# Where Things Roam: Uncovering Cellular IoT/M2M Connectivity

Andra Lutu
Telefonica Research

Byungjin Jun
Northwestern University

Alessandro Finamore
Telefonica Research

Fabián E. Bustamante
Northwestern University

Diego Perino
Telefonica Research

## ABSTRACT

Support for "things" roaming internationally has become critical for Internet of Things (IoT) verticals, from connected cars to smart meters and wearables, and explains the commercial success of Machine-to-Machine (M2M) platforms. We analyze IoT verticals operating with connectivity via IoT SIMs, and present the first large-scale study of commercially deployed IoT SIMs for energy meters. We also present the first characterization of an operational M2M platform and the first analysis of the rather opaque associated ecosystem.

For operators, the exponential growth of IoT has meant increased stress on the infrastructure shared with traditional roaming traffic. Our analysis quantifies the adoption of roaming by M2M platforms and the impact they have on the underlying visited Mobile Network Operators (MNOs). To manage the impact of massive deployments of device operating with an IoT SIM, operators must be able to distinguish between the latter and traditional inbound roamers. We build a comprehensive dataset capturing the device population of a large European MNO over three weeks. With this, we propose and validate a classification approach that can allow operators to distinguish inbound roaming IoT devices.

## 1 INTRODUCTION

The infrastructure established by Mobile Network Operators (MNOs) over the last 20 years for person-to-person communications is being leveraged to enable Internet of Things (IoT) and Machine-to-Machine (M2M) services. In particular, support for "things" roaming internationally has become critical for IoT verticals, from connected cars to smart meters, and explains the commercial success of M2M platforms.

M2M platforms benefit from the extensive global network infrastructure that international carriers (e.g., incumbent tier-one operators such as Vodafone, Tata, Telefónica, or Orange) have been shaping for the past decades. They created the so-called Subscriber Identity Module (SIM) for things (or global IoT SIM), which is a SIM provisioned by a single (home) MNO, but operational anywhere in the world through roaming. This is attractive for IoT verticals, as using M2M platforms (i) can result in more stable connectivity/coverage, (ii) allow to avoid the cost of establishing technical and commercial relationships with operators in the countries they deploy, and (iii) application logic is handled in a centralized manner (all SIMs have a single home country) which can simplify management. However, this ecosystem remains largely unexplored in our community.

In this paper, we present the first characterization of the global footprint of an operational M2M platform and the first analysis of IoT SIM deployments in the wild, as part of this rather opaque roaming ecosystem. To do so, we take two different perspectives, as follows.

First, we present a characterization of the global footprint of an operational M2M platform managed by Telefónica, supporting IoT verticals world-wide. Using an 11-day long dataset and comprising a sample of over 100k IoT SIMS, we show both the "centralization" adopted by M2M platforms, as well as the breadth of their operations (Section 3).

Second, we take the perspective of a (visited) MNO – O2 UK– whose role (in this context) is to honor its roaming partnerships and (blindly) connect the IoT SIMs operating (i.e, roaming) in the UK. Our goal is to analyze the impact of roaming things on the MNO resources. We build a dataset that captures both real users and M2M/IoT devices of the large European MNO over a period of 3 weeks (Section 4). Our analysis quantifies the adoption of roaming by M2M platforms and the impact they have on the underlying visited MNOs. Out of 39.6M devices active across the 3 weeks, we find 26% (10.1M) being M2M related, with 75% devices being international roamers (Section 5).

The exponential growth of roaming IoT are key for MNOs, as it could lead to increased stress on the infrastructure shared with traditional roaming traffic. Managing the growing stress of M2M communication would not be a new problem for MNOs, if M2M traffic showed similar characteristics to that of phone traffic (and brought comparable revenues). However, M2M traffic exhibits significantly different features than phone traffic in a range of aspects from signaling, to uplink/downlink traffic volume ratios to diurnal patterns [21]. In other words, though these devices occupy radio resources in MNOs networks and exploit the MNOs interconnections

in the cellular ecosystem, they do not generate traffic that would allow MNOs to accrue revenue (Section 6).

To manage the network and financial impact of M2M traffic, operators must be able to distinguish between this and traditional inbound roaming traffic. This requires some ingenuity, and to support such task, the GSM Association released a binding permanent reference document [2], recommending home networks and carriers to provide transparency of their outbound roaming M2M traffic by sharing information on the dedicated Access Point Name (APN)s or dedicate International Mobile Subscriber Identity (IMSI) ranges they use. In fact, if it is true that MNOs should be able to identify their native devices, *i.e.,* IoT devices that carry an MNO's SIM and connect to the MNO's infrastructure, without a common policy IoT devices identification and classification is not an easy task. In this work, we propose and validate a method based on both device properties, traffic use, and APN strings. We demonstrate this approach for the case of IoT SIMs deployed for energy smart meters (Section 7).

This paper makes the following contributions:

- *We present the first characterization of mobile roaming support for M2M communication.* We describe how M2M platforms build on top of cellular infrastructure (§ 2), and showcase the operation of a large M2M platform (§ 3). We illustrate the sheer size of these platforms with an analysis of the population of IoT SIMs activated/managed by it to support IoT verticals over 4G networks.
- *We show the impact of cellular IoT on visited MNOs.* We build a vast dataset to capture the roaming status of devices connected to a large European MNO for a period of 22 days (§ 4). We introduce an approach for classifying devices into M2M, smartphones, and feature phones. We present general population characteristics, and show that the majority of IoT devices connecting to the MNO's network are roaming (§ 5).
- *We analyze IoT verticals operating with global IoT SIMs, and present the first large-scale study of commercially deployed IoT SIMs for energy meters.* We confirm that IoT SIMs' traffic patterns greatly differ from those of smartphones (§ 6). We focus on smart energy meters, and present the largest (more than 3 million devices) measurement study of smart meters in real-world deployments (§ 7).

# 2 THE ROLE OF ROAMING IN IOT/M2M CONNECTIVITY

In this section, we expose how roaming supports cellular IoT/M2M communications. We close the section with a brief overview of related work in this space.

## 2.1 Roaming Overview

Roaming is one of the fundamental features of the cellular networks ecosystem. It enables clients of one MNO to use the network of another MNO when traveling outside the provider's area of coverage, nationally or internationally.

To support customers of an Home Mobile Network Operator (HMNO) roaming in the network of a Visited Mobile Network Operator (VMNO) both networks must have a commercial agreement. With a technical solution in place, commercial roaming is then possible and MNOs' customers can use their respective partners' networks to extend coverage. MNOs generate roaming revenue by charging their roaming partners as a function of the data/voice/SMS the partner's users (inbound roamers) generate on the visited network. The roaming partners must each record the activity of roaming clients in a given VMNO. Then, by exchanging and comparing these records, the VMNO can claim revenue from the partner HMNO.

In terms of business agreement solutions, the most popular option for MNOs is a standard *bilateral agreement* where the two parties involved define terms and conditions of their cooperation. However, new bilateral roaming agreements for roaming are costly and are generally of lower value today. Even more, smaller and newer operators have great difficulty entering this market and extending their international coverage even for basic voice services.

These challenges have motivated a new model that relies on *roaming hubs*. In this model, operators connect to a hubbing solution provider to gain access to many roaming partners, externalizing the roaming interworking establishment to the roaming hub provider. Hubs are then interconnected to further expand potential operator relationships. The roaming hub solution does not preclude the existence of bilateral agreements between MNOs, and can be viewed as a complement to the bilateral roaming model.

When a commercial agreement exists between two MNOs, there are multiple network configurations to enable roaming between the two networks. Figure 1 presents a set of architecture configurations that can be used for roaming in a mobile network – home-routed roaming (HR), local breakout (LBO) and IPX hub breakout (IHBO). When a mobile node is at home (Fig. 1, left), the *home user*'s traffic will take a short path inside the network to reach a suitable Packet Data Network Gateway (PGW) to the Internet. The traffic of a *roaming user* (Fig. 1, right) is directed to an egress PGW whose location depends on the roaming architecture. Previous work has found that the default roaming configuration currently used in the majority of MNOs in Europe is the HR roaming [15].

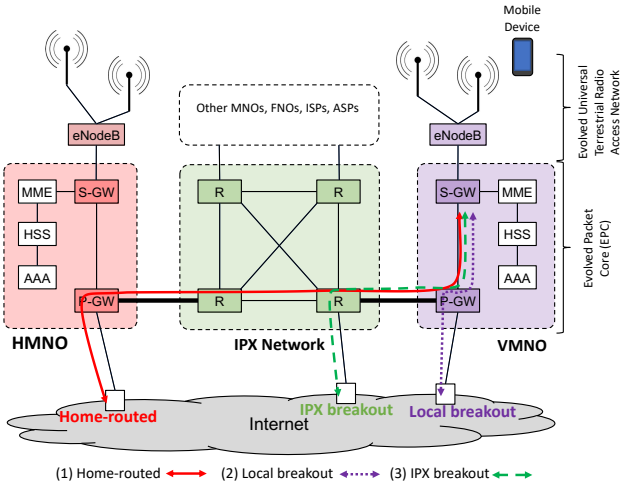


**Figure 1: Network configurations for roaming and interconnection of MNOs through a roaming hub (i.e., IPX Network).**

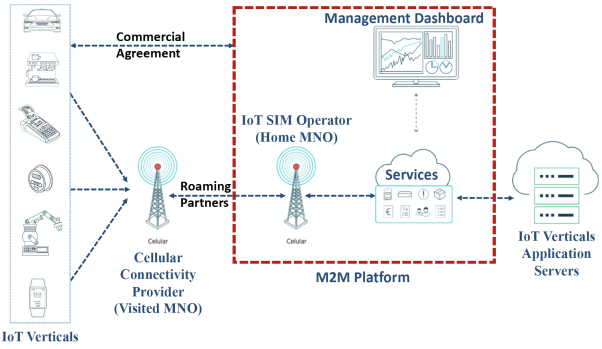


**Figure 2: Overview of the M2M ecosystem and the role of roaming for cellular IoT.**

## 2.2 Roaming for IoT/M2M

Most IoT device manufacturers need a (global) connectivity solution. This motivates them to evaluate communication providers who can ensure data connectivity across the globe, such as cellular connectivity providers. Roaming is thus an essential service for IoT verticals. In Figure 2, we give an overview of the ecosystem around managed connectivity for cellular IoT. We show the main players and detail next their role in connecting different types of IoT devices.

Depending on the use case (e.g., automotive, logistic tracking, smart meters), roaming may be required occasionally or persistently/permanently. Different *IoT verticals* come with potentially different requirements – while logistics services, for instance, may prioritize international roaming to track assets in flux, payment services depend on signal reliability, where terminals always connect, and select an alternative network in the event the first one fails.

*M2M platforms* usually rely on one or several (home) MNOs to provision the global IoT SIM. They then leverage the infrastructure international carriers have been creating through their strategic positioning as MNO interconnection providers and 'roaming hubs' (Fig. 1), to provide IoT businesses with the managed cellular connectivity they require. International carriers bring the potential to serve IoT players in every sphere and bridging the gap to seamless roaming. For example, BICS, one of the largest players in this space, interconnects with about 500 operators and carries about 25% of worldwide roaming traffic, by its own estimates [4]. These global carriers have an important role to establish reliable connectivity — so every vertical can access every place in the world through mobile connectivity, and manufacturers can produce a device in one part of the world that will connect to radio networks in another.

The M2M platform offers a transparent solution to IoT companies, which use the global IoT SIMs to deploy devices in any country without the need to interact with any local cellular operators. However beneficial for the IoT companies, this may prove challenging for the *local VMNO*, whose lack of specific knowledge regarding which inbound roamers represent M2M devices does not allow them to optimally cater to these types of customers. The service-specific levels of support required by roaming smart metering applications may differ from those for an e-book reader, or Low Power Wide Area (LPWA) devices. To support the applications efficiently, a VMNO requires visibility of inbound roamers representing M2M customers, dependent on what device or application is being used, so that it can assess the appropriate service impacts to support that M2M roamer and manage the network efficiencies for M2M. Currently, transparency is provided by the M2M APN, IMSI ranges (full or partial) and, for Narrow Band IoT (NB-IoT) (and other dedicated LPWA platforms), the Radio Access Technology (RAT).

Despite their growing importance, we have a limited understanding of the operational reality of M2M platforms dynamics, and how MNOs support the IoT/M2M communications. A key contribution of this work is illuminating these aspects by analyzing two real-world datasets from an operational world-wide M2M platform, and from an MNO that hosts (i.e., as a VMNO) many devices whose connectivity is provided by different global M2M platforms.

## 2.3 Related Work

Standardization bodies and different working groups have been defining both network structure and services for M2M platforms [10, 11, 24, 25]. Considering mobile networks, two opposite trends currently coexist, one pushing towards repurposing 2G/3G to serve M2M, and the other adopting 4G/5G [8, 12, 16]. Differently from this literature, we take a data-driven approach focusing on the technologies we see deployed in live networks.

Furthermore, we core our analysis of roaming dynamics. Prior literature on cellular network traffic has focused

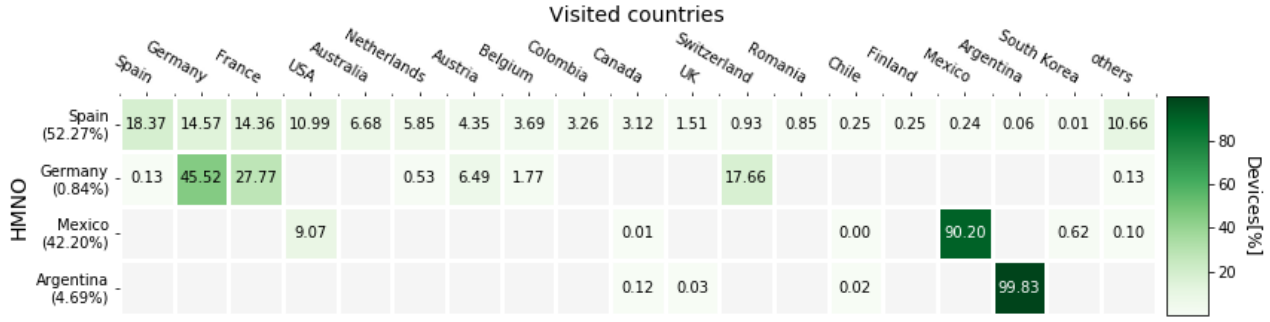


**Figure 3: Percentage of M2M devices per visited country.**

on traditional, people-to-people communication or M2M communication within a single MNO [9, 13, 21, 22]. Vallina-Rodriguez et al. [23] analyzes roaming, primarily national roaming, using crowdsourced measurements. A recent study by Mandalari et al. [15] presents an in-depth characterization of international roaming in Europe, extending the work by Michelinakis et al. [17] limited to two operators. These past efforts have focused on roaming on traditional communication. There is however some literature regarding modeling M2M traffic [14, 19, 20], but it is intrinsically orthogonal with respect to our aim. Hence, to the best of our knowledge, this is *the first study focused on roaming for M2M communication.*

## 3 DYNAMICS OF AN M2M PLATFORM

In this section, we focus on the operational system of the global M2M platform (delimited by the dotted red block in Figure 2) operated by Telefónica, through its extensive roaming partners network.[1] This platform builds on top of an underlying international carrier, and offers the service of global IoT SIM. The global IoT SIM is a SIM from a (home) MNO that operates inside IoT devices world-wide through roaming (provided roaming is allowed in the country of operation). M2M platforms exploit roaming and the underlying carriers to give global connectivity to IoT providers, which ship their devices internationally (from smart meters to wearables and cars) with pre-arranged cellular service.[2] For example, Telefónica's M2M platform provides continuous connectivity for General Motors vehicles in Mexico [6].

The carrier that supports the M2M platform under consideration operates a large infrastructure worldwide, interconnecting directly with MNOs from 19 countries through 40 Points of Presence (PoP), with a predominant presence in Europe and Latin America. It further interconnects with other carriers to extend its footprint to the rest of the globe, and allow roaming on visited networks that are not directly interconnected to its PoPs. We discuss next the main characteristics of an M2M platform, focusing on 4G/LTE connectivity (i.e., we do not capture traffic for 2G or 3G in the dataset). For this, we use a dataset of passively collected signaling activity from IoT devices connected to global networks through the operational M2M platform.

### 3.1 M2M dataset

The M2M dataset we analyze spans 11 days (November 19-29, 2018) and contains 14 million transactions generated by a sample population of over 100,000 4G-enabled IoT devices. The monitoring probes reside in the M2M platform at the service level (see Fig. 2) and capture control plane information, focusing specifically on the attach/detach procedures, as generated by devices connected to the VMNO radio network. Given that few HMNOs issue the global IoT SIMs, the monitoring probes reside close to the infrastructure of the HMNOs. We are using a commercial solution for monitoring the devices that are being provisioned via the M2M platform. The solution is vendor-independent, allowing for the deployment of non-intrusive probes in these routers that form part of the M2M platform.

The dataset does not provide visibility into the data plane traffic, nor do we capture information on the specific IoT vertical served by the M2M platform. Our goal here is to expose the reliance of the M2M platform on roaming to support IoT verticals on top of 4G/LTE networks. This provides visibility on the stress imposed by the dynamics of M2M devices on mobile networks, whose infrastructures offer the basic technological support for IoT/M2M services.

Each transaction in the signaling dataset reports an event generated by an IoT device attempting to connect the 4G radio network of a (visited) MNO, and the dataset represents a sampled view of world-wide M2M infrastructure traffic. More specifically, each transaction reports a unique device ID (a one-way hash), a timestamp, SIM country code ( Mobile Country Code (MCC)) and network code (Mobile Network Code (MNC)), visited country code and mobile network code (VMNO MCC-MNC), message type (either *authentication*,

[1] Telefónica M2M Platform: https://iot.telefonica.com/en/solutions/connect/kite-platform/.

[2] In contrast to an approach where IoT providers make local arrangements to obtain connectivity in *each* country where their devices operate.

*update location* or *cancel location*), and a message result (e.g., OK, RoamingNotAllowed, UnknownSubscription, etc.). In the remainder of the section, we evaluate the passive traces in this dataset to characterize the footprint of (4G) IoT device dynamics globally.

## 3.2 Overall Dynamics

The records captured in the 11-days-long M2M dataset show that there are 4 main HMNOs that support M2M communications through the underlying infrastructure. We denote them by their home countries, namely Spain (ES), Germany (DE), Mexico (MX), and Argentina (AR).

Figure 3 presents the overall distribution of IoT devices on each of the HMNOs. Each row shows one of the HMNOs and each column corresponds to the different visited countries where the IoT devices operate. We breakdown countries having at least 0.1% of devices, and we group the rest into a single class "Other". We normalize cell values by row, while the y-axis labels report the share of devices for each HMNO. We find that two HMNOs support the majority of the M2M communications. In particular, the MNO from Spain provides the SIM cards for 52.3% of all the IoT devices in our sample dataset. Overall, during the entire period of analysis, the devices enabled by the M2M platform with SIMs of Spain were active in 77 different countries, connecting to over 127 VMNOs through the M2M platform. We note that the Spanish MNO is active in a region where free-roaming has been promoted intensively through regulation [1].

The second most important HMNO supporting the operations of the M2M platform is Mexico, with 42.2% of all devices operating with a SIM card belonging to this MNO. These IoT devices spread in 7 countries and connect to 10 VMNOs overall. Note, however, that the large majority (90%) operate in their home country and are not roaming. This is due to the local restrictions on roaming in countries in Latin America. Argentina, much in a similar manner to Mexico, has 4.7% of devices (with 6 visited networks), and nearly all of its traffic is not roaming.

The fourth HMNO we identify, the German MNO, supports a relatively small number of devices in our sample (around 1,000), but the number of visited networks is large (18 VMNOs). This might be explained by the requirements of the specific IoT vertical. For example, connected cars have high mobility requirements that would explain the need for seamless coverage, thus generating numerous signaling procedures from the devices and requiring alternative connectivity from multiple networks [7].

Given that the Spanish MNO supports a large portion of IoT devices in our dataset, we continue our analysis on the dynamics of the M2M platform by capturing only IoT SIMs this HMNO provides, which is either local (non-roaming) or global (roaming). For the Spanish network, roaming extends coverage over 76 countries and also generates large amounts of signaling traffic (81.8% of all signaling traffic in our dataset comes from ES-powered IoT devices). We verify that 92% of these messages are triggered while devices are roaming. Conversely, only 8% of the signaling traffic we capture from these devices occurs when they attach to the HMNO, even though the fraction of non-roaming devices is relatively high (18%). This suggests that IoT devices active in their native home country are potentially less mobile than the roaming ones and are connected over longer periods of time.

For the roaming devices supported by the Spanish MNO (82%), we find that 75% of the signaling traffic comes from 62% of devices. This covers operations over only 5 visited countries and 10 visited MNOs. The geographical distances between the HMNO and the VMNO are not always small (e.g., Spain to Australia), pointing to potential performance penalties in the case of HR roaming [15]. In this case, however, the M2M platform uses different roaming configurations in order to optimize the performance of IoT devices roaming in very far destinations. This analysis is, however, outside the scope of this work.

## 3.3 Device-level Dynamics

We now focus our analysis on the device-level signaling traffic patterns of IoT devices connecting with a global IoT SIM for the Spanish provider. Specifically, we look at the frequency of three procedures we monitor (Update Location, Authentication, and Cancel Location). Each record has a status message associated, describing the outcome of the procedure, such as "OK," or the error message when it failed. We find that in this IoT device population, 40% of devices trigger failed signaling procedures against the 4G/LTE networks with the following error messages; Feature Unsupported, Roaming Not Allowed or Unknown Subscription. For the rest of 60% IoT devices connecting through the Spanish MNO, we register at least one successful procedure in our dataset. This is a non-negligible number of IoT devices generating traffic through the 4G signaling infrastructure by attempting (and failing) to use 4G connections.

We further investigate the amount of signaling traffic per roaming IoT device, the distribution of the number of VMNOs used, and the frequency of inter-VMNO switches (see Figure 4). First, we note that the distribution of the number of signaling records per device has a long tail, showing the wide range of signaling patterns the M2M dataset captures (Fig. 4-left). We show this distribution for all IoT devices, and for devices successfully connected to the 4G network (4G devices), roaming devices, and non-roaming (native) devices. From the distribution on all devices, the average load is of 267 signaling records overall, with 97% of

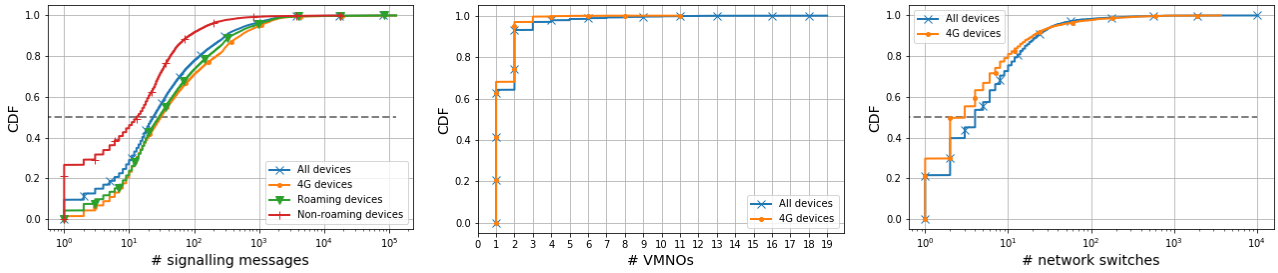


**Figure 4: M2M Platform dynamics (left) distribution of total number of signaling records (center) VMNOs used; (right) inter-VMNO switches.**

devices triggering less than 2,000 signaling procedures over the 11 days period, and a very small fraction of IoT devices flooding the signaling network with as many as 130,000 messages. We note the difference between roaming and native devices, with the former generating 10 more procedures than the latter in the median.

Figure 4-center shows the number of VMNO the roaming IoT devices use over the observation window. We find that 65% of roaming IoT devices use only one VMNO, while more than 25% roaming IoT devices switch between two VMNOs. Only 5% of roaming devices require coverage from more than three VMNOs. Interestingly, for some of the IoT devices with only failed signaling procedures, we find that the maximum number of attempted VMNOs is as high as 19 mobile networks. This shows high international mobility requirements and the need for reliable seamless coverage, which are indeed difficult to guarantee with only 4G/LTE connectivity in some regions. These devices fall-back on 2G/3G coverage (which our sample dataset does not capture).

For those IoT devices with at least two VMNOs (35% of IoT devices), we examine the number of inter-MNO switches. Figure 4-right shows a mixed result. For approximately 50% of IoT devices, we register a maximum of two VMNOs switches during the total 11 days. However, for 20% of devices, the inter-VMNO switches happen at least once a day. Approximately of 3% devices present high frequency in switching between VMNOs, namely from 100 times to 3,000 times during the period we monitor. Again, we do not have visibility into the IoT vertical using these devices, but their high mobility and their requirements for reliable coverage are clear.

# 4 VIEW FROM AN MNO

In this section, we focus our analysis on from the point of view of an MNO (i.e., visited MNO in Figure 2) that (blindly) supports a large number of IoT devices as inbound roamers, from multiple different M2M platform around the world (including Telefónica's M2M platform). Specifically, we analyze the device population of O2 UK, a large MNO in the UK.

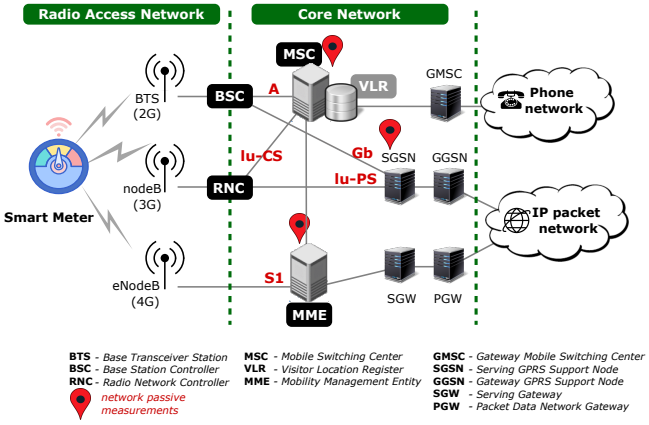


**Figure 5: High-level architecture of the measurement infrastructure integrated in the cellular network.**

## 4.1 MNO dataset

The cellular network we study supports 2G, 3G, and 4G mobile communication technologies. In Figure 5, we illustrate a high-level schema of the MNO architecture. Such a network can be simplified to consist of three main domains: (i) the cellular device (in our case, the smart meter), (ii) the Radio Access Network (RAN), and (iii) the Core Network (CN). Our passive measurement approach relies on commercial solutions integrated within the MNO's infrastructure. The red pins in Fig. 5 mark the network elements that we monitor, namely the Mobility Management Entity (MME), the Message Sequence Chart (MSC) and the Serving GPRS Support Node (SGSN). We collect control plane data on the total population of devices connected to the MNO's radio network. This includes both native devices (operating with a SIM card provisioned by the UK MNO) or inbound roaming devices (operating with a SIM card provisioned by a foreign MNO, from outside the UK, such as the global IoT SIM provided by the M2M platform).

The dataset we use provides a comprehensive view of the entire population of mobile devices connected to the MNO's network over a period of 22 days in April 2019 (from April 5th to April 26th, 2019). These include both smartphones and

feature phones, as well as devices deployed for IoT verticals (e.g., smart meters). This population also integrates users with different roaming status, including MNO's *native* users (active either in the home country or abroad), the users of MVNOs that operate on top of the MNO's infrastructure, and foreign users that belong to other MNOs (national or international), but that use the radio network of the MNO under analysis.

We start introducing the raw data collected by the operator, and then we discuss how we label devices based on their roaming category, and how we identify IoT devices.

**Radio interfaces.** We process logs reporting on activities on IuCS, IuPS, A, and Gb radio interfaces. Those carry events generated by the devices connecting to the radio sectors, and requesting resources for either data or voice communications. Each event carries the anonymized user ID, SIM MCC and MNC, Type Allocation Code (TAC)[3], the sector ID handling the communication, timestamp, event type, event result code. Such events are captured for all connected devices, except for outbound roamers (in this case, radio signaling for outbound roamers is carried over the visited country network only).

**Service usage.** We use Call Detail Records (CDRs) and eXtended Detail Records (xDRs) to provide aggregate service usage for calls and data. Each record reports the anonymized user ID, MCC, and MNC codes for both device SIM and visited country, timestamp, duration, and bytes consumed. Data records also report APN strings, which usually encode information about the specific service/business they relate to. Notice that differently from radio logs, CDRs/xDRs contain traffic also for outbound roamers. These are usually used by the roaming partners to trigger the process of revenue retrieval from roaming (see § 2).

**Device properties.** We also consider a commercial database provided by GSMA. This catalog maps the device TAC to a set of device properties such as device manufacturer, brand and model name, operating system, and radio bands supported.

**Daily devices-catalog.** We combine the three data sources to create a daily list of active devices and associated properties and traffic characteristics. We refer to this aggregate view as devices-catalog. Each record reports a device ID, total number of events, calls, bytes seen, SIM MCC/MNC, list of visited MCC-MNC, a list of APN strings, device manufacturer, device model, device Operating System (OS). We summarize the radio activity into *radio-flags*, a series of three 1-bit flags which are set to 1 if the device has successfully communicated with 2G, 3G, 4G sectors respectively on radio interfaces. Finally, we compute mobility metrics for each device; namely, we calculate the gyration, which captures the distance devices travel throughout a day.

[3]The first 8 digits of the device IMEI, which are statically allocated to device vendors.

## 4.2 Roaming Labels

To capture the roaming status of the MNO's population, we label each device as either *native*, *inbound* roamer, or *outbound* roamer. A device is native if it carries an MNO's SIM and connects to that same MNO's radio network. When such devices connect to a different operator network (either within the same country, or when traveling outside the country) they become outbound roamers (national or international, respectively). Conversely, an inbound roamer is a device operating with a SIM card not belonging to the MNO whose radio network is actually using.

To capture these variants, we tag each record in the devices-catalog with a *roaming label* `<X:Y>`, where `X` relates to the device SIM, and `Y` to the visited network. Specifically, given a SIM card, we assign to `X` four possible values: **H** (*home*, the SIM belongs to the MNO we analyse), **V** (*virtual*, the SIM belongs to an MVNOs enabled by the MNO we analyse), **N** (*national*, the SIM belongs to another MNO in the same country as the current MNO), **I** (*international*, the SIM belongs to an MNO in a country different than the one of the MNO under study). Instead, we assign to `Y` only two values: **H** (*home*, the SIM is attached to the current MNO), **A** (*abroad*, the SIM is attached to a foreign MNO outside the country of the MNO under study).

Overall, we define 6 different roaming labels. For example, the `H:H` label reflects a device that uses the MNO's SIM card and is attached to the MNO's network (i.e., native user), while the `I:H` label refers to devices connected to the MNO under study but with a SIM of operators of different countries than the MNO one (i.e., inbound roamers). Using these labels, we further breakdown devices per roaming status. As expected, we find that the majority of devices are native, i.e., either MNO (about 48% per-day) or MVNO (about 33% per-day) devices connected to their home MNO. However, we find that the third largest population is international inbound roamers (about 18% per-day).

## 4.3 M2M Device Classification

As previously mentioned, the devices-catalog includes all devices connected to the MNO network. This encompasses devices used by people as their main personal device (e.g., smartphones, feature phones), as well as devices IoT verticals use to support their applications (e.g., car manufacturers, energy companies). Supporting our analysis from the point of view of the M2M platform (§ 3), we aim to establish whether IoT devices are usually roaming internationally. To do so, we to split the devices into three classes: *smart* (for smartphones), *feat* (for feature phones), and *m2m* (for IoT/M2M devices). Prior work [21] demonstrated that using device properties one can perform classification, especially to spot M2M devices, at the cost of some manual verification. The

GSMA database already offers a device classification label, but devices other than smartphones are mostly marked as "modem" or "module" which might not necessarily imply an M2M/IoT application. Furthermore, across the 22 days, we observe 2,436 device vendors, and 24,991 device models across the whole population (i.e., a manual classification as operated in [21] is not feasible).

One possible approach to reduce the classification complexity is to focus on "big players" only. For instance, Gemalto, Telit, and Sierra Wireless are among the top device vendors with a combined 75% of all roaming devices in the dataset. Similar considerations can be made to identify smartphones and feature phones, but we argue that this is a naïve approach as still requires to investigate a large number of devices to validate the classification.

APNs string can be a significant aid for strengthening the confidence of the classification as they hint the vertical used by a device. For instance, *smhp.centricaplc.com.mnc004.mcc204.gprs* hints to Centrica[4], a company working in the energy vertical, i.e., the devices using such APN are possibly smart meters. Notice also the MCCMNC revealing the home country and operator (20404 = Vodafone Netherlands).

We find a total of 4,603 APN strings in the dataset. However, ranking the APNs by number of devices using it, we identified 26 "keywords" in the APN string which we mapped to M2M/IoT verticals using information found online (e.g., *scania* - automotive company, *rwe* - energy company, *intelligent.m2m* - global IoT SIM provider). Using these 26 keywords we obtained 1,719 APNs, while the other are either generic labels related to mobile operators (2,178 labels likely related to consumer services), or other IoT services we could clearly identify.

Finally, we classify the devices combining both APNs and device properties. We start marking as *m2m* all devices using the validated APNs. Then, we extend the *m2m* class to all devices having the same properties of the devices using the validated APNs. For *smart* and *feat* we still use a set of APN strings, but we take advantage of 2 labels properties defined in the GSMA database (e.g., device manufacturer, and operating system). Specifically, we classify a device as *smart* if declared to be using a major smartphone OS (android, iOS, blackberry, windows mobile) and use a consumer APN (e.g., a string contain keywords such as payandgo). Instead, we classify a device as *feat* if the GSMA database declares it to be a feature phone or uses a consumer APN.

Out of the 39.6M devices active across the 22 days, we find 24.4M (62%) *smart*, 3.1M (8%) *feat*, and 10.1M (26%) *m2m*. We label the remaining 2M (4%) as *m2m-maybe* as the device properties suggest they are neither smartphones nor feature phones, but we don't have APNs for them, i.e., those devices only use voice services (the APN is provided only when the device connect for data services). This does not preclude them from possibly being M2M related, but based on the information available we are not able to provide a final classification. Hence, we do not consider those devices for the remainder of the analysis.

Differently from [21], using APNs is useful to increase the classification confidence, and reduce the number of manual investigations. However, when used in isolation, APNs are not enough as we find about 21% of the devices in the dataset not having any APN. This justifies our multi-steps classification process (keywords→APNs→device properties).

## 4.4 Smart Meters Dataset

In this section, we describe the dataset we build specifically for smart meter devices connected to the radio network of the MNO we monitor. We use this dataset to focus our analysis of the roaming things on a specific vertical (i.e., energy), and compare it with the analysis of the general population of devices, or other vertical (i.e., connected cars). Smart Grid applications have received increasing attention in the past years, with regulation pushing for mandatory deployment of metering devices in consumer premises. Specifically, the UK Government is committed to ensure that every home and small business in the country is offered a smart meter by 2021, with more than 12 million devices already deployed at the end of 2018[18] as part of the Smart Metering Implementation Programme (SMIP).

The mobile operator we study provides connectivity for a set of smart meters in the UK under the SMIP framework (*i.e.,* SMIP native devices). Based on private communications, we learned that the MNO uses a dedicated IMSI range for the SIMs installed in smart meters. Moreover, the operator has dedicated resources for these SIMs. The rationale of this choice is to control the impact of such devices on the native users, as well as better control the performance of the smart meter network. We further denote this dataset as *SMIP native*. In other words, the MNO considers SMIP as a special class of devices. However, we cannot generalize this setup to other operators and countries. Indeed, depending on contract agreements and provisioning preferences, an energy provider might prefer a global M2M platform to serve smart meters, which, in turn, would result as actual roaming devices for the visited MNO.

We further investigate whether there are other devices connecting to the MNO's radio network that are smart meters. Specifically, we aim to identify Global IoT SIMs that connect smart meter devices.

To identify these devices, we rely on the classification we explained above. Specifically, in the APN strings of these inbound roaming devices, we are able to identify patterns

[4]https://www.centrica.com/about-us/what-we-do/our-strategy

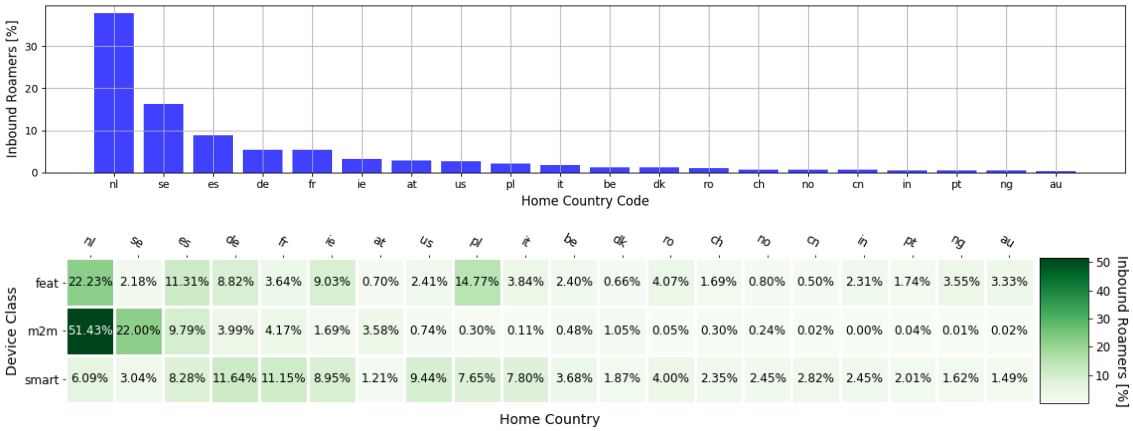


**Figure 6: Home country of inbound roaming devices (top) overall share; (bottom) device class breakdown.**

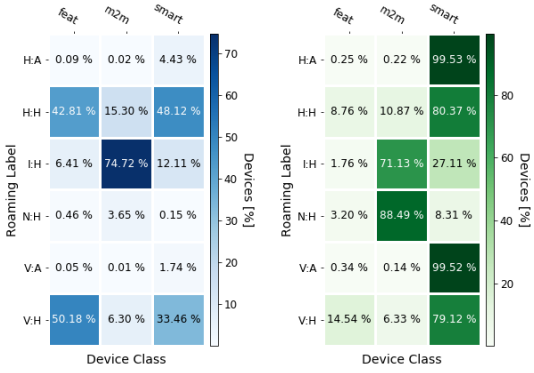


**Figure 7: Devices breakdown (left) Device class -vs- Roaming label; (right) Roaming label -vs- Device class.**

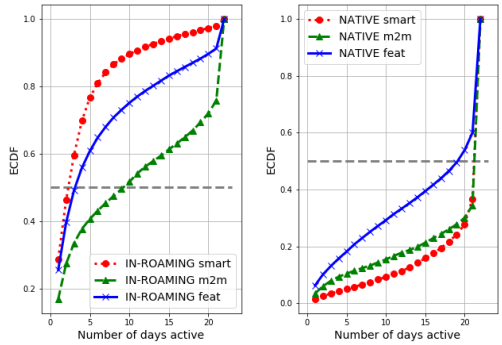


**Figure 8: Number of days devices are active (left) inbound roaming; (right) native.**

that confirm the use of these devices as smart meters by energy companies in the UK. We are able to identify different patterns in the Network Identifier part of the APN string that relate to large energy companies in the UK, including Elster, RWE, Centrica PLT, or General Electric. Using these, we are able to separate the inbound roaming devices that are smart meters. Surprisingly, all the Subscriber Identity Moduless (SIMss) we identify are provisioned by the same cellular operator in the Netherlands. To further validate our inference, we use the Global System for Mobile communications (GSM) Association (GSMA) TAC data catalog to identify the manufacturers of these devices. We find that these devices map to only two manufacturers mainly specialized in M2M modules, namely Gemalto and Telit. We further denote this dataset as *SMIP roaming*.

# 5 M2M POPULATION PROPERTIES

With the processed dataset, in this section, we investigate the home country of the devices, if they are constantly connected to the (visited) MNO network, and if they are stationary or moving. To better highlight those properties, we contrast M2M devices against smartphones and feature phones.

## 5.1 Device Class and Roaming Label

Figure 7 shows heatmaps of the distribution of devices per roaming label and per device class, normalized by device class (Fig. 7-left) and by roaming category (Fig. 7-right). Considering inbound roamers (I:H), 71.1% are M2M device, while 27.1% are smartphones (right heatmap). This further supports the popularity of supporting IoT verticals on top of the roaming infrastructure of cellular providers (§ 3).

Considering the device classes (left heatmap), 74.7% of M2M are inbound roaming, while the rest are either native (H:H) or related to the MVNO (V:H). Instead, for smartphones and feature phones the trend is almost reversed: only 12.1% and 6.4% respectively are inbound roaming, while those device classes are either native of MVNO related.

## 5.2 Home Country

Figure 6 shows the distribution of inbound roamers with respect to their home country. We first break down the whole population per home country, regardless the device class (Fig. 6-top). The top 20 home countries contribute more than 93% of all inbound roaming devices, with the top 3 (Netherlands, Sweden, and Spain) accounting for about 60%.

Figure 6-bottom further detail the breakdown of each home country with respect to the different device classes. Columns are ordered to match the histogram in Fig. 6-top. We normalize the values by device class (i.e., per row), using the total number of inbound roaming devices per class, although we show only the top 20 countries, discarding the long tail of the distribution. We see that 83% of M2M devices use SIMs from the operator from either the Netherlands, Sweden, or Spain; for smartphones and feature phones is 17% and 35% respectively. In other words, the distribution per home country for the M2M devices is significantly more skewed than for the other two classes, further corroborating the "centralization" of M2M platforms.

## 5.3 Spatio-Temporal Dynamics

We further analyze how long M2M devices are active in our dataset. For this, we count the overall number of days the device is generating data, voice, or signaling traffic. Figure 8-left plots the empirical CDF of the number of active days for two device classes, M2M and smartphone devices in the inbound roaming class. Considering inbound roamers (left plot), IoT devices (category *"m2m"*) are active 4.5x longer than smartphones as a median (9 days for M2M devices and 2 days for smartphones), while the 2 device types present similar properties if they are native devices (right plot). When aggregating this information regardless of the roaming category, we note that M2M devices spend less time connected to the network than smartphones. We conjecture that this can be due to the roaming nature of those devices which enables them to switch network if/when needed, or could be the application logic itself which uses the network only when needed. Unfortunately, the dataset does not offer sufficient details to unravel this aspect.

In Figure 9 we investigate the mobility of the different device classes. For this, we evaluate the radius of gyration for the device, capturing the area the device usually travels. We use the physical coordinates of the cell sectors to which devices connect to as a proxy of the actual device position, and compute a centroid (an aggregate representation of where in the country the device was located) and the gyration radius (indicating how far from the centroid the device was moving). Both are weighted based on the time spent connected to each cell sector by the devices. We compute daily metrics

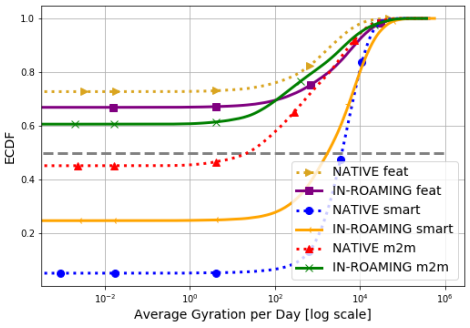


**Figure 9: Radius of gyration comparison.**

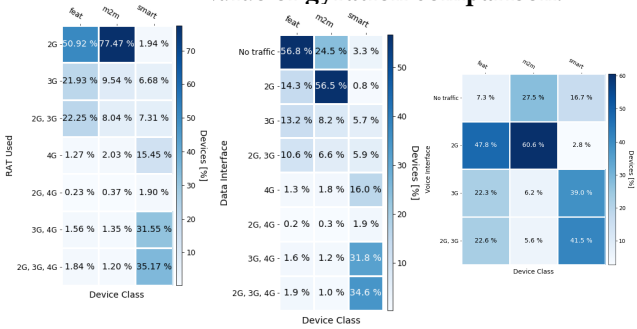


**Figure 10: Devices share with respect to services (left) connectivity; (center) data traffic; (right) voice traffic.**

and present averages across days. Results confirm expectation, i.e., the M2M inbound roaming devices are in majority stationary, with only 20% devices present a gyration larger than 1km (some likely due to cell reselection, rather than actual movements).

# 6 M2M TRAFFIC ANALYSIS

In this section, we continue the analysis of the MNO dataset, investigating M2M communication patterns. We present next how "things" are actually using the cellular network, which is the radio technology on which they depend most and how much traffic they generate.

## 6.1 Device Network Usage

Using the device activity on the data or voice interfaces per RAT, we generate a view of the patterns for each of the device classes (see Fig. 10-left). We find that the vast majority of M2M devices (77.4%) are active on the 2G network only, while smartphones have mostly 3G and/or 4G capabilities. Similar to M2M devices, feature phones are also dependent mostly on the 2G radio network (50.9%).

When focusing on M2M devices voice usage (Fig. 10-right) we find that 60.6% use the 2G network, but 27.5% do not generate any voice traffic. Furthermore, when checking just the activity of the devices in the three different classes on the data interfaces (Fig. 10-center) per RAT, we find that

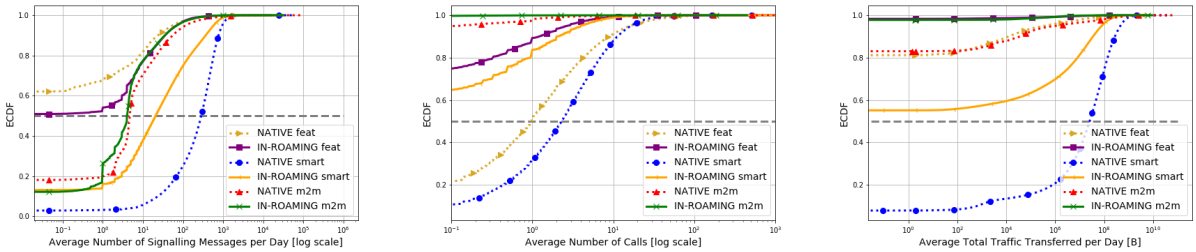


**Figure 11: Traffic analysis for in-roaming and native M2M devices (left) signaling; (center) calls; (right) data usage.**

56.7% of all the M2M devices are indeed only active on the 2G data interface. Interestingly, we note that 24.5% of the M2M devices are actually not active on the data interfaces of the cellular network, relying only on voice communications. Notice also how 56.8% of feature phones do not generate any data traffic, but only 7.3% of them do not generate voice traffic, i.e., as expected, those type of devices are commonly used for calls. The sustained dependency of M2M devices and also features phones on the 2G network bring to light the discussion around the need of MNOs to keep maintaining the legacy technology. Some MNOs (e.g., AT&T) already shut down 2G services, while others announced their target dates.[5]

## 6.2 Traffic Volumes

We analyse the amount of radio resource management signaling events from devices, the number of voice calls, and the amount of data traffic the M2M devices generate compared to smartphones (Fig. 11). Over the three weeks, we find that the number of resource management events the M2M devices trigger is much smaller than the traffic generated by smartphone devices, regardless of their roaming configuration (Fig. 11-left). This is partially explained by the fact that IoT devices are more stationary compared to smartphone devices (Fig. 9). Feature phones, however, generate less signaling traffic than even M2M devices, most likely due to the lack of data services usage.

We further check the average number of voice calls per day for the different device categories (i.e., native M2M, inbound roaming M2M, native smartphones, and inbound roaming smartphones). In Fig. 11-center we show that, although for the majority of M2M devices we do not find any calls registered, there is a small fraction for which the number of voice calls is non-null (regardless of their roaming configuration). We conjecture these might be due to M2M security applications (e.g., emergency elevator services, home security).

Finally, we analyse the total volume of data traffic the different categories of devices transferred in different roaming configurations (namely, native and inbound roaming). Figure 11-right shows that inbound roaming M2M devices generate a very small amount of data traffic, similar to inbound roaming feature phones. Some native M2M devices show non-null data traffic usage (20% of devices generate more than 1 Byte of data on average per day). We note that they have a very similar pattern of data traffic usage with feature phones. There is a clear difference in people's behavior while roaming, which we extract from the comparison of smartphones native to the home country of the MNO and the inbound roaming smartphones. We assume that the decreased volume of traffic for inbound roamers is due to fear of potential bill shock the users might incur when traveling outside their home country (non-EU).

[5]https://1ot.mobi/blog/2g-and-3g-networks-are-shutting-down-globally

# 7 THE CASE OF SMART METERS

As requirements differ between IoT verticals, in this section, we investigate two of the most prominent: smart meters, and connected cars. In particular, we study their connectivity and mobility, signaling volume, and data volume.

## 7.1 SMIP Device Activity

Smart meters are the next generation of gas and electricity meters, and their deployment aims to deliver a much needed digital transformation of the energy system. Different from traditional meters, they record consumption with high frequency (e.g., hourly) and report such information to smart grid infrastructure over a cellular connection. Though stationary, these devices require reliable service from connectivity providers, making the M2M platform a perfect solution. To guarantee reliability and prepare the network for increased deployment, it is important to characterize the traffic patterns of smart meters.

We measure SMIP signaling activity looking at the patterns of Attach, Routing Area Update, and Detach signaling procedures that we capture from the passive monitoring of MSC and MME elements (see Figure 5). Such traffic is known as "background traffic", and it does not bring profit to the service provider, while leading to possible overheads, which we show is significant for inbound roaming devices.

Figure 12(left) reports on the number of days SMIP devices have been active (i.e., they triggered at least one signaling event per day, either on data or voice interfaces). We split

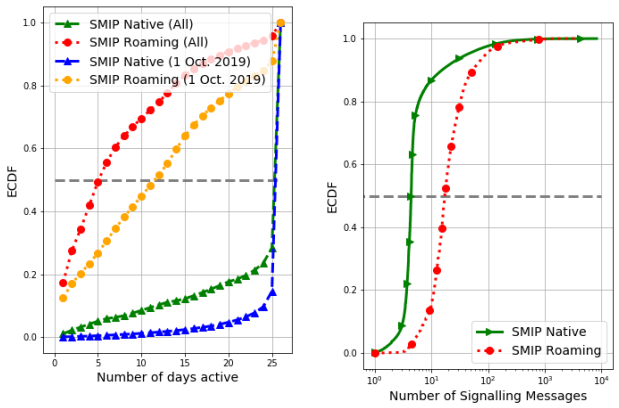


**Figure 12: Device activity for SMIP Native and SMIP Roaming groups, 1-26 October 2019: a) Total number of days SMIP devices were active during the period we study. We show the active time for the total set of devices detected in October 2019, as well as the active time of the devices detected on October 1 across the entire period of analysis. b) Average number of signaling messages per SMIP device per day.**

devices between native and inbound roamers, and report on their activity across the whole period, as well as on the devices being active from the first day of the time period. We can see that native devices have long-lasting connectivity (73% are active for the whole period), while the opposite is true for roaming devices (50% are active only up to 5 days). We conjecture that this is a side-effect of the fundamentally different manner in which they connect to radio resources: roaming devices are free to connect to any UK operator, while native devices rely exclusively on the MNO we study. Figure 12(left) also shows the effect of the ongoing deployment of SMIP devices. Notice indeed how the fraction of constantly active native devices increases to 83% when considering the ones active on the first day of the dataset.

Figure 12(right) reports instead on the generated background traffic. Interestingly, notice how roaming SMIP generates on average ten times more signaling messages than native ones. This considers all the signaling events associated with the smart meters, regardless of these procedures being successful or not. When considering only the failed events, only 10% of all SMIP devices registered to the MNO during October 2019 had at least one failed signaling message, but this increase to 35% when considering roaming devices. Unfortunately, we do not have sufficient data to understand if the increased background traffic is a side effect of roaming, or the roaming itself is a symptom of something deeper, such as network coverage issues.

Looking at the supported radio technologies (see Section 4), all SMIP roaming devices are only 2G capable; this is confirmed also looking that the RAT used by the devices. Conversely, native SMIP support both 2G and 3G, but 2/3 of them use only on 3G, while the rest uses both 2G and 3G connectivity.

## 7.2 Traffic Analysis for IoT Verticals

Using the exposed APN information from inbound roaming IoT devices in our MNO dataset, we separate devices mapping to connected cars. We further use this dataset to contrast against the traffic patterns of smart energy meters. In Figure 13 our analysis shows that connected cars are very similar to normal inbound roaming smartphones, with high mobility patterns (left) large volume of signaling traffic (center) and data traffic (right). At the same time, smart energy meters are IoT devices with completely different behavior. As expected, they are stationary devices that generate very little signaling traffic as well as data traffic, when compared to the connected cars. These patterns validate our intuition on the manner in which these two groups of IoT devices use the visited network. However different, their similar need for reliable connectivity makes them major customers for M2M Platforms.

# 8 DISCUSSION

Our analysis combined two datasets – one from an operational M2M platform, and the other from an operational MNO – to shed light on the dynamics around roaming for M2M communications. This is, to the best of our knowledge, the first analysis of how roaming supports IoT/M2M connectivity world-wide, complementing prior work. Despite this, our view is still limited to the footprint of the system we analyze. Thus, our analysis has a strong European focus, where roaming is heavily used, and regulation efforts are pursuing the activation of "silent roamers" [5].

Although different technical roaming configurations (i.e., HR, LBO, IHBO) might be used for different IoT verticals (allowing the M2M platform to respond to specific QoS requirements), from the M2M platform dataset we currently lack visibility into these details. We complement this analysis investigating the traffic of more than 3 million UK smart meters comparing the ones configured in HR roaming with ones native to the MNO. This IoT vertical account for the largest number of devices compared to the other IoT verticals we were able to identify (e.g., connected cars). Previous work characterized different verticals such as wearables [13] or connected cars [7] highlighted the difference in terms of requirements of these applications and the corresponding devices used, but did not highlight their reliance on roaming.

Finally, our analysis of the global M2M platform relies on 4G signaling information we collected from more than 100,000 IoT devices. As MNOs across the world move to phase out 2G/3G support, IoT verticals will likely rely on

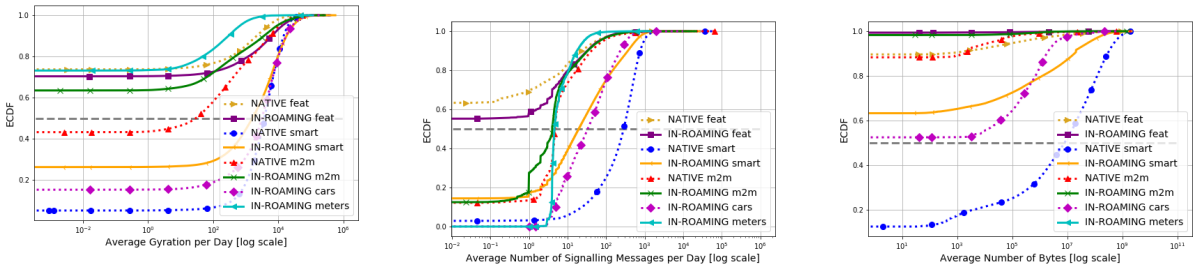


**Figure 13: Connected cars and smart meters traffic patterns (left) mobility; (center) signaling; (right) data usage.**

more sustainable technologies such as 4G/LTE, driven in part by sectors like the connected automotive industry, in which seamless, cross-border, ultra-reliable, low-latency connectivity is of paramount importance. In countries such as Japan, South Korea, Singapore, or Australia, MNOs have already switched off 2G. MNOs in Europe are reportedly planning to retire their legacy 2G/3G networks starting 2020. However, our results from characterizing the IoT devices connected to the European MNO show that IoT devices such as smart meters are currently active mostly in 2G or 3G networks. We have been engaging in private communications with the operator about the implications of these findings. These results opened the discussion within the operator around billing models for IoT devices (especially for "permanent roamers"), given their impact on the cellular ecosystem. Within the GSMA, there is an active working group debating the need for such billing models [3].

Given its power to support the commercial success of IoT, roaming is coming to other IoT technologies. For example, NB-IoT is a low-power wide-area network technology developed for the huge concentration of connected “things” that receive and transmit only small amounts of data, but do so over long periods of time, such as smart meters. The GSMA announced the first international NB-IoT roaming trial back in June 2018, with numerous others having taken place since. The planned deployment of NB-IoT coupled with roaming support will likely create a powerful environment to support the growth of IoT. Moreover, NB-IoT will enable visited MNOs to easily detect the inbound roaming IoT devices, a task that currently is challenging.

## 9 CONCLUSIONS

In this paper, we reported on the role of roaming in enabling IoT, and offered the first characterization of a global M2M platform. We showed how such solutions leverage the maturity of the roaming infrastructure to provide reliability and ubiquity to sustain consistent communications for IoT verticals, such as smart energy meters. Despite its exponential growth, IoT also translates into increased stress on the infrastructure of the (visited) MNOs to which they connect. Our analysis of the device population of a UK MNO showed that M2M devices account for 26% of connected devices across 3 weeks, out of which, 75% are inbound roamers. Though they are connected for longer periods than people roaming, these devices generate very little traffic. In other words, these devices occupy radio resources in the MNO’s network and exploit the MNO’s interconnections in the cellular ecosystem, but they do not generate traffic that allows the MNO to retrieve the corresponding revenue. In a market expected to reach 75.44 billion connected devices worldwide by 2025, (i.e., almost 10x the estimated world population) this puts in perspective the importance of the M2M platform and the corresponding international carrier in supporting the relationships between VMNOs and IoT verticals.

## A ETHICAL CONSIDERATIONS

Both datasets used in this work are collected from operators and covered by NDAs prohibiting any re-sharing with 3rd parties even for research purposes. Raw data has been reviewed and validated by the operators with respect to GPDR compliance (e.g., no identifier can be associated with a person), and all analyses performed to report on aggregated metrics only. The data collection and retention at network middle-boxes and elements are in accordance with the terms and conditions of the operators and the local regulations.

## B ACKNOWLEDGMENTS

We thank the IMC anonymous reviewers, and our shepherd, Nina Taft, for their helpful comments and guidance. We also thank Daniel Hidalgo Pazos (Telefonica Business Solutions) for his invaluable help collecting and analyzing the dataset from the M2M platform; and Javad Kangosstar (O2 UK) for his continued support while processing the MNO’s dataset. The work of Andra Lutu was supported by the EC H2020 Marie Curie Individual Fellowship 841315 (DICE).

# REFERENCES


[1] [n. d.]. European Commission: New Rules on Roaming Charges and Open Internet. https://ec.europa.eu/digital-single-market/en/news/new-rules-roaming-charges-and-open-internet. ([n. d.]). [Online; accessed 06-March-2018].

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

[3] [n. d.]. GSMA Billing and Charging Evolution. https://www.gsma.com/aboutus/workinggroups/interoperability-data-specifications-and-settlement-group/billing-and-charging-evolution. ([n. d.]). [Online; accessed 25-September-2020].

[4] 2017. *BICS Global IoT Whitepaper.* Technical Report. BICS. https://bics.com/download/global-iot-wp/

[5] 2019. *"Roam Like at Home" Impact Explained.* Technical Report. Juniper Research. https://www.juniperresearch.com/document-library/white-papers/roam-like-at-home-impact-explained

[6] 2020. *Case Study: Telefónica's KITE Platform is delivering a wide range of IoT services to GM vehicles in Mexico.* Technical Report. Telefónica. t.ly/mjSe

[7] Carlos E Andrade, Simon D Byers, Vijay Gopalakrishnan, Emir Halepovic, David J Poole, Lien K Tran, and Christopher T Volinsky. 2017. Connected cars in cellular network: a measurement study. In *Proceedings of the 2017 Internet Measurement Conference*. ACM, 235–241.

[8] Andrea Biral, Marco Centenaro, Andrea Zanella, Lorenzo Vangelista, and Michele Zorzi. 2015. The challenges of M2M massive access in wireless cellular networks. *Digital Communications and Networks* 1, 1 (2015), 1 – 19. https://doi.org/10.1016/j.dcan.2015.02.001

[9] Tiago de Andrade, Carlos Astudillo, and Nelson da Fonseca. 2015. Impact of M2M traffic on human-type communication users on the LTE uplink channel.

[10] ETSI. 2013. *Machine-to-Machine Communications (M2M), Functional Architecture.* Technical Report TS 102 690 V2.1.1. ETSI. https://www.etsi.org/deliver/etsi_ts/102600_102699/102690/02.01.01_60/ts_102690v020101p.pdf Work in Progress.

[11] V. Gazis. 2017. A Survey of Standards for Machine-to-Machine and the Internet of Things. *IEEE Communications Surveys Tutorials* 19, 1 (2017), 482–511.

[12] M. Klymash, H. Beshley, O. Panchenko, and M. Beshley. 2017. Method for optimal use of 4G/5G heterogeneous network resourses under M2M/IoT traffic growth conditions. In *2017 International Conference on Information and Telecommunication Technologies and Radio Electronics (UkrMiCo)*. 1–5.

[13] Harini Kolamunna, Ilias Leontiadis, Diego Perino, Sureanga Seneviratne, Kanchana Thilakarathna, and Aruna Seneviratne. 2018. A First Look at SIM-Enabled Wearables in the Wild. In *Proc. of IMC*.

[14] Filippo Malandra, Steven Rochefort, Pascal Potvin, and Brunilde Sansò. 2017. A Case Study for M2M Traffic Characterization in a Smart City Environment. In *Proceedings of the 1st International Conference on Internet of Things and Machine Learning (IML '17)*. Association for Computing Machinery, New York, NY, USA, Article 48, 9 pages. https://doi.org/10.1145/3109761.3109809

[15] Anna Maria Mandalari, Andra Lutu, Ana Custura, Ali, Özgü Alay, Marcelo Bagnulo, Vaibhav Bajpai, Anna Brunstrom, Jörg Ott, Marco Mellia, and Gorry Fairhurst. 2018. Experience: implications of roaming in Europe. In *Proc. of MobiCom*.

[16] Y. Mehmood, N. Haider, M. Imran, A. Timm-Giel, and M. Guizani. 2017. M2M Communications in 5G: State-of-the-Art Architecture, Recent Advances, and Research Challenges. *IEEE Communications Magazine* 55, 9 (2017), 194–201.

[17] Foivos Michelinakis, Hossein Doroud, Abbas Razaghpanah, Andra Lutu, Narseo Vallina-Rodriguez, Phillipa Gill, and Joerg Widmer. 2018. The Cloud that Runs the Mobile Internet: A Measurement Study of Mobile Cloud Services. In *Proc. IEEE INFOCOM*.

[18] Office of Gas and Electricity Markets (Ofgem). 2018. Smart Metering Implementation Programme: Progress Report for 2018. (2018).

[19] R. C. D. Paiva, R. D. Vieira, and M. Saily. 2011. Random Access Capacity Evaluation with Synchronized MTC Users over Wireless Networks. In *2011 IEEE 73rd Vehicular Technology Conference (VTC Spring)*. 1–5.

[20] R. Ratasuk, J. Tan, and A. Ghosh. 2012. Coverage and Capacity Analysis for Machine Type Communications in LTE. In *2012 IEEE 75th Vehicular Technology Conference (VTC Spring)*. 1–5.

[21] WMuhammad Zubair Shafiq, Lusheng Ji, Alex X. Liu, Jeffry Pang, and Jia Wang. 2012. A first look at cellular machine-to-machine traffic: large scale measurement and characterization. *ACM SIGMETRICS Performance Evaluation Review* 40, 1 (2012), 65–76.

[22] WMuhammad Zubair Shafiq, Lusheng Ji, Alex X. Liu, Jeffry Pang, and Jia Wang. 2013. Large-Scale Measurement and Characterization of Cellular Machine-to-Machine Traffic. 21, 6 (2013).

[23] 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*. ACM, 375–387.

[24] Pawan Kumar Verma, Rajesh Verma, Arun Prakash, Ashish Agrawal, Kshirasagar Naik, Rajeev Tripathi, Maazen Alsabaan, Tarek Khalifa, Tamer Abdelkader, and Abdulhakim Abogharaf. 2016. Machine-to-Machine (M2M) Communications. *J. Netw. Comput. Appl.* 66, C (May 2016), 83–105. https://doi.org/10.1016/j.jnca.2016.02.016

[25] Ming Zhao, Arun Kumar, Tapani Ristaniemi, and Peter Han Chong. 2017. Machine-to-Machine Communication and Research Challenges: A Survey. *Wirel. Pers. Commun.* 97, 3 (Dec. 2017), 3569–3585. https://doi.org/10.1007/s11277-017-4686-1