A Shortcut through the IPX: Measuring Latencies in Global Mobile Roaming with Regional Breakouts

Viktoria Vomhoff*, Marleen Sichermann*, Stefan Geißler*, Andra Lutu‡, Martin Giess¶, Tobias Hoßfeld*
*University of Würzburg, Germany
Email: {firstname.lastname}@uni-wuerzburg.de
‡Telefonica, Spain
Email: {firstname.lastname}@telefonica.com
¶emnify GmbH, Germany
Email: {firstname.lastname}@emnify.com

***Abstract*—In the rapidly evolving field of mobile communications, the demand for seamless worldwide connectivity for Internet of Things (IoT) devices introduces distinct challenges. With IoT verticals expanding and international roaming becoming crucial for numerous IoT applications, ensuring high Quality of Service (QoS) has become a critical issue, especially as IoT deployments largely rely on the IPX network for seamless global connectivity. To improve QoS in mobile roaming, we conduct an extensive study, analyzing QoS metrics for over 530,000 roaming IoT devices worldwide. Our novel measurement methodology provides unique insights into individual network segments, including visited and home networks, and the global IPX network. By integrating this data with additional sources, we offer a comprehensive view of the global mobile roaming ecosystem for the first time. Our research highlights the impact of regional breakouts on optimizing QoS for international roaming, demonstrating their significance with the analyzed datasets, which we make available to the research community.**

***Index Terms*—internet of things, mobile network, LTE, 4G, 2G, 3G, international roaming, measurement, dataset**

I. INTRODUCTION

In the era of the Internet of Things (IoT), global and ubiquitous connectivity is more important than ever. With more and more new devices coming online every year, the underlying mobile network architecture, initially designed to facilitate international roaming for mobile subscribers, has evolved into a critical infrastructure for a broad spectrum of verticals. New applications and use cases, particularly in the IoT domain, and the immense increase in Machine to Machine (M2M) communication, puts a considerable load on the ossified roaming infrastructure. As roaming has evolved from a seldom occurring edge case, into the de-facto standard mode of operation [1] for many IoT use cases, the current architecture that built on home-routed roaming increasingly reaches its limits. As IoT devices increasingly penetrate domains, from smart cities and agriculture to healthcare and industrial automation, they generate vast amounts of data. Exemplary use cases range from connecting offshore wind parks and agricultural sites to international asset tracking. Here, it is important to differentiate between Wi-Fi or LAN-based smart home devices, such as Smart TVs or IP Cams and SIM-equipped devices connecting to public mobile networks, where this work focuses on the latter. On a global scale, due to the aforementioned home-routed roaming, many devices are affected by roaming issues, independent of their location, mobility pattern or use case. Specifically, due to the way modern Mobile Virtual Network Operators (MVNOs) operate, 100% of the devices investigated in this work are actively roaming.

Optimizing network deployments for these roaming devices to ensure fast, reliable, efficient and secure communication on a global level becomes increasingly difficult. The depth of this challenge becomes emphasized by the fact that, in order to provide seamless global connectivity, the independent systems of many parties have to successfully interface with each other, to carry network traffic around the globe. The visited network operator provides access to the radio network, and the IP Packet eXchange (IPX) network serves as the pivotal backbone that is orchestrating the seamless flow of data across borders and continents. The home network then operates the breakout into the open internet, with each part along the way being critical to provide global connectivity with a high Quality of Service (QoS).

Despite its significance, due to its complexity, need for built-in security, global scale, and the architectural borders between different parties, the inner workings of the global roaming landscape remains largely opaque. While there are several studies dealing with the performance characteristics of operator networks, there is a critical lack of research on the dynamics within the global IPX network. Recognizing this gap in literature, we pose the research question of how the global mobile roaming architecture can be monitored. To this end, we make three key contributions.

- We introduce a novel method of capturing QoS metrics of individual network segments across the end-to-end transmission path for roaming IoT devices. By performing both active and passive measurements from the point of view of the home network, we dissect various transmission segments on the end-to-end path, including the visited network, the IPX network, and the internet facing breakout. Our results give clear empirical evidence of differences in data usage across different radio technologies and roaming destinations. By dissecting the obtained data, we highlight the importance

978-3-903176-64-5

of regional breakouts in global roaming scenarios.

- We investigate the routing dynamics within the IPX network itself, and distill various data sources into a graph representation of peering points within the IPX. We provide the first visualization of this opaque ecosystem that builds upon the tight interconnection of transit providers on the internet, and make it available to the community.
- We publish the preprocessed and anonymized datasets obtained in the scope of this work. This data encompasses detailed metrics on latency, throughput, and other Key Performance Indicators (KPIs) across various segments of the IPX, home and visited networks, offering unprecedented insights for researchers, network operators, and service providers alike. To the best of our knowledge, this is the first and only dataset of this scope and level of detail.

The remainder of this work is structured as follows. Section II provides a primer on the current international roaming architecture and the components involved therein. Related work on QoS studies in the mobile landscape and investigations of the IPX network is presented in Section III. Section IV presents the methodology applied and datasets gathered in this work. We evaluate and present key insights extracted from the datasets in Section V. A discussion of the results and their relevance, as well as directions for future research, is done in Section VI before Section VII provides concluding remarks.

II. MOBILE ROAMING ARCHITECTURE

As the demand for seamless global connectivity increases across the IoT landscape, the role of the underlying infrastructure becomes paramount in delivering dependable, high-performance connectivity across the globe. Catering to the diverse needs of IoT sectors—from connected vehicles and agriculture to industrial applications and consumer wearables—necessitates deploying devices globally with simplified management of device and customer connectivity. IoT managed connectivity providers such as Twilio/KORE, emnify, and Truphone, often acting as MVNOs, play a vital role in this ecosystem. They manage their own mobile core infrastructure without deploying their own Radio Access Network (RAN), instead relying on roaming agreements to provide global connectivity. Due to this fact, all devices in our datasets are roaming. This approach leverages the mobile networks' roaming capabilities, although roaming was originally conceived for occasional travelers, not the extensive roaming demands of IoT verticals.

In the following, we provide a primer on the components and challenges present in this globally distributed infrastructure. We highlight the individual roles of different stakeholders that form the global roaming architecture and discuss challenges when it comes to assessing the performance of this ecosystem. This architecture is underpinned by a complex network of components, protocols, and agreements that facilitate the continuity of mobile services outside the home network. Hence, we provide a high level overview of the individual segments traffic has to pass through on its transmission path from an IoT device to a server somewhere in a datacenter. Figure 1 schematically shows this transmission path across multiple domains, each operated by a different stakeholder. Note that 3GPP defines the concept of local breakouts [2], in which traffic is directly routed towards the internet by the visited network. However, due to technical, billing, and trust issues [3], most operators do not actively use local breakouts in their daily operation. Instead, the established and more widely used mode of operation is home routed roaming where traffic is routed towards the internet by the home network instead of the visited network.

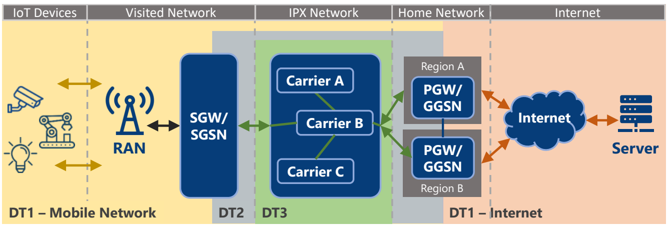

Figure 1: Roaming network infrastructure and data sources.

A. Visited Network

On the left of Figure 1, IoT devices worldwide connect to a local network, typically a Mobile Network Operator (MNO) with its own radio and core infrastructure. For visiting devices to access this infrastructure, two critical elements are needed: a roaming agreement between the local and the device's home operator, allowing service access across networks, and the physical transmission of signaling data for authentication, connection, and mobility management between networks. Roaming agreements are contractual arrangements between operators that allow users of one operator (the home network) to access services on another operator's network (the visited network) when they are outside their home coverage area. After successful connection, data can flow to its destination. Simplifying for clarity, we focus on the Serving GPRS Support Node (SGSN) and Serving Gateway (SGW) in Figure 1, key for data traffic management and routing. These components support connectivity and data transfer across networks, with the SGW tied to Long Term Evolution (LTE) (4G) and the SGSN to older General Packet Radio Service (GPRS) and Universal Mobile Telecommunications System (UMTS) networks (2G/3G). The SGSN/SGW's connection to the home network highlights the second major component in the roaming framework. Note that in the following, we only use SGW to represent the visiting network components.

B. IPX Roaming Exchange

The IPX Roaming Exchange is crucial for mobile roaming architecture, offering a secure, efficient route for exchanging IP-based traffic among global mobile operators and service providers. It supports various services, including MMS, data, VoLTE, and video calls. As a backbone network, the IPX ensures interoperability, QoS, and Service Level Agreements (SLAs) across different technologies, offering a global, private IP network for secure, efficient traffic exchange. Unlike public

Internet eXchange Points (IXPs), the IPX provides a controlled environment with strict standards for service compatibility and reliability. MNOs and MVNOs access this global network through carriers offering peering services, similar to internet peering. International carriers like Syniverse and BICS function as key roaming hubs within the IoT ecosystem, enabling operators to achieve broad roaming capabilities with hundreds of MNOs by connecting to one or more of these hubs.

C. Home Network

Finally, as indicated on the right side of Figure 1, data enters the home network that plays a crucial role in mobile roaming and is responsible for user authentication, routing data into the internet, and billing. Simplifying for clarity, we concentrate on the Gateway GPRS Support Node (GGSN) and Packet Data Network Gateway (PGW), key for managing data traffic among the roaming user, visited network, and the internet. These components, relevant to 2G/3G and LTE networks respectively, handle GPRS Tunneling Protocol (GTP) decapsulation and internet connectivity. The GTP protocol facilitates device-specific tunnels between the visited network's SGSN/SGW and the home network's GGSN/PGW. To offer localized services and reduce latency, some providers deploy multiple GGSN/PGW instances across different regions, enabling traffic rerouting based on device location or manual selection of internet entry points. In this work, we dissect a deployment with two regional breakouts, and discuss the performance implications of such a location-aware roaming setup. Note that in the following, we only use PGW to denote the home network components. Note that, while we focus on 2G/3G and LTE in this work, the developed methodology can be applied directly to 5G standalone deployments. Thereby, user plane traces can be obtained by monitoring at the User Plane Function (UPF) (DT1 equivalent) while the IPX delay or peering delay (DT2 equivalent) can be obtained by measuring RTTs between the home and visited network UPFs, respectively.

III. RELATED WORK

The related work surrounding this study falls into two main groups. First, research on QoS in mobile networks, considering both native users and those in roaming scenarios. Second, work focusing on mobile roaming and the IPX network, including studies on QoS and research addressing various other aspects. Notably, attributed to the IPX and roaming landscape's complexity and the challenges—both political and technical—in accessing necessary platforms and datasets, research in the latter category is scarce.

Caushaj et al. conducted an extensive evaluation of throughput and delay in 3G and 4G mobile architectures for native devices, illustrating the notable improvements in performance metrics as networks evolve from 3G to 4G technologies [4]. When it comes to roaming scenarios, Geissler et al. perform a detailed characterization of the signaling traffic emerging between the visited and home network [5]. Similarly, Vomhoff et al. investigate how the mobile signaling behavior in IoT focused networks changes when outages occur [6]. By evaluating the signaling datasets obtained in real production systems, they establish an underlying behavior that holds true across multiple different types of outages. An analysis of the data usage behavior, including the transmitted volume and duration of data tunnels, has been done by Raffeck et al. [7]. Their results show that devices tend to generate bursty load for the underlying architecture. By being active at specific times of the day, this leads to phases with high peak load, while the general load of the system is considerably lower.

In his survey paper [8], Moriya provides a first overview of the technical elements, services, and challenges of the IPX network. Later, Çakmak et al. compare interconnection models for ISPs and MNOs, considering technical and economic perspectives of the IPX [9]. However, strict security regulations mandated by the Global System for Mobile Communication (GSM) Association (GSMA) prevent any access to the IPX network from the outside [10], hindering researchers from gaining deeper insights. Nevertheless, Lutu et al. [11] present the first structural analysis of the IPX network, by investigating different datasets from an operational IPX provider. As a result, they provide insights into the peering behavior and topology of IPX providers, as well as conducting an analysis of data roaming signaling patterns. Building on this research, they further examine the IPX provider's operations, resulting performance implications, solutions, and investigate the range of devices used by their customer base, along with their emerging communication patterns [12]. As opposed to these works that investigate signaling traffic, we are capturing and evaluating data plane traces and metrics that allow us to assess the performance of individual network segments across the roaming ecosystem. While the publications by Lutu et al. mainly focus on the structure of the IPX network and peering in general as well as signaling patterns in mobile roaming, to the best of our knowledge, our work is the first to incorporate user plane traffic to evaluate the QoS of roaming devices.

To gain a comprehensive understanding of roaming in IoT, it is necessary to examine its various aspects. Therefore, Lutu et al. characterize the global roaming support for an operational M2M platform and analyze the impact of roaming IoT devices on the visited MNO [1]. In their study [13], Ballal et al. compare the performance of roaming IoT devices with those deployed in the home network, while also studying the impact of the underlying cellular IoT technologies.

The performance of MVNOs is of interest to various stakeholders, leading to several studies on this topic. Alcalá-Marín et al. compare three MVNOs with traditional MNOs [14], revealing that MVNOs are gradually transitioning to the regional breakout model to reduce latency. To investigate the differences in performance and behavior among four MNOs and MVNOs, Schmitt et al. analyze their Round Trip Times (RTTs) and route paths while accessing content from popular Content Delivery Networks (CDNs) [15]. To gain a better understanding of the MVNO ecosystem, Xiao et al. conduct a measurement study [16] of a national MVNO that operates on top of an MNO in China. They utilize different datasets to examine the MVNO's architecture, economics, as well as performance and interplay with its base carrier. Focusing on key applications such as web access, video streaming, and voice, Zarinni et al. compare the performance of two MVNO families [17]. In this context, our work is the first to combine real-world data plane traces with routing information and the assessment of multiple regional breakouts to develop a holistic view of the roaming landscape. In addition, the datasets captured and published in the context of this work are among the most detailed publicly available sources of real-world data to date.

IV. METHODOLOGY AND DATASET DESCRIPTION

This section presents the methodology and measurement procedures applied in this study, including the rationale behind the choice of measurement point. We discuss the data sources and datasets collected, referring to Figure 1 which illustrates the roaming architecture, data sets, and network segments involved. Each network segment color represents a part of the transmission path captured in our measurements. The data traces from Figure 1 (DT1, DT2) and additional secondary data sources (DT3, DT4, DT5) are summarized in Table I. DT1, DT2, and DT3 have been sourced in the network of emnify[1], a global MVNO specializing in IoT solutions. Based on this, we state the assumption that the dataset predominantly consists of IoT devices.

A. Measurement Setup

Capturing traffic in mobile networks is a complex challenge, with the choice of a suitable vantage point being crucial for the type of information gathered. In many cases, we must choose between observing network segments individually or the entire end-to-end path as a unified entity, such as through end device measurements.

To overcome these visibility challenges, we chose the PGW in the home network as our measurement point, offering comprehensive insight into the network's distributed segments. Positioned at the network's egress, as shown in Figure 1, the PGW enables observation of data traffic to and from the internet (DT1), interaction with the visited network's SGW, and region-specific data (DT2) through multiple breakout regions. This setup facilitates the measurement of route delays via the IPX network and the extraction of QoS metrics for network segments. Integrating these findings with secondary data sources (DT3, DT4, DT5) provides a comprehensive view of the roaming landscape and its performance characteristics.

In the following, Section IV-B explains the data plane traces obtained in DT1. Section IV-C details the delay measurements of the IPX network contained in DT2. Insights into the routing information within the IPX (DT3) are presented in Section IV-D before Section IV-E outlines the supporting data sources used to augment the primary data sources (DT4, DT5). Lastly, Section IV-F highlights the methodology of integrating these data sources. The datasets are made publicly available via zenodo [18].

Table I: Summary of data sources.

No.	Data	Type	Location	Sec.
DT1	Data Plane Traces	L2-4, MCC/MNC	Device - PGW - Server	IV-B
DT2	GTP Echoes	Response Times	SGW - PGW	IV-C
DT3	Routing Tables	AS Paths	IPX Network	IV-D
DT4	Peering Tables	Peering Points	IPX Network	IV-E
DT5	IR.21	Roaming Database	IPX Network	IV-E

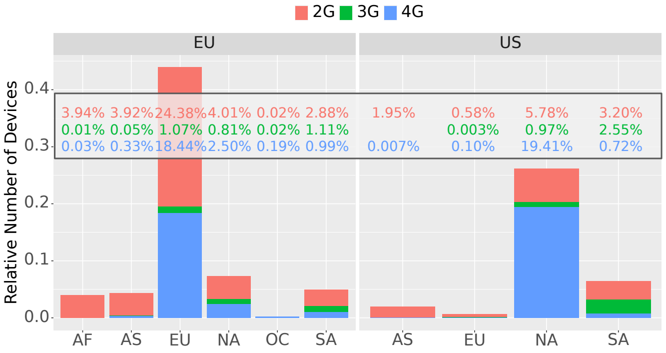

Figure 2: RAT type distribution per continent and breakout region (DT1). AF: Africa; AS: Asia, EU: Europe; NA: North America; OC: Oceania; SA: South America.

B. Data Plane Traces (DT1)

The data plane traces were collected at the PGW of the aforementioned breakout regions of a globally operating MVNO. Specifically, we evaluate traffic captured in two breakout regions located in the EU and US. We capture the raw packet data of both ingress and egress traffic while excluding the payload of packets to ensure user privacy. The data set of the breakout region EU was recorded in June 2023 from 13:40 to 16:10, while the trace in US was collected from 13:40 to 17:00. The dataset contains 510M packets across 24M unique flows generated by 530k unique devices. Note that there are no technical limitations when it comes to measurement duration and size of traces. The measurement durations in this work have been selected to obtain a meaningful amount of data. In a previous study conducted on data obtained from the same operator, we have shown that devices exhibit high periodicity based on hourly patterns [5]. Naturally, more long term studies are required to include day-night as well as workday-weekend cycles. These studies remain for future work, as we focus on establishing the methodology in this work.

Aside from timestamp, flow quintuple, frame length, and protocol information, such as Transmission Control Protocol (TCP) flags and sequence numbers, the data contains information about which visited network in which country a device is connected to. The source and destination IP addresses are salted and hashed with MurmurHash3 to preserve user privacy. IP space structures are not preserved due to privacy reasons. To showcase the scope of the data, we extract the Radio Access Technology (RAT) type distribution across various continents for both observed breakout regions. Figure 2 shows the relative

---

[1] https://www.emnify.com/

number of observed devices using 2G, 3G and 4G connectivity for each continent and breakout region along the y-axis. The facets show the different breakout regions EU and US, the x-axis shows the continent on which the visited network resides. The data shows that for most continents, 2G is still used by the majority of devices when looking at the EU breakout. Only for North America and the US breakout, 4G is used by more devices. This is in line with results from Lutu et al. [1], who find that 60% of roaming IoT devices were still only 2G/3G-capable in 2019. 4G takes the second place for the remaining continents, except for South America in the US breakout, where 3G is still used significantly. The IoT technologies LTE for Machines (LTE-M) and Narrowband IoT (NB-IoT) are also contained in the dataset, but are only used by a negligible number of devices, and are hence omitted here. Finally, while it might be expected that significant changes have occurred over the past five years, potentially reducing the prevalence of 2G/3G devices, it is common for devices to remain in use for extended periods. Those connected to 2G five years ago are likely still utilizing this technology, especially considering 2G's extensive coverage. Additionally, the discontinuation of 3G services may not have a substantial impact, as 3G devices are generally backward compatible with 2G, which remains widely accessible worldwide. Hence, it is expected to observe a significant fraction of devices relying on 2G connectivity, especially in the context of IoT deployments.

C. *GTP Echoes (DT2)*

From the same vantage point, we are exploiting the capabilities of the GTP protocol to obtain RTTs between the home network PGW and the visited network SGW. By sending GTP echo requests from the home network to all visited network currently hosting roaming devices, we are able to measure the isolated RTTs of the IPX network. To obtain the dataset, we extract a list of SGWs that currently hold active data tunnels. By repeatedly sending GTP echo requests, we can then measure the time until the GTP echo response arrives back at the measurement system. The dataset contains the timestamp of the GTP echo message, the RTT, Mobile Country Code (MCC), Mobile Network Code (MNC) for identification of the operator and country, the breakout region to which the measurement belongs and the target SGW IP address. Note that not all visited networks answer to GTP echo requests, and we consider a measurement timeout of 1000 ms. As these are active measurements, we are limiting the impact on the IPX network and the visited network by only sending requests at a limited rate. The specific rate depends on the current number of SGWs that hold active data tunnels, as we iterate over all targets before starting over. Overall, we took 890k measurements over 7 days across both breakout regions. In that time, we measured RTTs to 224 operators and 442 SGWs. The observations are distributed approximately equally among the breakout regions. The continents with the most measurements are, as in the data plane traces, Europe and North America. The measurement was taken in August 2023 at the same time for both breakout regions.

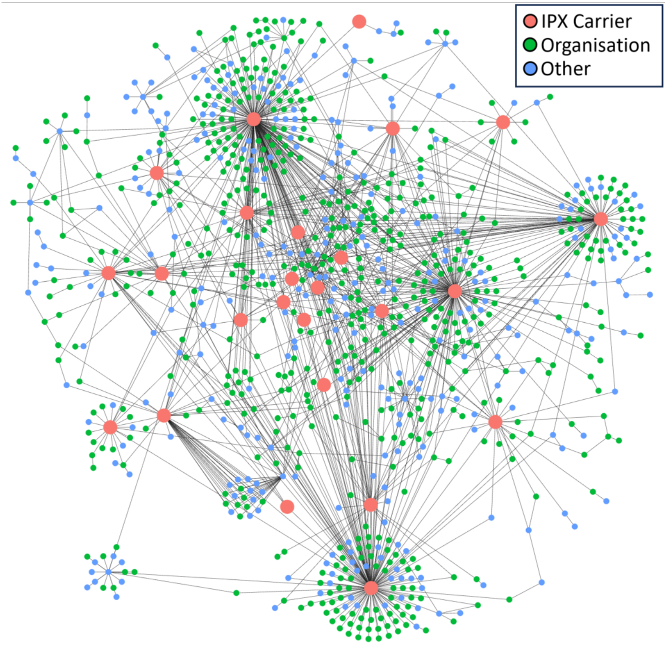

Figure 3: Representation of the global IPX network (DT3).

D. *IPX Routing Information (DT3)*

The first supporting data source is the routing information for the IPX network in the form of Border Gateway Protocol (BGP) routes. In total, we had four routing tables coming from four routers, two in Europe and two in the US, connecting the breakout regions to two IPX carriers. For each router, we see 27 000 routes with approximately 13 740 prefix destinations, since we usually have more than one route to a destination. A prefix in this context is an IP address range advertised by a router to its neighbors, specifying the destinations it can reach.

Without diving into the complexities of the BGP protocol, the most important information for this work are the Autonomous System (AS) path without prepending, the routing destination, and the prefix length. Similar to how it is done in the internet, prepending is used by Service Providers (SPs) and IPX carriers to manipulate the preference of a route. Thereby the AS path is artificially extended by appending certain AS Numbers (ASes) several times in sequence [11], making the route appear to have more hops. During data cleaning, we removed all prepending from the routes. The routing destination contains the network address of the prefix and the prefix length the associated subnet mask for the prefix. Based on this information, we can extract the path taken through the IPX network for each of the visited networks observed in both the passive dataplane measurements (DT1) and active GTP echo measurements (DT2).

Figure 3 illustrates the IPX network as a graph based on the obtained routing data. Vertices symbolize ASes with edges indicating BGP peerings. Red vertices are IPX carriers, green represent organizations (MNOs, MVNOs) as identified using the GSMA IR.21 database (see next section on supporting data sources), and blue vertices signify ASes either unlisted in

IR.21 or associated with multiple MNOs. This representation captures the IPX network from a singular operator perspective. In total, we identified 25 IPX carriers (red): 4 peer with over 100 entities with one carrier having 227 peerings, 12 carriers link between ten and 50 ASes, and 9 connect to fewer than ten. Additionally, we identified 686 organizations (MNOs and MVNOs) (green) and 240 Other (blue). These observations also reflect in the graph metrics of the IPX topology. A modularity of 0.67 and clustering coefficient of 0.11 indicate strong communal structures with a low degree of interconnectivity within communities. This is in line with communities forming around IPX carriers, where, from our perspective, members only peer with the carrier, but not among each other. Note that peerings among members may exist, but are not visible from our point of view. A density of only 0.002 and average shortest path length of 3.5 further indicates that many nodes only have few connections, but can reach other nodes via only few hops, emphasizing the strong position of IPX carriers and roaming hubs in the IPX.

E. Supporting Data Sources (DT4, DT5)

In addition to the datasets introduced above, we refer to two supporting data sources that are used to augment the obtained data by additional information. First, we assembled a list of IPX peering points (DT4). These records of which IPX carrier peers with other carriers at which geographical location allows us to better understand the peering behavior within the IPX and fill in the missing gaps in our graph representation. Peering between IPX carriers usually occurs at one or more IXPs [11], regional to the continent [19]. In the peering tables, however, this is simplified, and it is not specified at which IXP the carriers peer with each other, but whether they peer on a continent. This is done to, on one hand, reduce complexity, on the other hand, to work with flawed or incomplete data. Most carriers provide information of peering on a continental level, but not on IXP level.

Finally, we refer to the GSMA IR.21 roaming database (DT5) that contains roaming configuration information. The IR.21 document is a technical document maintained by mobile operators. It facilitates international roaming by providing the information necessary for operators to enable and manage roaming agreements and interoperability between different mobile networks. This document includes a wide range of technical specifications, such as the ASNs used by an operator, the corresponding IP prefixes, network configurations and service capabilities. In this work, we use this data source to extract a mapping between MCC, MNC and the ASes belonging to this respective operator. This allows us to identify the organization to which an autonomous system belongs in the graph representation of the IPX network.

F. Combining Data Sources

Combining these data sources allows us to map countries, MNOs, or SGW IPs to their respective latencies. This enables us to associate measurements with specific routes and both geographical and logical locations within the global IPX network. This, in turn, provides us with valuable BGP routing insights, including details like the AS path and the geographical trajectory of the measured route. The data sources are combined as presented in Figure 4. As a first step, DT3, DT4, and DT5 are combined based on their ASN and AS path. As the routing tables (DT3) only contain the ASNs and no description, we use DT4 to add information of the IPX carriers and DT5 to augment geographical data (country, continent) and the respective organization (MNO, MVNO, carrier). This newly created dataset (DT3*) is the base for the graph representation shown in Figure 3.

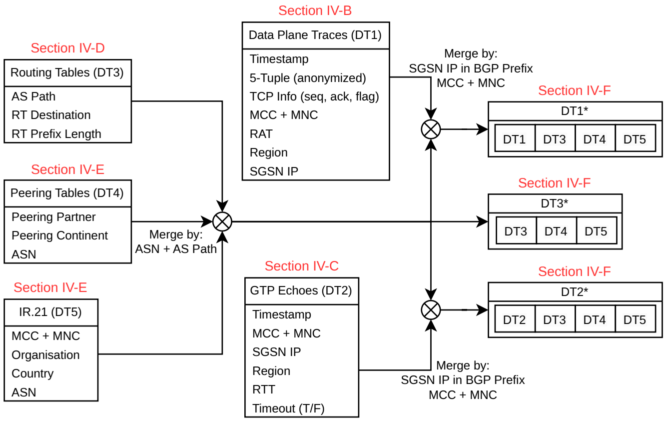

Figure 4: Processing pipeline for dataset augmentation.

In a second step, we extend both the dataplane and GTP echo measurements (DT1, DT2) by mapping the included SGW IP to a specific autonomous system based on DT3*. We label these augmented datasets DT1* and DT2*. Consequently, we can identify the exact trajectory of observations in both DT1 and DT2 by matching the SGW IP address from each measurement with those specified in DT3*. In cases where multiple routes are possible, we choose the one selected by BGP, which by default is the route with the longest prefix. Given that a single prefix can be served by more than one IPX carrier, we select the active route with higher priority.

V. EVALUATION

This section details the results of our measurement study, emphasizing the delay across network segments on the path from IoT devices to targets on the internet and observed throughput. We compare 2G and 4G usage patterns, noting that while the published datasets enable more fine-grained investigations, we omit a deep dive due to space limitations.

A. Throughput and Data Usage

We start by evaluating data usage across the two breakout regions US and EU for different mobile generations. Figure 5 shows the normalized throughput, meaning the total observed throughput divided by the number of unique devices for each of the breakouts. The normalization allows the direct comparison of the data points across the two regions, despite the presence of a differing number of devices. We make

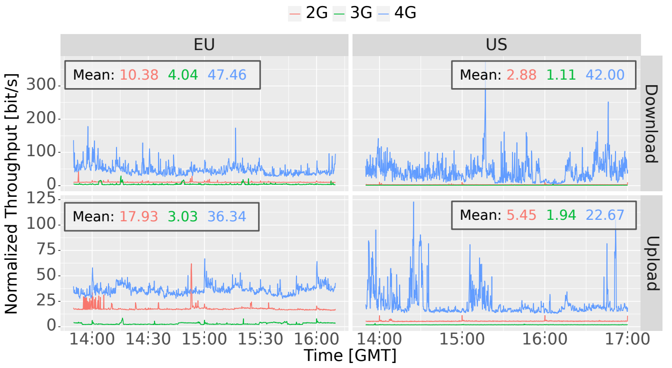

Figure 5: Dataplane throughput over time for different RATs and regional breakouts (DT1).

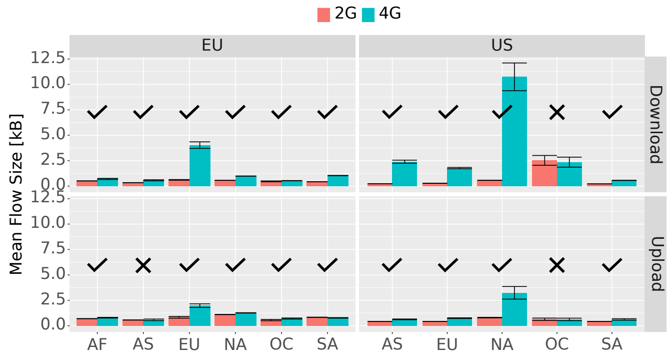

Figure 6: Mean flow size for different regional breakouts, continents and RATs (DT1). AF: Africa; AS: Asia, EU: Europe; NA: North America; OC: Oceania; SA: South America.

three key observations: (i) With 70% in the EU breakout and 85% in the US breakout, 4G is responsible for a significant fraction of the generated traffic. (ii) Despite its age, 2G is still contributing a significant amount of traffic, with 24% in the EU and 11% in the US. (iii) The traffic attributed to 3G is, despite its fade out, still contributing 6% and 4% for EU and US, respectively. The data shows further that 4G throughput fluctuates more over time, indicating its use for more bandwidth intensive applications, as opposed to the low bandwidth, continuous traffic observed in 2G and 3G. In addition, fluctuations are larger for the US breakout than the EU breakout. The coefficients of variation for EU are 0.60, 0.37, 0.62 for 2G, 3G and 4G, respectively. Analogously, the values for the US breakout are 0.16, 0.12, and 0.48. Finally, 2G is consistently used to upload more traffic than download (EU: 15% up, 9% down; US: 7% up, 4% down). For 4G, the measurements show that more data is being downloaded than uploaded (EU: 30% up, 40% down; US: 30% up, 55% down). It is rather unexpected that, contrary to common assumption, IoT devices upload less data than they download. The specific reasons for this phenomenon cannot be determined without detailed knowledge of the devices involved. However, the data shows that a significant portion, approximately 65%, of data downloaded over 4G networks occurs over HTTPS (port 443), while on 2G networks, HTTPS downloads only account for about 10%. A plausible explanation could be attributed to the spread of 4G technology in more modern devices, which typically receive more frequent firmware updates and are suitable for more modern, resource intensive use cases that often require significant downstream bandwidth, such as infotainment systems in cars. However, it is important to note that identifying the specific reason for this behavior would require detailed knowledge of the individual devices and the verticals they are deployed in, which is currently not available.

Based on the observations above, we omit 3G for the remainder of the paper, as its fade out is only a matter of time, and existing devices will likely switch over to 2G. Note that the dataset also contains data for LTE-M and NB-IoT, both of which are omitted here, as their contribution to the total throughput is negligible.

Digging deeper into the usage of data for 2G and 4G, Figure 6 shows the mean flow sizes in kB in both uplink and downlink direction for both breakout regions and all continents. The error bars show the 95% confidence interval. The data confirms that, for most continents, there is a significant difference in flow size between 4G and 2G traffic, even when neglecting parameters such as the flow duration. The markers for each bar group indicate whether the data supports a statistically significant difference between 2G and 4G (✓) or not (✗).

B. Roaming Packet Latencies

Moving on to the observed latencies, we separate the data by breakouts, continents, and RAT type. Figure 7 displays a time series of average radio network RTTs for roaming devices. It presents mean delays over 10-minute intervals with a 95% confidence interval which are computed using the student-t distribution and 1-minute aggregate data to highlight variance. The findings reveal lower delays in 4G compared to 2G across regions. In 4G, as is expected, EU devices experience less delay using the EU breakout, and devices in the USA benefit from the US breakout. This pattern is less evident in other continents and absent in 2G data, as the radio delay inherent to 2G dominates regional breakout advantages.

To calculate the RTTs shown in Figure 7, we exploit the three-way-handshake performed by each new TCP connection. From the perspective of the PGW, a connection from a roaming device to a server on the internet consists of an upstream SYN, downstream SYN-ACK and upstream ACK. Using this transaction, we can compute both the upstream and downstream RTTs between the device and the PGW and the RTT between the PGW and the server on the internet. Note that the data in Figure 7 shows the mobile network part between device and PGW. Figure 8 displays an ECDF of the latencies by continent and region, affirming the noticeable differences in delays between 2G and 4G, with regional breakouts significantly impacting 4G latencies. Additionally, the data reveals a multi-modal distribution, especially in Asia

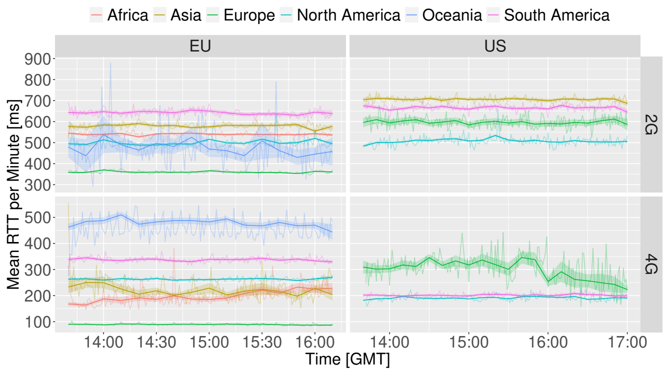

Figure 7: Mean observed RTT between device and home network over time across continents, regional breakouts and RATs (DT1).

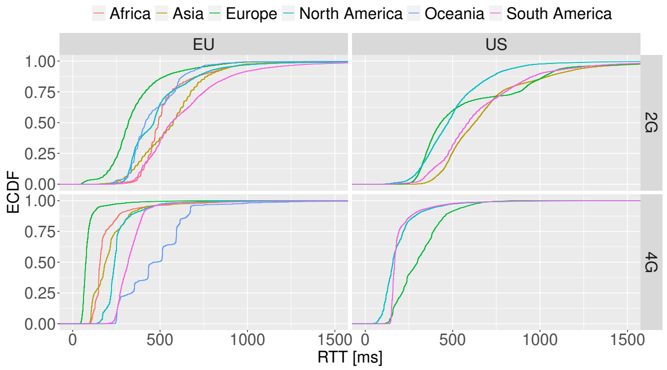

Figure 8: ECDF of RTTs between device and home network across continents, regional breakouts and RATs (DT1).

and Africa's 2G measurements, attributed to aggregating data across various operators with slightly different delay scales. However, this is not the case for the pronounced multi-modal behavior in Oceania and the EU breakout for 4G that is instead attributed to a single operator, suggesting it may result from that operator's specific technical setup.

Similarly, based on the RTT between PGW and the target on the internet, we can assess the impact of breakout regions on the delay experienced once the mobile network has been traversed. To highlight this, we show an exemplary scenario in Table II. We show the mean RTT and the 95% confidence interval for the internet facing network segment and a server located in North America. First, as expected, the used mobile technology has no impact on the internet facing delay. Second, the data shows a significant difference between the delays experienced via each of the breakouts. In a follow-up study, we plan to assess the impact of selecting a breakout close to the device over selecting a breakout close to the target service. The assumption here is that it is beneficial to select a breakout as close to the device as possible, to minimize delay. However, more in-depth analyses are required to verify this assumption.

Table II: Exemplary RTT between home network and internet (mean and 95% CI) across regional breakouts and RATs.

	EU	**US**
2G	87.32 ± 0.13	21.40 ± 0.02
4G	87.62 ± 0.19	21.52 ± 0.23

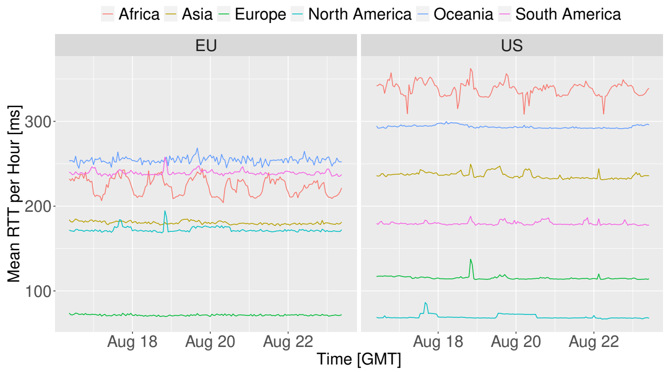

Figure 9: Mean RTT between visited and home network over time across continents and regional breakouts (DT2).

Finally, Figure 9 shows the mean IPX RTT per hour in ms over several days (DT2). As is expected, Europe in EU and North America in US have similar RTTs. The continents Africa, Asia, Europe, and Oceania have a lower RTT in EU than in the US breakout region. North America and South America have a lower RTT in the US breakout region. While one might assume that Asia benefits from low latency via the US breakout based on its geographical proximity, a closer examination of major carriers reveals a different picture. For instance, BICS has established low-latency routes connecting Asia, Africa, and Europe, which can explain the lower RTT [20]. Furthermore, the high variance observed in African visited networks correlates with the local day-night pattern, with higher RTT over the day, and lower RTT during the nights, indicating high load on either the targeted SGSNs, or the involved peering points. These fluctuations in the mean RTT indicate variations in network performance influenced by factors like network congestion and infrastructure quality. Analyzing these trends helps identify areas for network optimization and informs decisions regarding regional breakout selection.

VI. DISCUSSION

In the following, we discuss both the datasets gathered in the context of this work and the results obtained through their evaluation. We cover key insights, limitations as well as directions for future work.

Dataset Limitation and Regional Bias. The datasets obtained for this work exhibit a strong regional bias towards the US and EU (cf. Figure 2). This regional skew likely underrepresents the performance and delay patterns of devices using international roaming in other parts of the world, such as Asia, Africa, and South America, where network infrastructure and roaming agreements might differ substantially. However, even for those continents, the number of operators, devices, and samples is sufficient to obtain statistically significant results. We argue that further research is required in order to identify systemic differences in both device behavior and network QoS between geographic regions. Additionally, the insights derived from our datasets may not accurately reflect a generalizable global roaming experience, as we are investigating the global ecosystem from the point of view of a single operator and two regional breakouts. We argue that additional research is required to investigate the underlying impact of regional breakouts on the QoS of international roaming. Finally, the brief temporal snapshot provided by the datasets investigated in this work is likely insufficient to capture variability over time, such as peak vs. off-peak usage patterns, potentially overlooking factors that could significantly affect QoS. We argue that larger datasets over longer periods of time are required to faithfully capture these factors.

Key Insights and Practical Significance. We summarize the key insights gained by analyzing the latencies of individual network segments within mobile networks during international roaming for both operators and users. First and foremost, the real-world data shows that a significant amount of IoT devices are still relying on 2G/3G technology. In fact, with 58%, more devices use 2G/3G than LTE. However, only 24% of the total traffic volume is generated by those devices, while 75% of traffic is generated by 4G devices. The remaining 1% are attributed to NB-IoT and LTE-M. It is important to keep in mind that these numbers are specific to the MVNO investigated in this work. Given these findings on the prevalence of different access technologies for IoT devices, it is important to discuss the considerations behind the choice of access technology, with factors such as data requirements, resource constraints, and functional capabilities influencing the decision. While newer technologies like 4G and 5G provide faster data speeds and lower latency, older technologies such as 2G remain relevant for specific use cases. For resource-constrained devices that do not require high data transmission rates, 2G may be sufficient. Additionally, some IoT applications may be designed to operate optimally on 2G networks due to factors like available coverage or cost-effectiveness. However, the choice of access technology should be aligned with the specific requirements of the IoT application. For instance, applications that involve real-time data processing or require high bandwidth may necessitate the use of newer technologies like 4G or 5G. Ultimately, a comprehensive understanding of the use case and its requirements is essential to determine the most suitable access technology for IoT deployments. Addressing these considerations is part of our future work.

Regarding the structure of the IPX network, our data shows that the four largest carriers provide significant peering capabilities among global operators. In addition, most MNO-to-MNO paths are shorter than four hops, showing the dense peering fabric within the IPX. These findings regarding the IPX structure and usage of radio technologies are in line with the results obtained by Lutu et al. [1], [11] in 2020 from the point of view of another global operator. We argue that the validation of these previous results is an essential step towards a generalizable understanding of QoS during international roaming. However, further research is required to extend both the gathered datasets and developed methodology.

Finally, the obtained delay values highlight the importance of regional breakouts in international roaming, as both the delay between a device and the home network PGW, as well as the internet-facing delay after exiting the mobile network show significant differences between regions. Switching to the correct breakout region within any visited network can reduce delay by up to 74%. This value is obtained by identifying a public server on the internet based on its anonymized IP Address and searching for operators with two devices connecting to this exact server using the different investigated breakout regions. We then compare the latency difference between the two connections. Doing this for all available operators, the maximum observed improvement is 74%. We argue these insights are a critical first step towards understanding the impact of regional breakouts, but additional research, specifically broader measurement studies including active measurements involving specific devices, is required to establish a general understanding of the system.

VII. CONCLUSION

In this work, we present a comprehensive analysis of active and passive measurements of latencies present in individual network segments within the mobile network during international roaming. We outline both the developed methodology to obtain the required key metrics and provide insights into the expected QoS for devices roaming across the globe. By examining over 530,000 devices across a multitude of countries, continents, and operators, this study sheds light on the complexities and performance bottlenecks that IoT devices face in a roaming context. The specific focus on the delay values observed between the end device and the home network as well as the visited network and the home network highlights the potential for optimization in current mobile deployments. The findings, especially the comparative delay values from regional breakouts in the EU and the US, have significant practical applications. They offer network operators critical data to enhance the efficiency and robustness of mobile networks by pinpointing specific segments that may require infrastructure improvements or optimized routing protocols. By augmenting this information with routing data across the IPX network, we identify additional room for optimization on the operator end. For customers, particularly those relying on IoT solutions for critical operations, the study underscores the importance of selecting network partners that demonstrate superior performance in international roaming scenarios. Ultimately, these results lay the foundation for the development of more resilient and performant mobile networks, ensuring seamless global connectivity, reliability, and efficiency in an increasingly interconnected ecosystem. Finally, we invite the networking community to contribute their own investigations in the area and make the datasets gathered in the context of this work publicly available.

ACKNOWLEDGMENT

ORIGAMI project has received funding from the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation program under Grant Agreement No. 101139270.

REFERENCES

[1] A. Lutu, B. Jun, A. Finamore, F. E. Bustamante, and D. Perino, "Where things roam: Uncovering cellular iot/m2m connectivity," in *Proceedings of the ACM Internet Measurement Conference*, 2020, pp. 147–161.

[2] 3GPP, "Lte; service requirements for the evolved packet system (eps) (3gpp ts 22.278 version 15.4.0 release 15)," 3GPP, Tech. Rep. ETSI TS 122 278 V15.4.0 (2018-10), 2018.

[3] L. E. Chatzieleftheriou, M. Gramaglia, A. Garcia-Saavedra, S. Gebert, G. García-Avilés, S. Geissler, M. Fiore, P. Patras, A. Lutu, D. Tsolkas *et al.*, "Towards 6g: Architectural innovations and challenges in the origami framework," in *European Conference on Networks and Communications & 6G Summit*, 2024, pp. 1–6.

[4] E. Caushaj, I. Ivanov, H. Fu, I. Sethi, and Y. Zhu, "Evaluating throughput and delay in 3g and 4g mobile architectures," *Journal of Computer and Communications*, vol. 2, no. 10, p. 1, 2014.

[5] S. Geissler, F. Wamser, W. Bauer, M. Krolikowski, S. Gebert, and T. Hoßfeld, "Signaling traffic in internet-of-things mobile networks," in *2021 IFIP/IEEE International Symposium on Integrated Network Management (IM)*. IEEE, 2021, pp. 452–458.

[6] V. Vomhoff, S. Geissler, F. Loh, W. Bauer, and T. Hossfeld, "Characterizing mobile signaling anomalies in the internet-of-things," in *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*. IEEE, 2022, pp. 1–6.

[7] S. Raffeck, S. Geissler, M. Krolikowski, S. Gebert, and T. Hoßfeld, "Data usage in iot: A characterization of gtp tunnels in m2m mobile networks," in *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*. IEEE, 2022, pp. 1–6.

[8] T. Moriya, "Survey of ipx (ip exchange) as an emerging international interconnection between telecommunication networks," *IEICE transactions on Communications*, vol. 96, no. 4, pp. 927–938, 2013.

[9] G. Çakmak and H. Suomi, "A comparison of isp and mno interconnection models," in *2014 21st International Conference on Telecommunications (ICT)*. IEEE, 2014, pp. 431–436.

[10] GSMA, *Guidelines for IPX Provider networks (Previously Inter-Service Provider IP Backbone Guidelines) IR.34 Version 17.0*, GSM Association, 2021.

[11] A. Lutu, B. Jun, F. E. Bustamante, D. Perino, M. Bagnulo, and C. G. Bontje, "A first look at the ip exchange ecosystem," *ACM SIGCOMM Computer Communication Review*, vol. 50, no. 4, pp. 25–34, 2020.

[12] A. Lutu, D. Perino, M. Bagnulo, and F. E. Bustamante, "Insights from operating an ip exchange provider," in *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, 2021, pp. 718–730.

[13] K. D. Ballal, R. Singh, L. Dittmann, and S. Ruepp, "Experimental evaluation of roaming performance of cellular iot networks," in *2022 Thirteenth International Conference on Ubiquitous and Future Networks (ICUFN)*. IEEE, 2022, pp. 386–391.

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

[15] P. Schmitt, M. Vigil, and E. Belding, "A study of mvno data paths and performance," in *Passive and Active Measurement: 17th International Conference, PAM 2016, Heraklion, Greece, March 31-April 1, 2016. Proceedings 17*. Springer, 2016, pp. 83–94.

[16] A. Xiao, Y. Liu, Y. Li, F. Qian, Z. Li, S. Bai, Y. Liu, T. Xu, and X. Xin, "An in-depth study of commercial mvno: Measurement and optimization," in *Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services*, 2019, pp. 457–468.

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

[18] V. Vomhoff, M. Sichermann, S. Geissler, A. Lutu, M. Giess, and T. Hossfled, "Measurement data: Latencies and traffic traces in global mobile roaming with regional breakouts," 2024. [Online]. Available: https://doi.org/10.5281/zenodo.11065734

[19] B. Augustin, B. Krishnamurthy, and W. Willinger, "IXPs: Mapped?" in *Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement*, 2009, pp. 336–349.

[20] bics, "Scale capacity at will and enter new global markets," 2021. [Online]. Available: https://www.bics.com/wp-content/uploads/2021/07/BICS-Capacity-brochure.pdf

APPENDIX

This research has been conducted with the highest regard for ethical standards, prioritizing the privacy and confidentiality of the user data gathered in the context of this work. The collected data points have been conscientiously cleaned and anonymized before exporting the data for the research conducted in this work. No personal identifiers can be linked to individuals, and no contact information was available or used at any stage of the research. Hence, this work does not raise any ethical issues.