperform T-Mobile.

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. From the US perspective, we observed that Truphone was not able to achieve the same performance than in Spain. This is due to the PoP in the US, located in the middle-west and the domain analyzed located in the west coast. Again, for the Video measurements, it outperformed its benchmark MNO, namely AT&T, due to ISP-level Throttling on the MNO side.

As a conclusion, our experiments yielded insightful results concerning the performance of the analysed MNAs and

MNOs. This demonstrates the potential of MNAs to offer robust and efficient services that are on par with or even surpass those of established MNOs in certain aspects. In terms of web performance, our study revealed some interesting trends. The local MNOs often outperformed the MNAs, with the latter exhibiting slightly longer PLT. However, it is noteworthy that the differences in PLT were not so substantial as to impact the end-user experience severely. This suggests that MNAs can deliver satisfactory web performance despite operating on top of existing MNO infrastructures. The video performance tests provided some unexpected results. Despite the MNAs leveraging the infrastructure of MNOs, Truphone and Google Fi without VPN demonstrated superior video performance in certain scenarios. This could potentially be attributed to ISP-level throttling practices by certain MNOs. These findings underscore the importance of examining network performance under diverse real-world conditions, as well as the need for more transparent network management policies. While MNOs often exhibited superior performance metrics in our tests, MNAs proved to be robust alternatives that can offer competitive services. This reveals the potential of MNAs to disrupt the traditional mobile network market by offering flexible and efficient services.

VIII. Regional Breakout in 5G

The split of user/data plane from the control plane at both the radio front [28], [29] and the core [30] 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.

To explore this vision, we implement a proof-of-concept 5G core using CUPS across multiple AWS regions. This setup demonstrates how roaming users can experience improved latency and performance by dynamically relocating UPFs closer to their current location. Our work mirrors emerging commercial trends among MNAs, such as Twilio and emnify, which are leveraging regional breakout architectures on cloud platforms to overcome the latency penalties of HR [31], [11]. This section presents our deployment architecture and evaluates the performance implications of regional breakout in a roaming scenario, showing how it can approximate the quality of service users experience in their home networks.

A. Experimental setup

We utilise the edge (wavelength), local and regional deployments of Amazon global infrastructure [32] to deploy control and user plane functions of open-source 5G implementations (namely, Open5GS [33] and UERANSIM [34]).

Infrastructure. The setup we build aims to emulate different roaming configurations (see Section IV). 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 [33]. 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 IV). The selection of UPF to breakout is chosen by the DNN setting in the 5G User Equipment (UE). We use UERANSIM [34] 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 IV.

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 [35] and through virtual private cloud peering connections [36]. We show ping latency within the instances in the US and to UK (home) in Table V.

B. Pilot Results

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

Table IV: Locations of deployed 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 V: 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

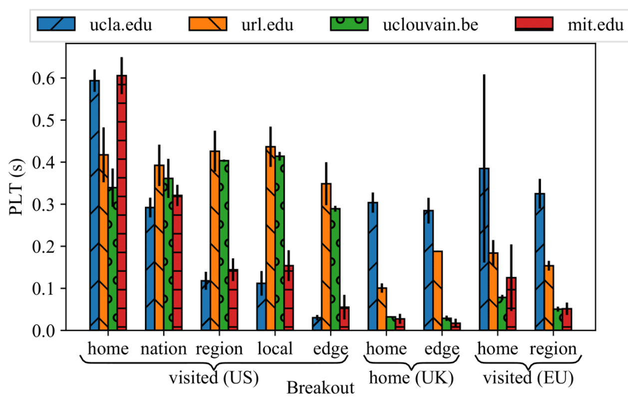

Figure 8: Page load time to different websites in function of the breakout point considered for the roaming configuration.

in the US, www.uclouvain.be in 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-user 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 result 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 V).

Our experimental findings support the growing consensus that regional breakout is a viable approach for improving roaming performance in 5G networks. The ability to instantiate UPFs across geographically distributed cloud regions allows operators to minimize latency while retaining centralized control over authentication and policy. These results align with large-scale measurement studies showing that regional breakout configurations can significantly reduce RTTs compared to HR, in some cases by over 70% [37]. As 5G deployments continue to evolve, cloud-based user plane relocation offers a practical and scalable path for enabling performance-aware, location-sensitive mobile services.

IX. Conclusions

The mobile communications industry is constantly evolving. Recently, a new type of global virtual operator, known as MNA, has emerged. In our research paper, we introduced a way to analyze the operations of MNAs and consider their potential development. We expanded the existing taxonomy of MVNOs to include the concept of MNA.

Our research involved measuring the roaming operations of three MNAs with different operational models: Twilio, Truphone, and Fi from Europe and the US. We evaluated their performance and measured the impact of their operational approach on various applications such as DNS, web browsing, and video streaming.

We discovered that, unlike MNOs and MVNOs which rely on HR roaming, MNAs are gradually adopting limited forms of LBO roaming. While current MNA operations may not enable LBO in the visited country, some implement regional breakout to keep traffic within the same continent/region and avoid long transoceanic links. This difference is clearly seen in the performance of the various applications we tested.

Our goal was to understand the operational models of MNAs and their impact on application performance, rather than to compare the different MNOs and MNAs tested. We also observed how the location of the breakout point significantly affects the end-user's QoE and how MNAs can outperform MNOs depending on their network configuration for certain applications.

Furthermore, we investigated the potential evolution of the MNA model and explored the performance gains that could result from fully utilizing LBO. Depending on the application, regional breakout can offer significant benefits, and there are limited additional advantages from implementing a full LBO.

Acknowledgment

This work has been partially supported by the ORIGAMI project, which has received funding from the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme with Grant Agreement No. 101139270. The work of Marcelo Bagnulo has been partially funded by the EU through the StandICT RXQ project.

References

[1] P. Schmitt, M. Vigil, and E. Belding, "A study of mvno data paths and performance," in *International Conference on Passive and Active Network Measurement*. Springer, 2016, pp. 83–94.

[2] Y. Li, J. Zheng, Z. Li, Y. Liu, F. Qian, S. Bai, Y. Liu, and X. Xin, "Understanding the ecosystem and addressing the fundamental concerns of commercial mvno," *IEEE/ACM Transactions on Networking*, vol. 28, no. 3, pp. 1364–1377, 2020.

[3] A. M. Mandalari, A. Lutu, A. Custura, A. Safari Khatouni, O. Alay, M. Bagnulo, V. Bajpai, A. Brunstrom, J. Ott, M. Mellia, and G. Fairhurst, "Experience: Implications of roaming in europe," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '18. New York, NY, USA: Association for Computing Machinery, 2018, pp. 179–189. [Online]. Available: https://doi.org/10.1145/3241539.3241577

[4] E. Caushaj, I. Ivanov, H. Fu, I. Sethi, and Y. Zhu, "Evaluating throughput and delay in 3g and 4g mobile architectures," *Journal of Computer and Communications*, vol. 02, pp. 1–8, 01 2014.

[5] S. Geissler, F. Wamser, W. Bauer, M. Królikowski, S. Gebert, and T. Hossfeld, "Signaling traffic in internet-of-things mobile networks," *2021 IFIP/IEEE International Symposium on Integrated Network Management (IM)*, pp. 452–458, 2021. [Online]. Available: https://api.semanticscholar.org/CorpusID:235420498

[6] A. Lutu, B. Jun, A. Finamore, F. E. Bustamante, and D. Perino, "Where Things Roam: Uncovering Cellular IoT/M2M Connectivity," in *Proceedings of the ACM Internet Measurement Conference*, ser. IMC'20. New York, NY, USA: Association for Computing Machinery, 2020, pp. 147 – 161.

[7] K. D. Ballal, R. Singh, L. Dittmann, and S. Ruepp, "Experimental evaluation of roaming performance of cellular iot networks," *2022 Thirteenth International Conference on Ubiquitous and Future Networks (ICUFN)*, pp. 386–391, 2022. [Online]. Available: https://api.semanticscholar.org/CorpusID:250939426

[8] P. Schmitt, M. Vigil, and E. M. Belding-Royer, "A study of mvno data paths and performance," in *Passive and Active Network Measurement Conference*, 2016. [Online]. Available: https://api.semanticscholar.org/CorpusID:16837847

[9] V. Vomhoff, M. Sichermann, S. Gieißler, A. Lutu, M. Giess, and T. Hoßfeld, "A shortcut through the ipx: Measuring latencies in global mobile roaming with regional breakouts," in *2024 8th Network Traffic Measurement and Analysis Conference (TMA)*, 2024, pp. 1–10.

[10] A. Lutu, B. Jun, F. E. Bustamante, D. Perino, M. Bagnulo, and C. G. Bontje, "A first look at the ip exchange ecosystem," *SIGCOMM Comput. Commun. Rev.*, vol. 50, no. 4, p. 25?34, Oct. 2020. [Online]. Available: https://doi.org/10.1145/3431832.3431836

[11] S. Alcalá-Marín, A. Raman, W. Wu, A. Lutu, M. Bagnulo, O. Alay, and F. Bustamante, "Global mobile network aggregators: taxonomy, roaming performance and optimization," in *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services*, 2022, pp. 183–195.

[12] F. Zarinni, A. Chakraborty, V. Sekar, S. R. Das, and P. Gill, "A first look at performance in mobile virtual network operators," in *Proceedings of the 2014 conference on internet measurement conference*, ser. IMC '14, 2014, pp. 165–172.

[13] A. Xiao, Y. Liu, Y. Li, F. Qian, Z. Li, S. Bai, Y. Liu, T. Xu, and X. Xin, "An in-depth study of commercial mvno: Measurement and optimization," in *Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys '19, 2019, pp. 457–468.

[14] N. Vallina-Rodriguez, S. Sundaresan, C. Kreibich, N. Weaver, and V. Paxson, "Beyond the radio: Illuminating the higher layers of mobile networks," in *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys '15, 2015, pp. 375–387.

[15] Z. Yuan, Q. Li, Y. Li, S. Lu, C. Peng, and G. Varghese, "Resolving policy conflicts in multi-carrier cellular access," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '18. New York, NY, USA: Association for Computing Machinery, 2018, pp. 147–162. [Online]. Available: https://doi.org/10.1145/3241539.3241558

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

[17] T. Kugler, "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.

[18] "Google Fi dialer codes," https://arkienet.com/2018/01/google-fi-dialer-codes/, accessed: 2021-12-10.

[19] "Twilio super sim," https://www.twilio.com/docs/iot/supersim, accessed: 2021-12-10.

[20] "Twilio super sim: Available networks," https://www.twilio.com/docs/iot/supersim/available-networks, accessed: 2021-12-10.

[21] "Programmable wireless: Gain global cellular iot connectivity," https://www.twilio.com/iot/wireless, accessed: 2021-12-10.

[22] M. Gharaibeh, A. Shah, B. Huffaker, H. Zhang, R. Ensafi, and C. Papadopoulos, "A look at router geolocation in public and commercial databases," in *Proceedings of the 2017 Internet Measurement Conference*, ser. IMC '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 463–469. [Online]. Available: https://doi.org/10.1145/3131365.3131380

[23] A. Lutu, M. Trevisan, A. Safari Khatouni, A. M. Mandalari, A. Custura, M. Mellia, O. Alay, M. Bagnulo, V. Bajpai, A. Brunstrom, J. Ott, and G. Fairhurst, "Measuring Roaming in Europe: Infrastructure and Implications on Users QoE," *IEEE Transactions on Mobile Computing*, pp. 1–1, 2021.

[24] M. contributors, "Navigation timing api," in *MDN Web Docs*, 2020, last accessed in 23/6/2021.

[25] M. Rajiullah, A. Lutu, A. S. Khatouni, M.-R. Fida, M. Mellia, A. Brunstrom, O. Alay, S. Alfredsson, and V. Mancuso, "Web experience in mobile networks: Lessons from two million page visits," in *The World Wide Web Conference*, 2019, pp. 1532–1543.

[26] F. Wamser, M. Seufert, P. Casas, R. Irmer, P. Tran-Gia, and R. Schatz, "Yomoapp: A tool for analyzing qoe of youtube http adaptive streaming in mobile networks," in *2015 European Conference on Networks and Communications (EuCNC)*. IEEE, 2015, pp. 239–243.

[27] F. Li, A. A. Niaki, D. Choffnes, P. Gill, and A. Mislove, "A large-scale analysis of deployed traffic differentiation practices," in *Proceedings of the ACM Special Interest Group on Data Communication*, 2019, pp. 130–144.

[28] X. Foukas, N. Nikaein, M. M. Kassem, M. K. Marina, and K. Kontovasilis, "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*, ser. CoNEXT '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 427?441. [Online]. Available: https://doi.org/10.1145/2999572.2999599

[29] X. Foukas, M. K. Marina, and K. Kontovasilis, "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*, ser. MobiCom '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 127?140. [Online]. Available: https://doi.org/10.1145/3117811.3117831

[30] I. Quintana-Ramirez, A. Tsiopoulos, M. A. Lema, F. Sardis, L. Sequeira, J. Arias, A. Raman, A. Azam, and M. Dohler, "The making of 5g: Building an end-to-end 5g-enabled system," *IEEE Communications Standards Magazine*, vol. 2, no. 4, pp. 88–96, 2018.

[31] V. Vomhoff, H. D. Jang, M. Varvello, S. Geißler, Y. Zaki, T. Hoßfeld, and A. Lutu, "Challenges and opportunities for global cellular connectivity," 2024. [Online]. Available: https://arxiv.org/abs/2411.19706

[32] "AWS global infrastructure," https://aws.amazon.com/about-aws/global-infrastructure/, accessed: 2021-12-10.

[33] "Open5gs v2.3.6," https://github.com/open5gs/open5gs, accessed: 2021-12-10.

[34] "Ueransim v3.2.4," https://github.com/aligungr/UERANSIM, accessed: 2021-12-10.

[35] L. Corneo, M. Eder, N. Mohan, A. Zavodovski, S. Bayhan, W. Wong, P. Gunningberg, J. Kangasharju, and J. Ott, *Surrounded by the Clouds: A Comprehensive Cloud Reachability Study*. New York, NY, USA: Association for Computing Machinery, 2021, p. 295–304. [Online]. Available: https://doi.org/10.1145/3442381.3449854

[36] "Amazon virtual private cloud," https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html, accessed: 2021-12-10.

[37] V. Vomhoff, M. Sichermann, S. Gieißler, A. Lutu, M. Giess, and T. Hoßfeld, "A shortcut through the ipx: Measuring latencies in global mobile roaming with regional breakouts," in *2024 8th Network Traffic Measurement and Analysis Conference (TMA)*, 2024, pp. 1–10.

Sergi Alcalá-Marín is a Research Scientist at University Carlos III of Madrid. He received his M.Sc. (2020) in Advanced Telecommunications Technologies from the Universitat Politècnica de Catalunya, Spain, and his Ph.D. (2025) in Telematics Engineering from the University Carlos III of Madrid and IMDEA Networks Institute. His research interests include wireless communications, network automation, and AI-driven network optimization.

Weili Wu received the Master of Science degree in Computer Science from Northwestern University, Evanston, IL, USA. He is currently pursuing the Ph.D. degree in Computer Science at the Georgia Institute of Technology, Atlanta, GA, USA. His research interests include Internet measurement, cybersecurity, and big data analytics.

Aravindh Raman is an Internet measurement researcher at Cisco ThousandEyes. His work focuses on identifying bottlenecks across the Internet, from access networks to the core, and designing next-generation access and content delivery architectures. Aravindh earned his PhD in Computer Science from King's College London, where he was a Professor Sir Rick Trainor Scholar. He is a visiting research fellow at the University of Surrey and an Internet Society (ISOC) Pulse research fellow.

Marcelo Bagnulo received the Electrical Engineering degree from the University of Uruguay and the Ph.D. degree in telecommunications from the Universidad Carlos III de Madrid (UC3M), Spain. Since 2008, he has been a tenured Associate Professor at UC3M. He has published more than 80 articles in the field of advanced communications in journals and congresses, including IEEE INFOCOM, ACM SIGCOMM, ACM Mobicom, ACM IMC, and IEEE/ACM TRANSACTIONS ON NETWORKING. He is the author of 21 RFCs in the Internet Engineering Task Force (IETF), including the Shim6 protocol for IPv6 multihoming and the NAT64/DNS64 tools suite for IPv6 transition. He has 26 H-index and 4381 total citations. His research interests include Internet architecture and protocols, interdomain routing, and security. From 2009 to 2011, he was a member of the Internet Architecture Board.

Ozgu Alay is a Full Professor at the University of Oslo. She received her B.Sc. and M.Sc. degrees in Electrical and Electronic Engineering from the Middle East Technical University, Turkey, and a Ph.D. degree in Electrical and Computer engineering from the Tandon School of Engineering at the New York University. She is the author of more than 100 peer-reviewed IEEE and ACM publications. Professor Alay actively serves on technical boards of major conferences and journals, and has participated as principal investigator in numerous EU-funded research projects in the area of networked systems and applications. Her research interests include mobile networks, low-latency networking, networked interactions and eXtended Reality.

Fabián E. Bustamante (Senior Member, IEEE and ACM) is a Full Professor of Computer Science at Northwestern University. He received his Ph.D. in Computer Science from the Georgia Institute of Technology in 2001. His research spans experimental systems, with a focus on Internet-scale networks and distributed systems. He is a recipient of the NSF CAREER Award, multiple Google Faculty Research Awards, and the E.T.S. Walton Visitor Award from the Science Foundation of Ireland.

Marco Fiore is a Research Professor at IMDEA Networks Institute, where he leads the Networks Data Science group, and co-founder and CTO at Net AI. He received MSc degrees from University of Illinois at Chicago and Politecnico of Torino, a PhD degree from Politecnico di Torino, and a Habilitation à Diriger des Recherches from Université de Lyon. Marco has held tenured positions at Institut National des Sciences Appliquées de Lyon and National Research Council of Italy, and has been a visiting researcher at Rice University, Universitat Politècnica de Catalunya, and University College London. Marco's research is at the interface of mobile networks and data science, and has received funding from the European Commission and national agencies in Spain, France and Italy, as well as a number of recognitions that include two best paper awards at IEEE INFOCOM. Marco is a former Marie Curie fellow and Royal Society visiting research fellow, and a Senior Member of IEEE and ACM.

Dr. Andra Lutu is a principal researcher and tech lead at Telefónica. With more than 10 years of post-PhD experience, her focus is on mobile network architecture design and implementation, integrating machine learning applications in the telco ecosystem. She has authored more than 70 articles in top venues, such as ACM Sigcomm, ACM MobiCom, ACM Infocom, ACM IMC, ACM CoNext, IEEE TNSM and IEEE TNET. She has contributed to numerous European projects, including SNS-JU ORIGAMI, H2020 DAEMON, H2020 MONROE, and she was a Marie Curie fellow in 2019/20.