Insights from Operating an IP Exchange Provider

Andra Lutu
Telefónica Research

Diego Perino
Telefónica Research

Marcelo Bagnulo
Universidad Carlos III de Madrid

Fabián E. Bustamante
Northwestern University

ABSTRACT

IP Exchange Providers (IPX-Ps) offer to their customers (e.g., mobile or IoT service providers) global data roaming and support for a variety of emerging services. They peer to other IPX-Ps and form the IPX network, which interconnects 800 MNOs worldwide offering their customers access to mobile services in any other country. Despite the importance of IPX-Ps, little is known about their operations and performance. In this paper, we shed light on these opaque providers by analyzing a large IPX-P with more than 100 PoPs in 40+ countries, with a particularly strong presence in America and Europe. Specifically, we characterize the traffic and performance of the main infrastructures of the IPX-P (i.e., 2-3-4G signaling and GTP tunneling), and provide implications for its operation, as well as for the IPX-P's customers. Our analysis is based on statistics we collected during two time periods (i.e., prior and during COVID-19 pandemic) and includes insights on the main service the platform supports (i.e., IoT and data roaming), traffic breakdown and geographical/temporal distribution, communication performance (e.g., tunnel setup time, RTTs). Our results constitute a step towards advancing the understanding of IPX-Ps at their core, and provide guidelines for their operations and customer satisfaction.

CCS CONCEPTS

• **Networks → Network performance analysis**; **Network measurement**; **Mobile networks**;

KEYWORDS

IPX Provider, Mobile Networks, International Mobile Roaming, Performance Analysis

ACM Reference Format:
Andra Lutu, Diego Perino, Marcelo Bagnulo, and Fabián E. Bustamante. 2021. Insights from Operating an IP Exchange Provider. In *ACM SIGCOMM 2021 Conference (SIGCOMM '21), August 23–27,2021, Virtual Event, USA.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3452296.3472930

1 INTRODUCTION

International mobile roaming is a key feature of cellular networks, enabling mobile subscribers to seamlessly use cellular services worldwide. It supports a growing number of international travellers [1], which can access data roaming at a limited or no cost [9, 10, 28], often using applications with stringent Quality of Experience (QoE) requirements (e.g., VoLTE, video streaming). Further, MNOs' infrastructure now offers the basic technological support for cellular Internet of Things (IoT) and boosts Machine-to-Machine (M2M) platforms as global connectivity providers [19]. Major Mobile Network Operators (MNOs) exploit international roaming to ensure world-wide connectivity to IoT providers, which ship their devices internationally (from wearables to cars and shipping containers) with pre-arranged cellular service (i.e., provisioned Subscriber Identity Module (SIM) card for the IoT device).

Under the IP Packet Exchange (IPX) model [4, 5], MNOs leverage IPX Providers (IPX-Ps) [26] for interconnecting with all other MNOs for roaming. Specifically, with only one connection and one agreement, IPX-Ps offer their customers (e.g., MNO) interconnection for worldwide data roaming, and support for a variety of emerging applications, including IoT verticals, VoLTE and video streaming. No IPX-P on its own is able to provide connections on a global basis (e.g., single-handily interworking with all MNOs). IPX-Ps peer to each-other to form the IPX Network, today composed of 29 active IPX-Ps peering using three major peering exchange points, and interconnecting about 800 MNOs worldwide [18]. This is an isolated network that bypasses the public Internet [3], ensuring global, secure, SLA-compliant services. In an earlier publication we provide a tutorial-style description of the IPX ecosystem [18].

In this paper, we present the first detailed analysis of operations in a real-world large IPX-P, and discuss performance implications. Despite IPX-Ps being at the core of today's international mobile ecosystem, little is know about how they operate to satisfy customer requirements. Our study is based on data we collected directly from the IPX-P's operational signaling and data roaming infrastructures for two weeks in December 2019 and July 2020. This allows us to capture the system status both prior to the COVID-19 emergency as well as the "new normal", given the significant change the pandemic brought to mobile network demand [20].

We study the main services the IPX provides, namely, the signaling services for data roaming, and zoom into the specific support of IoT customers (Section 4). **Our goal is to first establish which are the most popular solutions (i.e., corresponding to the different Radio Access Technologys (RATs)) the IPX-P offers, and which are the implications of the operational reality for the evolution of the ecosystem.** We observe the 2G/3G signaling infrastructure is one order of magnitude more loaded than the 4G one. The heavy reliance on 2G/3G incurs high costs to operators in maintaining legacy radio networks, and highlights lack of consistency in deploying latest generation technologies worldwide.

SIGCOMM '21, August 23–27, 2021, Virtual Event, USA

ACM ISBN 978-1-4503-8383-7/21/08…$15.00
https://doi.org/10.1145/3452296.3472930

We further capture the operational breadth of the IPX-P, and especially focus on how the underlying transit provider network impacts the operational presence of the IPX-P. We find that the IPX-P traffic is centered in few main mobility hubs where the IPX-P owns important trans-oceanic infrastructure. Nevertheless, operations provide coverage to more than 200 countries, highlighting the importance of the IPX Network, but can be impacted by local socio-economic mobility trends (e.g., Venezuela-Colombia migration). The analysis of the signaling error codes reveals that the IPX-P often uses these to implement specific routing policies for its customers (e.g., the non-negligible usage of steering of roaming practice via Roaming Not Allowed errors).

We also analyze the variety of device types that the IPX-P's customer base integrates. Specifically, we capture the impact of IoT devices that benefit from the IPX-P's global infrastructure. We find that most operate as permanent roamers, and their long roaming sessions contributes significantly more load to the IPX-P system than smartphone devices. Moreover, synchronous traffic patterns from IoT devices with similar behavior put a very high stress on the IPX-P platform, resulting in periodic high error rates and impact on the IPX-P performance. The design of the IoT devices (which likely ignores the GSMA standards around flow sequences for registration, retries, etc.) creates the synchronous pattern affecting the IPX Network. The large proportion of IoT devices within the IPX-P's customer base also explains why the mobility restrictions nations imposed to tackle the COVID-19 emergency did not heavily affect the IPX-P customer base (≈10% drop in number of devices active, compared to ≈20% MNOs reported [20]).

We further expose and evaluate the emerging patterns of data communication that IPX-P's end-users generate. For this, we focus our analysis on the dynamics of the data roaming service of the IPX-P (Section 5). We expose the large fraction of silent roamers from the South America region, which is the direct result of the high costs for roaming services that is still on offer from operators in the region (in contrast, for instance, with Europe and its *Roam Like At Home* regulation). Interestingly, traffic patterns of silent roamers are similar to IoT devices and generate traffic on the signaling infrastructure but very little or no data traffic.

We finally tackle one of the most important aspect of the operational analysis, namely, the performance of the IPX-P platform while enabling the data roaming service and fulfilling its main functionality (i.e., setting up and tearing down GPRS Tunneling Protocol (GTP) tunnels for data communication in roaming). For this, we analyze the data roaming dataset to reveal statistics and capture patterns of how the IPX-P platform activates GTP tunnels for data roaming communication its end-users request (Section 6). Majority of the data roaming traffic is TCP or UDP used for Web (i.e., HTTP/HTTPS) and DNS, respectively. In terms of performance, the quality of services strongly depends on the roaming configuration (i.e., home routed or local breakout), and is impacted by the geographical location of the users, or by the applications/IoT verticals and remote servers. We observe that the IPX-P takes full advantage of the flexibility if the IPX model, offering tailored solutions to its customers, to satisfy their requirements in terms of roaming configuration and quality of service.

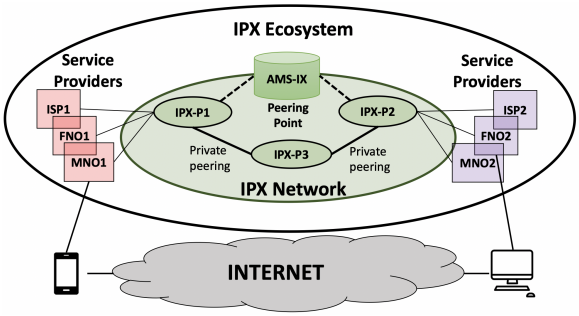

Figure 1: High level architecture of the IPX Ecosystem.

2 IPX ECOSYSTEM AND RELATED WORK

The mesh of interconnected IPX-Providers (IPX-Ps), and their Service Providers (SP) customers form the IPX Ecosystem (Figure 1). IPX-Ps are third party interconnection providers to Service Providers (SPs) such as MNOs, IoT providers or cloud providers. They provide support for global data roaming and a variety of emerging services, such as IoT, VoLTE or Rich Communication Services (RCS).

IPX-Ps peer (via private interconnects or public peering) with other IPX-Ps to extend their footprint worldwide forming the IPX Network – a private network, separate from the public Internet, that meshes together the infrastructures of the IPX-Ps. It enables the transport of global roaming data between networks, with interoperability of different implementations and standards.

SPs require a single connection and agreement with one IPX-P in order to connect to the IPX Network, and interconnect with partner SPs world-wide.[1] For instance, to enable data roaming, two MNOs must both have an agreement with an IPX-P in order to interconnect. For redundancy, a SP could establish physical connections to more than one IPX-P. Depending on the footprint of the IPX-P's infrastructure, SPs select one or more Point of Presences (PoPs) of the IPX-P to connect. In an earlier publication [18], we described the IPX ecosystem and how it enables data roaming.

2.1 Related Work

IPX model and network. The IPX model was first proposed by the GSMA in 2007 to replace the traditional, bilateral-agreement model for international roaming [5]. Despite the continuous evolution IPX-Ps and related parties [7, 33, 34] the topic has received little attention from the research community. Takaaki [23] provides an early survey of IPX and its technical requirements. Recently, in [18] we analyzed the IPX network and reported it is composed of 29 active IPX-Ps peering via three major peering exchange points, and interconnecting about 800 MNOs worldwide. [18] also showcases the radio signaling infrastructure of a commercial IPX-P and reports high-level trends. This paper presents the first in-depth analysis of a commercial IPX-P, providing implications for its operations and for its customers.

Roaming. Few studies have been conducted on roaming, possibly because its complex ecosystem and many involved parties bring about high costs and efforts for cooperation. Vallina *et al.* [31] studied national roaming between MNOs in France, and Michelinakis *et al.* [22] focused on international roaming between two

[1] Direct interconnection between SPs through leased lines or Virtual Private Network (VPN) is also possible but outside the scope of our analysis.

operators in Europe. More recently, Mandalari *et al.* [21] covered international roaming more extensively, diving into the traffic among 16 MNOs in 6 different countries. Differently, this paper focuses on the operational insights of a commercial IPX-P, with customers in 19 countries which receive inbound roamers from 215 countries and whose subscribers travel to 210 countries.

IoT/M2M. There is a large body of work on IoT/M2M traffic. In particular, numerous research shows that IoT devices generate traffic with significantly different patterns from human-driven mobile devices in the cellular networks [8, 15, 16, 27]. Based on such observations, Markus *et al.* [17] tried to fabricate M2M traffic models, and the other researchers design future systems that can efficiently handle these traffic [2, 24]. Lutu et al. [19] analyze IoT traffic from the point of view of an IoT provider and MNOs, typically customers of IPX-Ps. While this study does not focus on IoT, we complement previous work by analyzing the M2M service provided by a commercial IPX-P and provide operational insights. Our results are in line with recent work [14] that characterizes IoT signaling traffic from a network operator's point of view for the establishment of data connections at device level.

3 A LARGE IPX PROVIDER

In this section, we describe the IPX-P's underlying infrastructure, its functions and services, as well as the dataset we collect. The latter enable us to present a detailed view of the IPX-P's real-world operations.

The IPX-P we dissect is a Tier-1 Internet Service Provider operating one of the largest backbone networks world-wide. The carrier operates an IPX platform that runs on top of its vast Multiprotocol Label Switching (MPLS) transit network.[2] *The IPX-P infrastructure integrates more than 100 PoPs in 40+ countries with a particularly strong presence in America and Europe.* In terms of network connectivity, the IPX-P offers two types of interfaces, namely the IPX Access for clients (service providers) and the IPX Exchange for peering with other IPX-Ps. The main mobile peering points the IPX-P uses are those in Singapore, Ashburn and Amsterdam. By peering with other large Tier-1 carriers (via peering points or direct private interconnects), the IPX-P extends its footprint to regions where it does not own infrastructure.

The IPX-P we analyze implements in a flexible manner the model of multi-service connectivity solution. In other words, any customer can choose the optimal set of services that best fits their requirements. Specifically, the IPX-P provides a set of functions across all the different layers [29], including IPX Transport, SCCP Signaling, Diameter Signaling or GTP Signaling. Based on a tailored bundle of functions, the IPX-P then supports services such as Data Roaming, M2M and other roaming value added services (e.g., Steering of Roaming, welcome SMS, sponsored roaming, Data and Financial Clearing).

Overall, the IPX-P's customers are active in 19 countries and include MNOs, IoT/M2M service providers and cloud service providers. The majority of customers are MNOs that rely on the IPX-P for enabling data roaming for their end-users (≈75% of the customer base). IoT service providers also rely on IPX-Ps for their operations [19], and we note that these type of players constitute ≈20% of the customer base of the IPX-P we analyze.

[2]An IPX-P requires access to an underlying backbone network. The IPX-P may own its own MPLS network or alternatively, it might lease capacity on MPLS networks on which they deployed the infrastructure needed to deliver and manage inter-operable cross-network services.

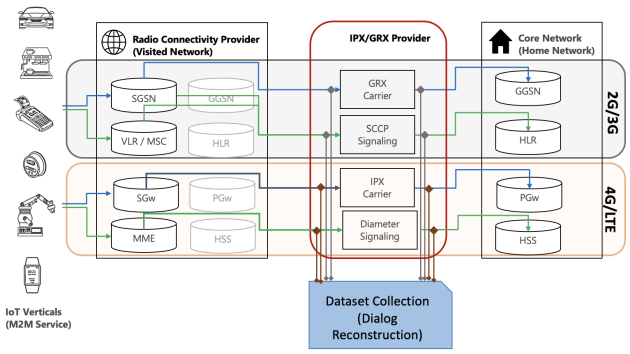

Figure 2: High level architecture of the IPX-P's monitoring. We build our dataset using a commercial software solution that processes the raw signaling traffic (SCCP, Diameter or GTP), and that rebuilds the dialogues between the different core network elements. We build datasets for 2G/3G as well as 4G/LTE.

Any customer for the data roaming service would implicitly need to use both the SCCP and Diameter signaling functions, as well as the corresponding GTP signaling function, in order to allow the different network elements from the home and visited networks to interact. Given that IoT service providers usually rely on the communication services of one (or several) MNOs, they also require access to the same type of functions that enable data roaming. However, due to the immense load they put on the IPX-P's platform, IoT providers usually have access to separate slices of the roaming platform. We detail these services and the core functions that support them next.

3.1 IPX-P Infrastructure and Monitoring

We monitor the IPX-P infrastructure that supports three core functions – SCCP Signaling, Diameter Signaling, GTP signaling (for the different RATs) – that enable two main services, namely Data Roaming and M2M service.

We capture in Figure 2 a schematic view on the manner in which we capture these corresponding datasets. We rely on a commercial software solution for capturing and analyzing in real time the raw signaling traffic, which we mirror from the signaling routers to a central collection point. In that central location, the commercial software re-builds the signaling dialogues between different core network elements in the visited and the home MNOs. We monitor for two representative periods before and during the COVID-19 emergency, from December 1st to December 14th 2019 and from July 10th to July 24th 2020, respectively. We integrate these periods to provide a longitudinal analysis of operations in the platform, and also control for any potential impact the mobility restriction to tackle the COVID-19 emergency may have had on the IPX-P's operations. However, we mention that a deep analysis on the impact of measures nations imposed to tackle the COVID-19 pandemic is outside the scope of our work. Table 1 summarizes the datasets we use to characterize the operations of an IPX-P with a large international footprint.

SCCP Signaling: This function provides access to the IPX-P's SS7 signaling network, satisfying the 2G/3G interconnection needs for international roaming of MNOs. The SCCP Signaling network of this particular IPX-P has a redundant configuration with four international Signaling Transfer Points (STPs) located in North America (Miami, Puerto Rico) and Europe (Frankfurt, Madrid).

To capture clients' activity across this signaling platform, we monitor the Mobile Application Protocol (MAP) protocol, which supports end-user mobility and allows major network elements (e.g., the Home Location Registry (HLR), Visiting Location Registry (VLR) or the Mobile Switching Center (MSC)) use to communicate. Figure 2 shows that by sampling the traffic from the SCCP Signaling platform and mirroring it to a central location, we are able to reconstruct the SCCP dialogues between different network elements, and build the dataset we use in this paper. We use a commercial solution for the raw data processing towards rebuilding the SCCP records that correspond to procedures devices in international roaming trigger. We collect traffic corresponding to the following procedures of each device belonging to one of the IPX-P's clients (outbound roaming) or to foreign devices that connect to the network of one of the IPX-P's clients (inbound roaming): i) location management (update location, cancel location, purge mobile device); ii) authentication and security (send authentication information); iii) fault recovery.

Diameter Signaling: This function provides the Diameter signaling capabilities necessary to enable 4G roaming for customers. The infrastructure of this particular IPX-P includes four Diameter Routing Agents (DRAs) meant to forward Diameter messages and simplify interworking between different network elements. It is application-unaware and does not inspect the messages it receives. The service also integrates Diameter Proxy Agents (DPAs), which include the functionality of the DRAs and can additionally inspect and route Diameter messages based on different parameters. Finally, by leveraging the Hosted Diameter Edge Agent (DEA) service, the IPX-P offers a infrastructure-as-a-service functionality to help operators expedite the launch of LTE roaming services. Thus, operators can use the dedicated customer virtual DEA from the IPX-P instead of deploying their own infrastructure. The LTE Diameter service integrates other value added services, including Welcome SMS, Steering of Roaming or Sponsored Roaming.

To monitor the activity of the IPX-P's customers, we monitor traffic across the geo-redundant signaling network with four DRAs located two in Europe (Frankfurt, Madrid) and two in North America (Miami, Boca Raton). The approach is similar to the case of SCCP Signaling we described above (Fig. 2). We collect traffic corresponding to events including Diameter Transactions.

Data Roaming: This service enables the IPX-P to connect MNOs with foreign roaming partners, to ensure the data transport required for data roaming in 2G/3G (Gn, Gp interfaces) and LTE (S8 interface). The data roaming service relies on the GTP function to build and manage tunnels between roaming partners, to transport data to and from end-users. Note that the service requires the use of the SCCP and the Diameter signaling functions.

For this paper, we collect statistics regarding the tunnels between the Serving GPRS Support Node (SGSN) and Gateway GPRS Support Node (GGSN) nodes for 2G/3G, and between Serving Gateway (SGW) and Packet Data Network Gateway (PGW) for LTE. The IPX-P we study deploys a commercial software solution for monitoring, which centralizes large amounts of data from the different network elements (namely, the SCCP, Diameter and GTP signaling points) that are part of the infrastructure they operate. We capture the Create/Delete Packet Data Protocol (PDP) context procedures that the devices trigger before/after a data communication, as well as metrics about the data sessions. Specifically, the monitoring solution generates one record for the Create Session Request/Response exchange and retains basic information, such as the tunnel ID. Additionally, the monitoring solution generates a record when a data session is completed, which captures statistics for the whole session, such as the total amount of bytes transferred or the RTT. Because of the high amount of traffic and processing that collecting and generating these statistics implies, we only collect this dataset for the inbound and outbound roamers for the IPX-P's customers connecting to PoPs in only a few selected countries (i.e., Spain, US, Brazil, Argentina, Colombia, Peru, Costa Rica, Uruguay, Ecuador).

M2M Service: An M2M platform operating on top of the IPX-P allows to avoid the cost of establishing technical and commercial relationships with every local operator, and can offer more stable connectivity/coverage services to IoT providers. By leveraging data roaming and the basic functions this requires (e.g., SCCP signaling, Diameter signaling, IPX/GRX carrier), an M2M platform can directed all traffic from its IoT devices to a single home country, no matter where the device is located in the world. M2M platforms are being leveraged by a growing number of industries, from health to automotive and logistics.

We monitor the activity of one specific M2M platform that relies on a Spanish MNO and on the IPX-P we analyze to support its business. We separate from the above-mentioned signaling and data roaming datasets only the traffic corresponding to the IoT devices this M2M platform operates. For this, we use the unique identifiers (i.e., encrypted Mobile Station International Subscriber Directory Number (MSISDN)) assigned to each device of the M2M platform. This allows us to capture the performance of the IPX-P solution from the point of "things" using the same system.

3.2 Ethical considerations

Data collection and retention at network middle-boxes are in accordance with the terms and conditions of the IPX-P and the local regulations, and only with the specific purpose of providing and managing the IPX service. The terms also include data processing for monitoring and reporting as allowed usages of collected data. Data processing only extracts aggregated information and we do we not have access to any personally identifiable information. We nevertheless consulted with the Institutional Review Board (IRB) office at our institution who confirmed that no IRB review was necessary as the study relies on the analysis of de-identified data.

4 SS7/DIAMETER SIGNALING

In this section, we provide insights into the complexity of the operations of the IPX-P by analyzing one of the base functions it provides, namely, signaling between mobile core network elements – SCCP

Table 1: IPX Datasets.

Dataset	Infrastructure	Procedures captured
SCCP Signaling	4 STPs (Miami, Puerto Rico, Frankfurt, Madrid)	MAP traffic, location management , authentication and security
Diameter Signaling	4 DRAs (Miami, Boca Raton, Frankfurt, Madrid)	Session Initiation Protocol (SIP) Registration, Voice over IP (VoIP) Call, Diameter Transaction, Domain Name Service (DNS) Query or RCS Session
Data Roaming	GTP-C control data and GTP-U data sessions.	Create/Delete PDP Context/Session; Flow-level metrics for data connections.
M2M Platform	IoT devices for specific M2M customer	SCCP Signaling, Diameter Signaling and Data Roaming.

Signaling and Diameter Signaling. These functions are mandatory for the correct operation of data roaming and multiple other IPX-P services. The SCCP Signaling and Diameter Signaling datasets we collect from the IPX-P comprise signaling information for 2G/3G and 4G/LTE radio technologies from all the devices that use the IPX-P's infrastructure (Table 1).

4.1 Signaling Traffic Trends

Figure 3 shows signaling activity of roaming mobile subscribers during the observation period in July 2020. We look at both Signaling System No. 7 (SS7) and Diameter signaling procedures. MAP is the most important application protocol in the SS7 stack, and handles the roamers' mobility between countries in 2G-3G. The same function is performed by the Diameter [13] signaling protocol in LTE.

Overall, we capture more than 120M devices active in the MAP dataset, and more than 14M devices active in the Diameter dataset.[3] This shows a slight decrease compared with December 2019 (likely due to mobility restrictions imposed to tackle the COVID-19 emergency [20]), when the total number of IMSIs we captured was more than 130M active in 2G/3G and more than 15M active in 4G/LTE. These results highlight that *the IPX-P 2G/3G infrastructure handles an order of magnitude more devices than the 4G infrastructure.*

Figure 3a shows the average number of records per IMSI calculated over all the IMSIs we observe in each one-hour interval (continuous line) during the July 2020 observation period, as well as the standard deviation of the number of records per IMSI calculated over all the IMSIs active in the same one hour interval (shaded area). We observe both the MAP procedures for 2G/3G (red color) and the similar Diameter procedures for 4G/LTE (green color). Each record in both of these datasets represents a signaling dialogue that two network elements have, corresponding to different standard procedures. For instance, from the MAP interface we capture mobility management routines, including location management and

[3]Note that there might be an overlap between these two sets, as some devices switch from 2G/3G to LTE connectivity. However, we aim to show here the load on the two different signaling infrastructures.

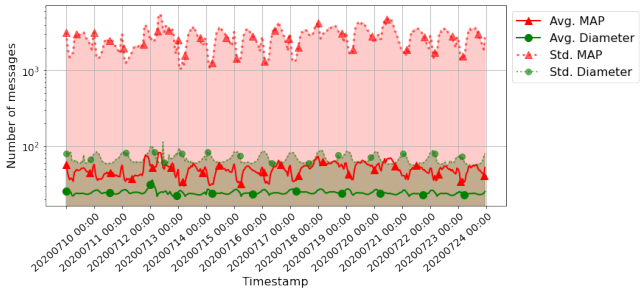

(a) Average and standard deviation of the number of MAP and Diameter messages per IMSI per hour.

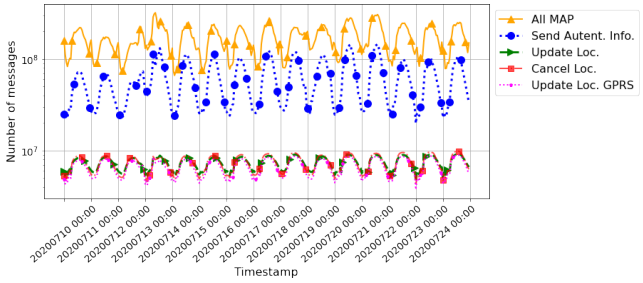

(b) MAP signaling traffic breakdown per type of procedure.

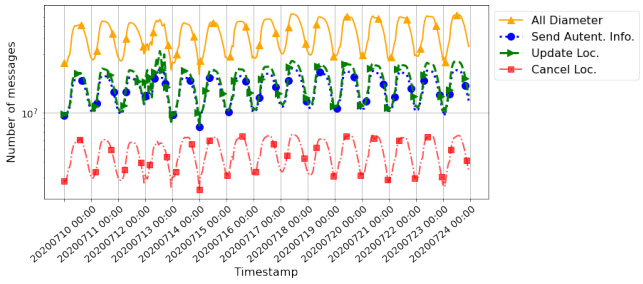

(c) Diameter signaling traffic breakdown per type of procedure.

Figure 3: Signaling traffic time series for a two-week observation period in July 2020.

authentication. While Diameter and MAP are different protocols, the underlying functional requirements (e.g., authenticating the user to set up a data communication) have many similarities in terms of the messages used for Diameter and the SS7 MAP protocol implementation. We note that the load in terms of average signaling records per IMSI is in the same order of magnitude (the continuous lines on the plot), regardless of the infrastructure the devices use; yet, *there are significantly more messages generated on average by an IMSI using MAP than an IMSI using Diameter, as Diameter is a more efficient protocol than MAP [13, 30].*

We further break down the signaling traffic on record type (or procedure) both for MAP (Figure 3b and Diameter (Figure 3c). Figure 3b shows the time series of signaling traffic broken down by type of procedure, including Update Location (UL), Cancel Location (CL) and Send Authentication Information (SAI) messages. The latter, SAI, represents the highest fraction of MAP signaling traffic; this is also the case for the Diameter signaling traffic. Indeed, according to the GSM standard definition, the visited network triggers the authentication of subscriber procedure upon IMSI attach, location

update or before starting data communication, thus explaining the larger volume of SAI messages.

Takeaway: We find that the number of devices using the IPX-P's 2G/3G infrastructure (MAP traffic) is an order of magnitude higher than those using 4G infrastructure (Diameter traffic). The volume of signaling traffic in the SCCP infrastructure is, correspondingly, more significant than in the Diameter infrastructure. This heavy reliance on 2G/3G is problematic, because of the high costs the maintaining legacy radio networks incurs to operators. This brings to light the lack of global consistency between operators in deploying the latest generation access technologies. Further, the use of less efficient protocols imposes a higher operational cost for both the IPX-P platform and its customers.

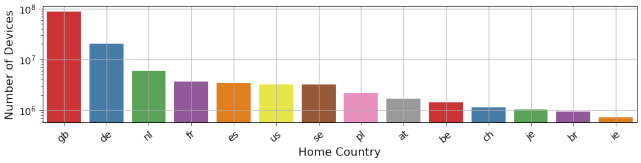

(a) Distribution of devices per home country.

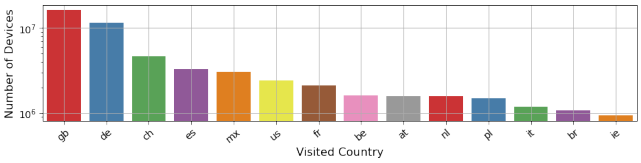

(b) Distribution of devices per visited country.

Figure 4: Distribution of device per home country and visited country; we include here all devices active in any of the two signaling datasets.

4.2 Operational Breadth

The goal the IPX-P is to offer global coverage to its entire customer base. This means allowing all customer devices to connect anywhere in the world, and, conversely, allowing anyone in the world to connect to their customers' networks. Overall, the IPX-P's infrastructure serves devices from MNOs from over 220 (home) countries, operating in more than 210 (visited) countries. In Figure 4 we show the distribution of mobile devices, and focus on top-14 home operators and top-14 visited operators in July 2020. We notice that *the distribution is fairly skewed to few operators, and the best represented countries correspond to the locations of the main IPX-P's customers, namely Spain, UK, Germany.*

Through the lens of the signaling dataset, we can further capture the (international) mobility of devices. Figure 5 shows the distribution of mobile devices, registered during each of the two observation intervals (December 2019 in Figure 5a and July 2020 in Figure 5b), that travel from their home country (column) to a visited country (row). In the following we comment on the December 2019 dataset. Overall, we find that the majority of subscribers using the IPX-P infrastructure - serving large European MNOs – comes from UK (≈8 million devices in December 2019), Germany (≈2 million devices) or Spain (≈2 million devices). Most of these devices tend to visit the UK (≈6.5 million devices), Germany (≈2.5 million devices) or the US (≈500,000 devices).

When clustered by geographic regions (i.e, Europe and the Americas), we see that the most popular destinations roamers visit include the UK in Europe and the US in America. Indeed, Figure 5a shows that the UK operators connected to the IPX-P we analyze receive 34% of all the devices from Germany (DE) visible in the system, 85% of the devices from the Netherlands (NL), and 45% of all devices from Spain (ES), among others. Interestingly, we verified with the British operator connected to our IPX-P and found that the inbound roaming devices from the Netherlands (≈ 7.8 million devices) are IoT devices deployed by energy providers (smart meters)[4]. US, Brazil and Mexico emerge as the most popular destinations in the Americas. Specifically, the US operators connected to our IPX-P accommodate 79% of all the outbound roaming devices from Mexico (MX) using the IPX-P infrastructure, 44% of all outbound roaming devices from El Salvador (SV), 17% of all outbound roaming devices from Colombia (CO) and 22% of all outbound roaming devices from Brazil (BR).

Finally, it is interesting to note how data from an IPX-P can capture *socio-economic patterns in international mobility*. Indeed, we can observe the migration between Venezuela and Colombia, with 71% of the subscribes from Venezuela (VE) traveling to Colombia (CO) during the period we capture. Inversely, we find that 56% of all Colombian outbound roamers travel to Venezuela (VE). The Venezuela-Colombia border is one of the most active in the world [11], as Colombia is the primary destination of most Venezuelan migrants, which capture Venezuelans with different status ranging from economic migrants to refugees.

We observe in both the observation periods a fraction of devices that operate within their home countries. For instance, in July 2020 (Figure 5b), we note that 39% of all the UK devices operate within their home country, or 47% of Mexico device operate within Mexico. These usually belong to Mobile Virtual Network Operatorss (MVNOs) enabled by the IPX-P we analyze, to operate on top of MNOs that are already customers of the MNO. The increased ratio in July 2020 is a side-effect of reduced international mobility in July 2020, compared to December 2019.

Takeaway: The IPX-P underlying infrastructure impacts its operational breadth. Specifically, given that the IPX-P leverages access to important trans-oceanic infrastructure connecting the Americas and Europe (e.g., Brusa subsea cable connecting Brazil and USA, Marea subsea cable connecting the US and Spain, or the SAm-1 subsea cable with various landing points from US to Argentina), we note that these are the main markets where it operates. Specifically, US, UK, MX and BR emerge as the main mobility hubs for devices that depend on this particular IPX-P to operate and infrastructure needs to be provisioned accordingly. At the same time, IPX-P operations need to provide coverage beyond this core infrastructure to more than 200 countries, and can be impacted by specific socio-economic trends.

4.3 Steering of Roaming

Every dialogue we capture in our signaling dataset corresponds to a roaming procedure and includes, apart from the requested

[4]This is consistent with deployment of the Smart Metering Implementation Programme of the U.K. government. In 2019, there were 15.4 million smart meters deployed in the UK [32].

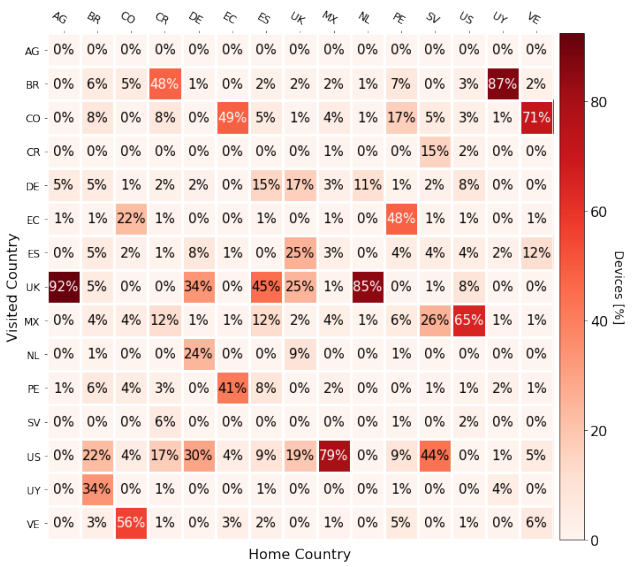

(a) December 2019.

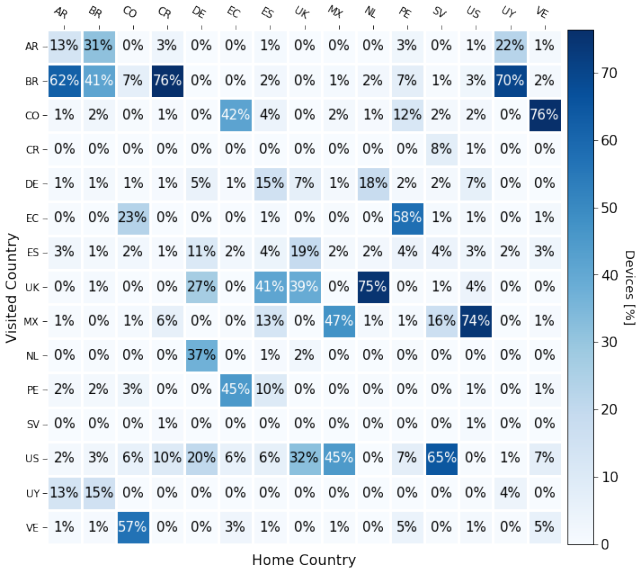

(b) July 2020.

Figure 5: Mobility dynamics based on SCCP signaling (devices using 2G/3G) and Diameter signaling (devices using 4G/LTE).

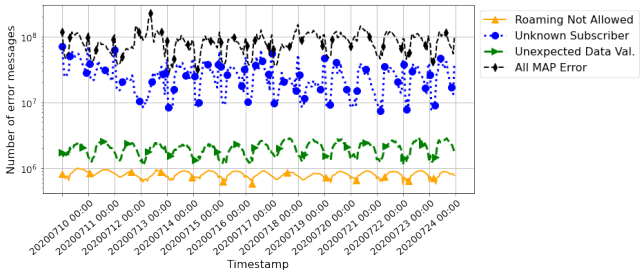

Figure 6: Breakdown of the MAP Error Codes (July 2020).

operation code, the result of the operation. For example, in the case of roamer authentication in the visited network, the requested operation code is "Send Authentication Information", to which the home network replies with the requested information. This is the most frequent procedure we capture in our dataset (Figure 3). In the event of a failure, the response from the home network may contain an error code showing why the procedure failed. Such errors for the SAI operation include Unknown Subscriber (There is no allocated IMSI or no directory number for the mobile subscriber in the home network), and for the UL operation include Unexpected data value (The data type is formally correct, but its value or presence is unexpected in the current context.).

Figure 6 shows the time-series of errors in the MAP dataset, regardless of the type of the operation that triggered them, broken down per type of error, for the July 2020 dataset. We note that the most frequent error is Unknown Subscriber, pointing to a numbering issue during the SAI procedure.

Another frequent error code we observe is Roaming not Allowed (i.e., the home operator is baring the roaming of the device), which corresponds to an Update Location procedure. Often, operators use this error code to implement different routing policies for the mobile user, such as *Steering of Roaming (SoR)* [6]. In a general manner, this may bring an increase of the signaling load between 10% and 20% [6].

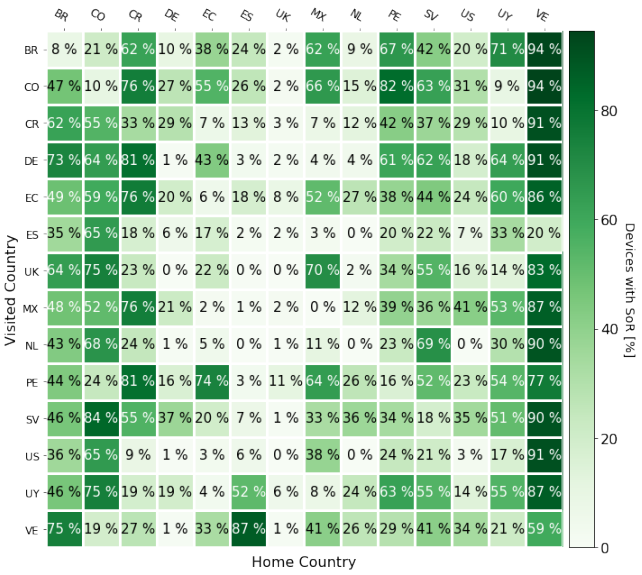

Figure 7: Steering of Roaming Service Analysis: Percentage of devices that travel from home country to visited country for which the IPX-P triggered *Roaming Not Allowed* at least once over a period of two weeks in December 2019.

By using SoR, an MNO can specify the preferred roaming partner in a given visited country and allow the IPX-P to enforce those preferences. With this in place, if a roamer device traveling outside its home network (HMNO)[5] attempts to attach to a less preferred roaming partner, the IPX-P will force the *Roaming Not Allowed* response code (RNA, error code = 8) to an Update Location (UL) message intercepted from the visited network (VMNO). The IPX-P will then try steering the roamer to one of the HMNO's preferred

[5] An MNO can use the IPX-P within the same home country for national roaming or for enabling virtual operators, as well as in foreign countries for international roaming.

roaming partners after forcing four UL attempts from the roamer to fail, unless no preferred roaming partner is available in the area (in which case, the SoR platform triggers an exit control to avoid the risk of the roamer not receiving service at all).

Figure 7 shows the percentage of end-user devices roaming from the home country (column-wise) to a visited country (row-wise) for which we registered at least one RNA error code for the UL procedure. *We observe a non-negligible number of Roaming Not Allowed operational code as a result of the UL request from the VMNO, typically due to the use of the SoR service the IPX-P provides its customers.*

One notable exception is Venezuela. We note the prevalence of this error code for mobile subscribers traveling from Venezuela abroad, regardless the visited country. Because of the volatility of Venezuelan currency, mobile operators in Venezuela suspended international roaming as they said they lacked enough foreign currency to pay roaming partners in foreign countries. The reason this is allowed for Spain (where we only note that 20% of subscribers from Venezuela receive a RNA message) is because of internal agreements between operators that belong to the same international corporation.

On the other end of the spectrum, we see that the fraction of UK users (marked GB in Fig. 7) affected by this error code is very small, regardless of the country they visit. This is because the IPX-P's customer in the UK does not use the SoR service from the IPX-P, but instead handles the steering of its subscribers separately. Thus, the RNA errors we capture are due to the HMNO from UK not allowing its subscribers to roam (e.g., because of billing issues).

Takeaway: Operators use forced errors to implement different policies for their subscribers when these are roaming abroad. One example is the Steering of Roaming, which the IPX-P offers as a service for its customers, at the cost of increasing the signaling load on the roaming platform.

4.4 Impact of IoT Devices

In the following paragraphs we focus on the traffic corresponding to the IoT devices that the IPX-P M2M platform operates for the December 2019 dataset. Although not reported, analysis of the July 2020 dataset leads to similar takeaways.

Figure 8 shows the time series of average number of signaling messages per device for 2G/3G (Fig. 8a) and 4G/LTE (Fig. 8b), as well as the 95th percentile calculated over one hour intervals. To put this figures in context, we include statistics from a similar number of smartphones using the same radio technology. We selected the set of smartphones leveraging the device brand information, which we retrieve by checking the International Mobile Equipment Identity (IMEI) and the corresponding Type Allocation Code (TAC) code, and included only iPhone and Samsung Galaxy devices (the two most popular smartphones) in the pool. Figure 8 shows that *IoT devices generally trigger a higher load on the signaling infrastructure, regardless of the infrastructure they use (Diameter or SS7). This holds when checking either the average number of messages per device across time as well as the 95% percentile per hour across all devices.*

We also compare the duration of roaming sessions (i.e., the total number of days a device sent at least one signaling message while in roaming) for both IoT devices and smart phones. Figure 9a

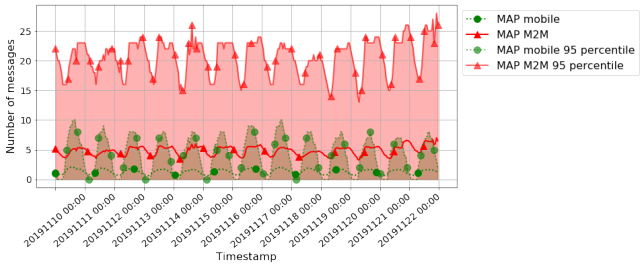

(a) 2G/3G M2M/IoT devices and smartphone devices.

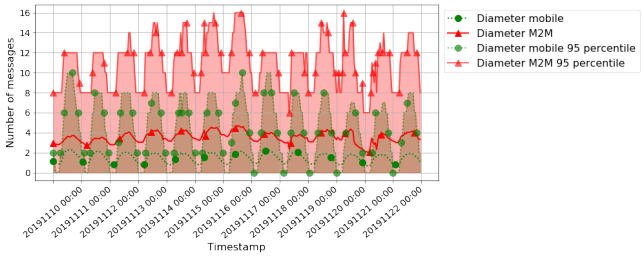

(b) 4G/LTE M2M/IoT devices and smartphone devices.

Figure 8: Signaling traffic statistics for IoT/M2M devices and smartphones. We generate the average per device (continuous line) and 95th percentile (shaded area) for each one hour interval over a period of two weeks in December 2019.

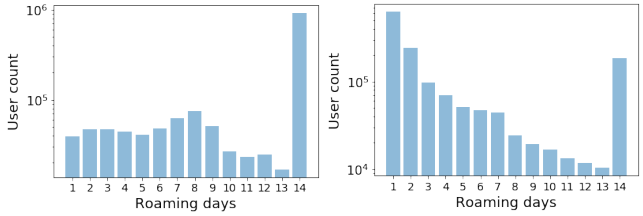

(a) IoT devices. **(b) Smartphones.**

Figure 9: Roaming session duration for (a) IoT devices and (b) smartphones (December 2019).

shows the number of days an IoT device was active during the two-week period we analyze. We note that the majority of IoT devices have long roaming sessions, which in our case cover the entire observation period. This is very different to what we observe for smartphones (Figure 9b), whose roaming session lengths are shorter. This is expected, since *IoT devices are meant to provide services in the country where the IoT provider deploys its services during long periods of time, thus becoming "permanent roamers" in the visited country. At the same time, this also translates into significant signaling load on the VMNO infrastructure from inbound roaming IoT devices.*

Takeaway: The M2M service of the IPX-P is very popular and has different operational requirements than the other services. IoT devices operate as permanent roamers with long roaming session, and generate more signaling traffic than smartphones. This contributes significantly more traffic to the IPX-P signaling system than smartphone devices.

5 GTP-C SIGNALING

In this section, we capture dynamics of data roaming, focusing on the GTP tunnels the IPX-P manages between roaming partners to enable data communications for the users.

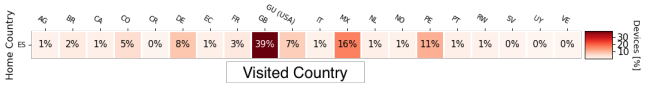

(a) Breakdown per visited country.

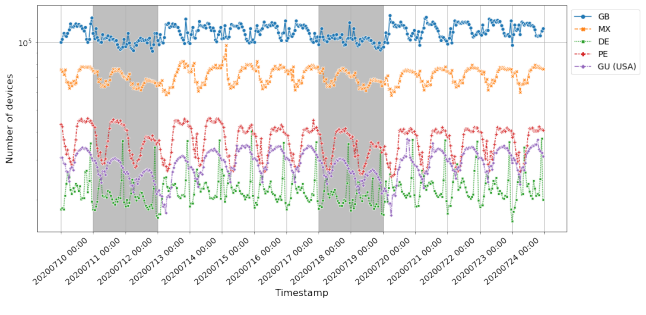

(b) Time-series of number of active devices per hour, for the top five visited countries.

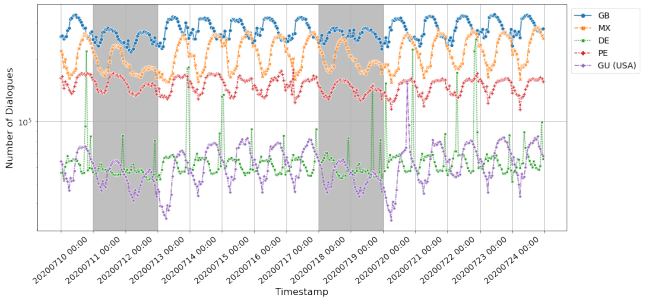

(c) Time-series of number of GTP-C dialogues per hour, for the top five visited countries.

Figure 10: Breakdown of active devices in the data roaming dataset from Spain per visited country (July 2020). Grey areas indicates weekend days.

5.1 Data Roaming Dataset

The GTP-C protocol is used for setting up and tearing down GTP tunnels for user data across the IPX-P platform. The data roaming dataset (see Table 1) we collect includes information for a subset of devices that we previously captured in the signaling dataset (Section 4). For the observation period of July 2020, we capture the GTP-C data records from over 3.3M devices operating world-side, in over 170 (visited) countries. Majority of these devices uses SIM cards from operators in Spain (≈2,3million devices) or in Brazil (≈600k devices). Given that the devices from Spain (corresponding to the same IPX-P customer, an IoT service provider) represent approximately 70% of all the devices in this dataset, we focus our analysis on these, to characterize the operations of the IPX-P offering the data roaming service.

Figure 10a shows the breakdown of this set of devices per visited country. We note that all these devices are IoT devices, serving different verticals. We observe that the main areas of operation for this set of devices include the UK (40%), Mexico (16%), Peru (11%) and Germany (8%). We notice that the main area of activity is Europe and the Americas (where the IPX-P has important trans-oceanic infrastructure), which is consistent with the previous observations from Section 4.

Figure 10b-10c shows the number of active devices and total number of GTP-C dialogues they trigger per hour in the top five visited countries, respectively. We notice a *daily pattern for both metrics. Also, during the weekend the number of active devices and overall data roaming activity decreases* (the grey area in the time-series plots).

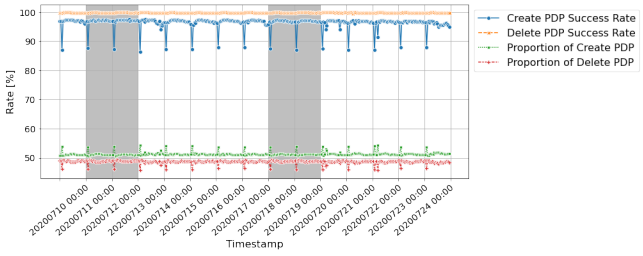

(a) Success rate of create/delete PDP context requests.

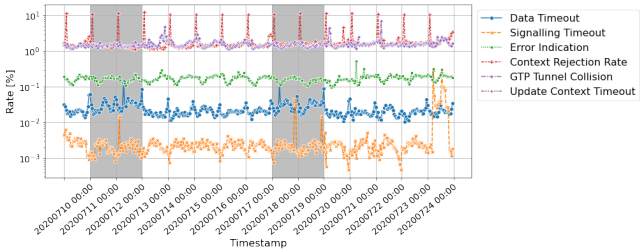

(b) Error rate for GTP-C dialogues.

Figure 11: Time-series of the result of PDP create/delete requests in the data roaming dataset (July 2020).

The two main types of GTP-C dialogues correspond to the procedures to setup and tear down tunnels, namely create/delete PDP context requests. Figure 11 captures the success rate and the error rate of these GTP-C dialogues. The distribution of dialogues on the type of request (create/delete PDP context) is symmetrical, with slightly higher ratio of create PDP context requests (Figure 11a). Interestingly, we notice that *many of the devices from the Spanish operator request data roaming connections at the same time, putting a high load on the platform. The synchronicity of the devices comes from the fact that they are IoT devices with pre-determined behavior by the IoT vertical providers* (e.g., they might be smart energy meters the energy companies deploy). This brings an important challenge to the IPX-P, since the platform is not dimensioned for peak demand. This results in a decreased success rate (in Figure 11a we notice that the success rate drops below 90% every day at midnight), and an overall larger number of create PDP context requests. Overall, the delete PDP context requests have close to maximum success rate.

We further investigate the different errors the unsuccessful dialogues include (Figure 11b). The Signaling timeout error has the lowest rate (affecting 1 in 1000 GTP-C requests), showing that it is rare that a Create PDP Context request remains unanswered and

times-out. Once a data communication is successfully established, it may be terminated because of lack of data transfer, generating a Data Timeout error. This error does not imply that there is something technically wrong with the data communication. We see this occurs for approximately 1 in 100 data communications. Interestingly, we note a clear increase of this type of error during the weekends (corresponding to the grey areas in the time-series). The Delete PDP Context request may result in an "Error Indication" result, when the operation is unsuccessful. This affects 1 in 10 such requests, and shows a clear daily pattern. Finally, the Context Rejection presents the same pattern with the Create PDP Context time-series, confirming that the IPX-P cannot respond the synchronized behavior of groups of IoT devices.

Takeaway: Signaling traffic for data communications over the IPX-P shows daily and weekly patterns. Synchronized PDP context requests from devices with similar behavior (e.g., IoT devices such as smart energy meters) put a very large stress on the IPX-P platform, resulting in high Context Rejection (≈10% of requests are rejected).

5.2 GTP-C Performance

Leveraging the data roaming dataset (Table 1) from December 2019 we now characterize the performance of the GTP tunnel management of the IPX-P. Specifically, in Figure 12a we investigate the tunnel setup delay (the time between a PDP Create request and its reply) and the total GTP tunnel duration (the time between a PDP Create and the corresponding PDP Delete event) as they have a strong correlation with the load on the IPX-P, HMNO and VMNO systems. We use the data from December 2019 to avoid the impact of the travel restrictions imposed to tackle the COVID-19 pandemic.

The tunnel setup delay is an indicator for the amount of processing involved at different network elements for Create PDP messages (as well as the general processing load). In Figure 12a (green line), we notice the average setup delay (≈150ms) depends on the total number of devices requesting a data connection at a moment in time. This value is consistent with the setup delay values we capture in July 2020. We also note that, in 80% of cases, we measure a tunnel setup delay below 1 second.

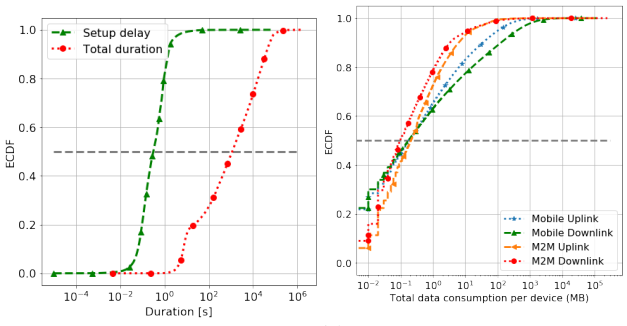

(a) GTP Tunnel analysis. **(b) Data volume per roaming session.**

Figure 12: Analysis of how roamers between countries in Latin America use the IPX-P system: (a) GTP Tunnel setup delay and total GTP tunnel duration; (b) average amount of traffic roamers generate during data sessions, comparing with IoT devices provisioned by an IoT provider from Spain (December 2019).

A decrease in the average tunnel duration will increase the number of total tunnels and thus also the volume of signaling messages and the necessary processing for these messages. Conversely, longer tunnel duration cause an increased overall memory footprint in the involved nodes to store the PDP Contexts. When verifying the total tunnel duration, we note that on median, the duration of the GTP tunnel is approximately 30 minutes (Figure 12a red line). Private conversation with operational teams confirmed us that these *values for these metrics are an indication of healthy systems*, i.e., processing and storage load at IPX-P, MVNO, HMNO elements are under normal operational conditions.

One likely important factor that influences both these metric is the device type, e.g. phone or IoT Operating System (OS). For instance, the OS implementation decides when the device should establish a mobile data connection, how long the connection is held, or which mobile technology takes preference. Since this ecosystem is extremely varied, we are here interested in the aggregated impact on the IPX-P and MNO systems. The granular analysis of the device type impact on these metrics is outside the scope of this analysis.

Takeaway: The load on the platform, in terms of number of tunnels and PDP Create/Delete requests, impacts the speed to bring up tunnels for new data communications that customers request. The IPX-P maintains a healthy system with similar values for both analyzed datasets.

5.3 Silent Roamers

Despite the dynamic global movement of mobile subscribers, not all might be active in terms of data communications. Data communications while roaming have often generated bill shock for mobile subscribers or kept roamers silent (i.e., they do not trigger data communications over cellular networks). Thus, when traveling to a foreign country, mobile subscribers often turn off the data communication capabilities of their devices to avoid high bills. Even if this may no longer be the case for roamers inside Europe [9], we find that majority of roamers within Latin America are still silent.

By contrasting the mobility of users from signaling dataset (regardless the radio access technology) with the activity we register in the data roaming dataset (active GTP tunnels, see Section 4), we are able to quantify the amount of silent roamers. For the first two weeks of December 2019, we capture the signaling activity of ≈ 2 million subscribers roaming between the Latin American countries where the IPX-P has significant volume of subscribers (Brazil, Argentina, Colombia, Costa Rica, Ecuador, Peru and Uruguay). Out of these, we find that ≈400,000 mobile devices only are using data services while traveling abroad within Latin America. For these, we observe in Figure 12b that the amount of total traffic volume per session (uplink or downlink) is no more than 100KB, in average, per device.

Even more, when focusing on *inbound* roamers in Latin America (regardless the home country), we also capture ≈2,5 million IoT devices provisioned by one of the IPX-P's M2M customers. The latter provisioned the IoT devices operating in Latin America with connectivity from a Spanish MNO. We compare the amount of traffic each roamer within Latin America generates with the amount of traffic from IoT devices (Figure 12b). *We find that, although "things"*

generate very little traffic, mobile subscribers within Latin America have a very similar behavior (though tend to transfer slightly larger data volumes than IoT devices). We conjecture this is the result of the lack of regulation on roaming within the region, as well as the socio-economic landscape, which keeps the cost of roaming data communications prohibitive.

Takeaway: Silent roamers are still a phenomenon we observe, especially in Latin America, where roaming charges are high. Their traffic patterns are similar to IoT devices, generating signaling traffic, but very little or no data traffic.

6 DATA ROAMING TRAFFIC

In this section, we explore the data roaming dataset (see Table 1), and give further insights into the type of traffic flowing through the IPX-P's platform and its performance.

6.1 Roaming Traffic Breakdown

Each data record captures different applications/protocols that correspond to the subscriber activity over the period we monitor. The data record integrates performance parameters per roaming data communication, including Round-Trip Time (RTT), packet retransmissions or volume of bytes transferred (uplink and downlink). We find that *the majority of the traffic we capture is TCP (40%) or UDP(57%)*, with a small fraction of ICMP (2%) and other protocols. The destination port breakdown for the TCP records we capture reveals that *the largest amount of traffic is web traffic (HTTP, HTTPS), accounting for 60% of TCP traffic.* The breakdown of UDP traffic shows *that majority of traffic we capture is DNS over port 53 traffic (more than 70%).* This is (largely) due to the procedure the MNOs roaming partners implement in order to allow for international context creation (tunnels) over the IPX Network (i.e., control traffic). The Visited Mobile Network Operator (VMNO) uses the IPX to resolve the Access Point Name (APN) associated to the mobile subscriber to an actual IP address corresponding to the home network GGSN (or PGw for EPC), which the Home Mobile Network Operator (HMNO) performs.

6.2 Performance Implications

In this section, we investigate the quality of the roaming service over the IPX-P's network. We focus our analysis on devices operating with IMSIs from a Spanish operator that supports multiple IoT verticals (e.g., energy sensors, fleet tracking, wearables, etc.) over the world. We zoom into the top countries in terms of number of devices (namely, UK, Mexico, Peru, US and Germany), and monitor the session duration, RTT uplink and downlink, and the connection setup delay for all the TCP data communications the IPX-P supports during the period of analysis (July 2020). The session duration (Figure 13a) varies largely, depending on the country where these IoT devices roam and, likely, the usage dictated by the IoT provider deploying these devices. We note that the devices in Germany have significantly longer average session duration (≈45s) than devices in the UK (≈150s).

For the TCP traffic, we further check the RTT distribution the mobile subscribers experiment, broken down per visited country. The uplink RTT (Figure 13b) captures the RTT between the sampling point (i.e., Miami) and the application server, capturing the impact of the PGw (or the GGSN) and the latency over the Internet path towards the application server. The downlink RTT (Figure 13c) captures the RTT between the sampling point within the IPX-P's infrastructure and the mobile subscriber, thus capturing the impact of the visited network (including the radio access network) and the SGw (or the SGSN, respectively). For both metrics, we notice that the lowest values are for devices operating in the US. This is due to the use of a different roaming configuration, namely local breakout roaming configuration, in this visited MNO. We note that, in the case of home routed roaming configuration, the uplink RTT might vary with the distance between the home country (in this case, Spain) and the visited country. This is reflected in the distance between the PGw (within the Spanish operator's network) and the application server (likely within the visited country). The downlink RTT shows similar pattern and rank between the groups of devices per visited country.

The connection setup delay (Figure 13d) represents the time in milliseconds between TCP SYN (first packet sent by the mobile subscriber) and TCP ACK (last packet in the three handshake procedure). We observe that this metric does not follow the same trends of the RTTs. This highlights the applications/IoT verticals and of remote servers play a dominant role in the connection setup delay.

Takeaway: Majority of the data roaming traffic is TCP or UDP used for Web (i.e., HTTP/HTTPS) and DNS over port 53, respectively. In terms of performance, the RTTs strongly depend on the roaming configuration (i.e., home routed or local breakout), and is then impacted by the geographical location of the users, while the connection establishment delay and session duration is dominated by the applications/IoT verticals and remote servers. The IPX-P leverages the flexibility of the IPX model to offer tailored connectivity to different customers, depending on their requirements.

7 CONCLUSIONS

IPX-Ps are at the core of the IPX ecosystem, allowing their customers to achieve global connectivity for their end-users. A view into this opaque ecosystem using public available data is not possible, as IPX-Ps intentionally separate their operations from the public Internet. In this paper, we provide the first deep dive analysis into the operations of a real-world IPX-P, with a large platform serving customers in 19 countries. As services across different IPX-Ps are generally consistent, and rely on the same basic functions that we explore in this paper, we argue that our insights provide a valuable peak into this hidden operational ecosystem.

We show how the IPX-P benefits from the flexibility of the IPX model to create tailored solutions for its customers, which include IoT providers or MNOs. We build a complex dataset to capture operations over the basic functions of the IPX-P, including SCCP signaling, Diameter Signaling, and GTP signaling. These allow us to dissect the data roaming service, both for MNOs and IoT providers. We characterize the traffic and performance of the main infrastructures of the IPX-P, and provide implications for its operations, as well as for the IPX-P's customers.

Our analysis leaves several open questions for the community to consider. Though the IPX ecosystem was meant to come with intrinsic security (via the deliberate separation with from the public Internet), there are many well-known weaknesses in the current SS7

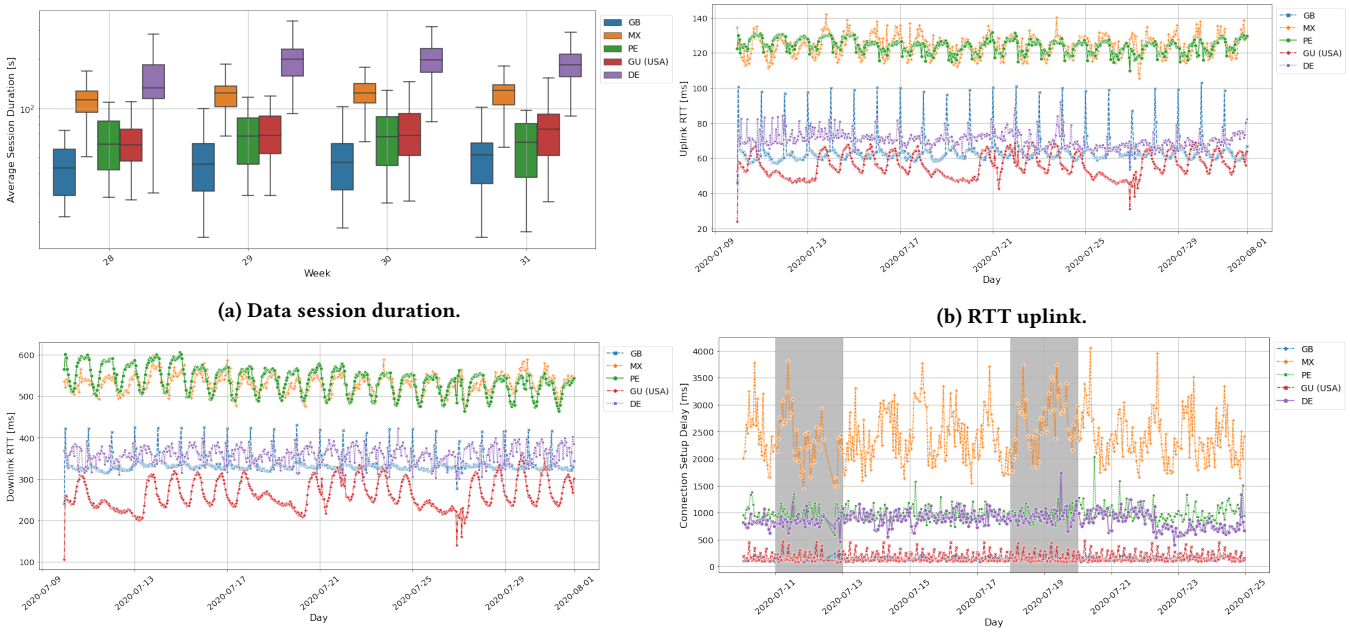

(a) Data session duration.

(b) RTT uplink.

(c) RTT downlink.

(d) Connection Setup Delay.

Figure 13: Service quality of TCP data connections over the IPX-P (July 2020).

and Diameter signaling platforms (e.g., roaming signaling equipment unsecured in the public Internet [12], advanced IPX network protocol vulnerabilities [25]) that translate into attacks on end-user privacy or on critical IoT platforms. This brings the obvious challenge of addressing these vulnerabilities in current operational systems, as well as building upon this knowledge to design better solution for next generation signaling platform for data roaming in 5G and beyond. Specifically, the 5G System architecture specifies a Security Edge Protection Proxy (SEPP) as the entity sitting at the perimeter of the MNO for protecting control plane messages, thus replacing the Diameter or SS7 routers from previous generations. The SEPP is meant to enforce inter-MNO security on the N32 interface, and tackle many of the existing vulnerabilities of the existing signaling systems. As we start deploying operational 5G networks, ensuring that the specified requirements for these proxies are met is an important challenge. Privacy in the IPX ecosystem is of paramount importance, especially as cellular IoT devices often underpin critical services that should be protected. These requirements for security, privacy and confidentiality both within the IPX ecosystem (between MNOs and IPX-Ps), and between the IPX ecosystem and the wired Internet bring to light the need for proactive approaches to monitoring the health of the ecosystem, thus tackling anomalies, malicious or unintended.

Furthermore, our work also brings to light the need for novel business models within the IPX ecosystem. We argue that the cellular ecosystem needs to draw from the success of the peering fabric within the public Internet, where the benefits of peering are well known among Internet Service Providers (ISPs) and Content Delivery Networks (CDNs), particularly when it comes to public peering via an Internet Exchange Point (IXP). This established practice in the wired Internet has not yet been fully translated to the mobile Internet [26], where currently only two major IXPs (i.e., AMS-IX and Equinix) offer the mobile peering service, even when more people are connecting to the Internet over cellular connections than fixed broadband. We highlight the need to build a new dynamic of interaction within the ecosystem that would ensure trust among MNOs to guarantee optimal performance for the end-user (e.g., enable local breakout roaming), as well as privacy and confidentiality.

A ACKNOWLEDGMENTS

We thank the SIGCOMM anonymous reviewers, and our shepherd, Z. Morley Mao, for their helpful comments and guidance. We also thank Daniel Hidalgo Pazos (Telefónica Global Solutions) and Carlos Gamboa Bontje (Telefónica Global Solutions) for their invaluable help collecting and analyzing the dataset from the IPX-P system. The work of Andra Lutu was supported by the EC H2020 Marie Curie Individual Fellowship 841315 (DICE). This project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement no.101017109 "DAEMON". The work of Marcelo Bagnulo has been partially funded by the EU Project PIMCITY.

REFERENCES

[1] 2019. *International Tourism Highlights, 2019 Edition.* World Tourism Organization, Madrid, Spain. https://doi.org/10.18111/9789284421152
[2] A. Ali, W. Hamouda, and M. Uysal. 2015. Next generation M2M cellular networks: challenges and practical considerations. *IEEE Communications Magazine* 53, 9 (2015), 18–24. https://doi.org/10.1109/MCOM.2015.7263368
[3] Mostafa Ammar. 2018. Ex Uno Pluria: The Service-Infrastructure Cycle, Ossification, and the Fragmentation of the Internet. *SIGCOMM Comput. Commun. Rev.* 48, 1 (April 2018), 56–63. https://doi.org/10.1145/3211852.3211861
[4] GSM Association. 2012. IPX White Paper. (2012). Retrieved Dec 1, 2019 from https://www.gsma.com/iot/wp-content/uploads/2012/03/ipxwp12.pdf
[5] GSM Association. 2013. IR.34-Guidelines for IPX Provider networks, Version 9.1. (2013). Retrieved Dec 29, 2019 from https://www.gsma.com/newsroom/wp-content/uploads/2013/05/IR.34-v9.1.pdf
[6] GSM Association. 2020. Official Document IR.73 - Steering of Roaming Implementation Guidelines. https://www.gsma.com/newsroom/wp-content/uploads/

/IR.73-v5.0-7.pdf. (May 2020).
[7] Rajesh Bhalla. 2012. Quality of service (qos) over network-to-network interfaces for ip interconnection of communication services. (2012). Retrieved Feb 4, 2020 from https://patents.google.com/patent/US20120218924 [US Patent US20120218924A1].
[8] A. Biral, Marco Centenaro, Andrea Zanella, Lorenzo Vangelista, and Michele Zorzi. 2015. The challenges of M2M massive access in wireless cellular networks. *Digit. Commun. Netw.* 1 (01 2015), 1–19.
[9] European Commission. 2015. New Rules on Roaming Charges and Open Internet. (2015). Retrieved Mar 6, 2018 from https://ec.europa.eu/digital-single-market/en/news/new-rules-roaming-charges-and-open-internet
[10] European Commission. 2017. Roaming charges| What has the European Commission done so far? (2017). Retrieved Dec 29, 2019 from https://ec.europa.eu/digital-single-market/en/roaming-charges-what-has-european-commission-done-so-far
[11] Shannon Doocy, Kathleen R Page, Fernando de la Hoz, Paul Spiegel, and Chris Beyrer. 2019. Venezuelan migration and the border health crisis in Colombia and Brazil. *Journal on Migration and Human Security* 7, 3 (2019), 79–91.
[12] Tobias Engel. 2014. SS7: Locate. track. manipulate. In *FTP: http://events. ccc. de/congress/2014/Fahrplan/system/attachments/2553/or iginal/31c3-ss7-locate-track-manipulate. pdf.*
[13] Victor Fajardo, Jari Arkko, John A. Loughney, and Glen Zorn. 2012. Diameter Base Protocol. RFC 6733. (Oct. 2012). https://doi.org/10.17487/RFC6733
[14] Stefan Geissler, Florian Wamser, Wolfgang Bauer, Michael Krolikowski, Steffen Gebert, and Tobias Hoßfeld. 2021. Signaling Traffic in Internet-of-Things Mobile Networks. In *IFIP/IEEE International Symposium on Integrated Network Management (IM 2021)*. Bordeaux, France (Virtual Conference).
[15] F. Ghavimi and H. Chen. 2015. M2M Communications in 3GPP LTE/LTE-A Networks: Architectures, Service Requirements, Challenges, and Applications. *IEEE Communications Surveys Tutorials* 17, 2 (2015), 525–549. https://doi.org/10.1109/COMST.2014.2361626
[16] Harini Kolamunna, Ilias Leontiadis, Diego Perino, Suranga Seneviratne, Kanchana Thilakarathna, and Aruna Seneviratne. 2018. A First Look at SIM-Enabled Wearables in the Wild. In *Proceedings of the Internet Measurement Conference 2018 (IMC '18)*. Association for Computing Machinery, New York, NY, USA, 77–83. https://doi.org/10.1145/3278532.3278540
[17] M. Laner, P. Svoboda, N. Nikaein, and M. Rupp. 2013. Traffic Models for Machine Type Communications. In *ISWCS 2013; The Tenth International Symposium on Wireless Communication Systems*. 1–5.
[18] Andra Lutu, Byungjin Jun, Fabián E. Bustamante, Diego Perino, Marcelo Bagnulo, and Carlos Gamboa Bontje. 2020. A First Look at the IP EXchange Ecosystem. 50, 4 (2020). https://doi.org/10.1145/3431832.3431836
[19] Andra Lutu, Byungjin Jun, Alessandro Finamore, Fabián E. Bustamante, and Diego Perino. 2020. Where Things Roam: Uncovering Cellular IoT/M2M Connectivity. In *Proceedings of the ACM Internet Measurement Conference (IMC '20)*. Association for Computing Machinery, New York, NY, USA, 147–161. https://doi.org/10.1145/3419394.3423661
[20] Andra Lutu, Diego Perino, Marcelo Bagnulo, Enrique Frias-Martinez, and Javad Khangosstar. 2020. A Characterization of the COVID-19 Pandemic Impact on a Mobile Network Operator Traffic. In *Proceedings of the ACM Internet Measurement Conference (IMC '20)*. Association for Computing Machinery, New York, NY, USA, 19–33. https://doi.org/10.1145/3419394.3423655
[21] Anna Maria Mandalari, Andra Lutu, Ana Custura, Ali Safari Khatouni, Özgü Alay, Marcelo Bagnulo, Vaibhav Bajpai, Anna Brunstrom, Jörg Ott, Marco Mellia, et al. 2018. Experience: Implications of roaming in europe. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*. 179–189.
[22] F. Michclinakis, H. Doroud, A. Razaghpanah, A. Lutu, N. Vallina-Rodriguez, P. Gill, and J. Widmer. 2018. The Cloud that Runs the Mobile Internet: A Measurement Study of Mobile Cloud Services. In *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*. 1619–1627. https://doi.org/10.1109/INFOCOM.2018.8485872
[23] Takaaki Moriya. 2013. Survey of IPX (IP eXchange) as an Emerging International Interconnection between Telecommunication Networks. *IEICE Transactions on Communications* E96.B (04 2013), 927–938. https://doi.org/10.1587/transcom.E96.B.927
[24] N. Nikaein and S. Krea. 2011. Latency for Real-Time Machine-to-Machine Communication in LTE-Based System Architecture. In *17th European Wireless 2011 - Sustainable Wireless Technologies*. 1–6.
[25] K Nohl and L Melette. 2015. Advanced Interconnect Attacks. Chasing GRX and SS7 Vulns. https://media.ccc.de/v/camp2015-6785-advanced_interconnect_attacks.. (2015).
[26] Hideyuki Sasaki. 2017. Peering culture can improve mobile Internet. (2017). Retrieved Jan 22, 2020 from https://blog.apnic.net/2017/05/12/peering-culture-can-improve-mobile-internet/ [APNIC Blog Online].
[27] M. Z. Shafiq, L. Ji, A. X. Liu, J. Pang, and J. Wang. 2013. Large-Scale Measurement and Characterization of Cellular Machine-to-Machine Traffic. *IEEE/ACM ToN - IEEE Conference on Transactions on Networking* 21, 6 (2013), 1960–1973. https://doi.org/10.1109/TNET.2013.2256431
[28] Kelsey Sheehy. 2019. Best International Cell Phone Plans 2019. (October 2019). Retrieved Feb 3, 2020 from https://www.nerdwallet.com/blog/utilities/best-international-cell-phone-plans/
[29] Telefonica Global Solutions. 2019. Making IPX deployment simple and personalized to your needs. https://www.wholesale.telefonica.com/en/information-centre/multimedia/making-ipx-deployment-simple-and-personalized-to-your-needs/. (2019).
[30] Tina Tsou, Ruibing Hao, and Tom Taylor. 2013. Realm-Based Redirection In Diameter. RFC 7075. (Nov. 2013). https://doi.org/10.17487/RFC7075
[31] Narseo Vallina-Rodriguez, Srikanth Sundaresan, Christian Kreibich, Nicholas Weaver, and Vern Paxson. 2015. Beyond the Radio: Illuminating the Higher Layers of Mobile Networks. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '15)*. Association for Computing Machinery, New York, NY, USA, 375–387. https://doi.org/10.1145/2742647.2742675
[32] Claire Volkwyn. 2019. Ed's note: The soap opera that is the UK smart meter rollout. (Dec. 2019). Retrieved Feb 4, 2020 from https://www.smart-energy.com/industry-sectors/smart-meters/the-soap-opera-that-is-the-uk-smart-meter-rollout/
[33] Richard Xu. 2015. Method and Apparatus for Managing Communication Flow in an Inter-Network System. (2015). Retrieved Feb 4, 2020 from https://patents.google.com/patent/US20150222554 [US Patent US20150222554].
[34] Richard Xu, Hwan Jang Tang, and Ajay Joseph. 2014. Method and System For Hub Breakout Roaming. (2014). Retrieved Feb 4, 2020 from https://patents.google.com/patent/US20140169286 [US Patent US20140169286A1].