

Internet Usage and Performance in GEO Satellite Networks: A Large-Scale Study Across Europe and Africa

Original

Internet Usage and Performance in GEO Satellite Networks: A Large-Scale Study Across Europe and Africa / Merlach, G., Perdices, D., Perna, G., Trevisan, M., Giordano, D., Mellia, M.. - In: COMPUTER NETWORKS. - ISSN 1389-1286. - ELETTRONICO. - 282:(2026). [10.1016/j.comnet.2026.112244]

Availability:

This version is available at: 11583/3009628 since: 2026-04-05T17:33:18Z

Publisher:

Elsevier

Published


DOI:10.1016/j.comnet.2026.112244

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)



Internet usage and performance in GEO satellite networks: A large-scale study across Europe and Africa

Gabriele Merlach ^{id a,*}, Daniel Perdices ^{id b}, Gianluca Perna ^c, Martino Trevisan ^{id a}, Danilo Giordano ^{id c}, Marco Mellia ^{id c}

^a University of Trieste, Italy

^b Universidad Autónoma de Madrid, Spain

^c Politecnico di Torino, Italy

ARTICLE INFO

Keywords:

Satellite communications

Performance

Passive measurements

ABSTRACT

Satellite Communication (SatCom) offers internet connectivity where traditional infrastructures are too expensive to deploy. When using satellites in a geostationary orbit, the distance from Earth forces a round-trip time of at least 550 ms. Coupled with the constrained capacity of the physical link, this challenges the traditional internet access quality we are used to.

In this paper, we present a complete passive characterization of the traffic carried by an operational SatCom provider. With this unique vantage point, we observe the performance of the SatCom technology, as well as the usage habits of subscribers in different countries in Europe and Africa. We highlight the implications of such technology on Internet usage and functioning, and we pinpoint technical challenges due to the CDN and DNS resolution issues, while discussing possible optimizations that the ISP could implement to improve the service offered to SatCom subscribers. We complete the characterization of the adoption and performance of newer protocols with a focus on IPv6 and QUIC.

1. Introduction

While 5G and Fiber-To-The-Home (FTTH) technologies give us access capacity on the order of Gb/s [1–4] and Content Delivery Networks (CDN) can guarantee end-to-end delay of less than a millisecond [5–7], there are significant parts of the world where economic and technological constraints force people to rely on solutions that provide far more constrained access to the Internet. These include mountainous and rural areas in developed countries, as well as the entire territory of underdeveloped countries, where even the supply of stable electricity can be problematic. In such scenarios, Satellite Communications (SatCom) offers a practical connectivity solution. Among the available SatCom technologies, geostationary (GEO) satellites are the oldest and most widely used solution [8], with the first offerings dating back to the early 2000. Here, a satellite orbits the Earth at an altitude of about 36000 km, and moves at an angular velocity equal to the Earth's rotational speed. To an observer on Earth, the satellite appears immobile, making it easier to establish a communication link. A single GEO satellite can cover entire continents, and the directional beams enable efficient space and fre-

quency multiplexing, with each beam providing a total channel capacity on the order of 10 Gb/s [9].

In addition to limited shared capacity, GEO satellite communications suffer from high propagation latency, which is about 550 ms for a round trip (including two passes through the satellite). For this reason, complicated Medium Access Control (MAC) and scheduling protocols coordinate the access and sharing of the satellite's uplink and downlink, while traffic shapers, Performance Enhancing Proxies (PEP), and TCP optimization solutions attempt to mitigate the effects of end-to-end delay and limited capacity [10–12].

In this paper, we have the unique opportunity to present the first large-scale passive characterization of a global GEO SatCom Internet access solution. Through passive instrumentation of the satellite ground station, we observe a significant amount of traffic, 634 TB of network packets and 5.6 billion flows, from tens of thousands of customers in more than 20 countries in Europe and Africa, covering the period from June 2023 to July 2023. On the one hand, this allows us to characterize the different Internet usage habits in the different scenarios, if any. On the other hand, we observe the impact of

* Corresponding author.

E-mail addresses: gabriele.merlach@dia.units.it (G. Merlach), daniel.perdices@uam.es (D. Perdices), gianluca.perna@polito.it (G. Perna), martino.trevisan@dia.units.it (M. Trevisan), daniilo.giordano@polito.it (D. Giordano), marco.mellia@polito.it (M. Mellia).

<https://doi.org/10.1016/j.comnet.2026.112244>

Received 17 June 2025; Received in revised form 16 March 2026; Accepted 20 March 2026

Available online 28 March 2026

1389-1286/© 2026 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

SatCom technology on performance and identify possible optimization strategies.

The main observations can be summarized as follows:

- In Africa, chat and social media applications consume 100 and 10 times more data than in Europe. This is due to the presence of community WiFi points that share SatCom access, with dozens or even hundreds of users sharing a single subscription in community-based hotspots.
- Since these applications are accessed throughout the day, the typical peak time in African countries is anticipated in the morning.
- Satellite channel protocols and solutions increase the Round Trip Time (RTT) much more than just the propagation delay. Link channel quality and congestion (if any) can actually add seconds to the end-to-end RTT. It has implications on the QUIC protocol, an encrypted connection-oriented transport layer protocol over UDP. Since PEPs are ineffective due to the encrypted QUIC header, QUIC flows struggle to reach high throughput compared to the PEP-accelerated TCP traffic.
- In SatCom networks, all traffic must pass through the same ground station, in our case in Europe. This impacts local popular services. For example, services in Africa suffer from the additional delay caused by traffic being routed back and forth through the ground station.
- To complicate the picture, most customers use open DNS resolvers, some of which are located in China and Africa. This increases DNS response time to hundreds of milliseconds and jeopardizes the server selection policies of CDNs and DNS resolvers.
- We observe that the traffic to major content providers relies on the newest protocols, with the majority of traffic being QUIC and IPv6. However, in Africa, the transition proceeds at a slower pace compared to Europe.

We believe that the characterization offered in this paper contributes to understanding the complexity of the Internet by providing a novel perspective on SatCom networks and customers. We also discuss possible technical improvements that SatCom providers might consider to improve the quality of service offered to their customers. These include the use of additional ground stations to route traffic more efficiently and the control of DNS requests or responses to limit the impact of incorrect server selection. To allow the research community to conduct experiments with an emulated GEO SatCom connection and compare it to other connection technologies (including the novel Starlink connection with data from [13]), we have created a data-driven model for our ERRANT network emulator tool [14] and make it available at <https://github.com/SmartData-Polito/errant>.

This paper extends our initial work [15] in several directions. First, we characterize IPv6 adoption (Section 4.2) and leverage it to study user habits in terms of subscription sharing (Sections 5.3 and 5.4). Second, we focus on the QUIC protocol (Section 6.3), studying its performance (Section 7.5) in a SatCom environment where a PEP is in place.

In the remainder of the paper, we introduce the SatCom technology and monitoring infrastructure (Section 3) and provide an overview of our dataset (Section 4). We then illustrate our results in terms of volumetric traffic distribution (Section 5.2), user habits and service consumption (Section 6), and performance indicators (Section 7). Finally, we discuss related work (Section 2) and draw conclusions (Section 8).

2. Related work

Since the Internet was born, characterization of network traffic has been an important topic to understand its evolution and usage. Seminal works analyze trends of Internet traffic trends in the 1990s and early 2000s [16,17]. Since then, there has been a large body of research using passive measurements to understand Internet operations and usage patterns. Most works focus on characterizing Internet traffic generated

by cable/fiber customers [1,5,18,19], while others analyze traffic captured in the backbone network, mixing users with different access technologies [20,21]. Recently, researchers have focused on analyzing the impact of the Covid-19 pandemic on Internet traffic, noting sudden and remarkable changes [22–24]. Our work is in this area and uses a similar methodology for collecting and analyzing passive measurements. We note that most measurements are related to Europe and North America, while we also provide insights into African Internet traffic, similar to Johnson et al. [25]. In addition, in this work, we have a unique opportunity to specifically study traffic from SatCom subscribers. To the best of our knowledge, we are the first to conduct a longitudinal study of SatCom traffic and present a comparison of traffic from different countries on two continents.

atellite-based consumer Internet access [8] was launched in 2003 and has been the subject of a wide corpus of literature. Most of the work targets the design and optimization of communication channels, and a few studies successfully cover the topic [9,26,27]. Other work focuses on measuring and improving the performance of Internet protocols in SatCom. Tropea et al. [28] evaluates different TCP versions on geostationary satellite links, while Peng et al. [29] and Muhammad et al. [30] focus on the interaction of TCP and PEPs. QUIC is an encrypted, connection-oriented transport protocol running over UDP, standardized by IETF [31], whose reduced handshake latency and encryption of transport headers may significantly alter its interaction with SatCom-specific optimization mechanisms. More recently, the performance of QUIC and its interaction with PEPs have been studied [32–34]. For a complete benchmark of SatCom performance, see Deutschmann et al. [35], which provide results on SatCom latency, throughput, and web page load time.

Recently, Michel et al. [13] provided a first characterization of Starlink, a low Earth orbit satellite system, analyzing its impact on user-perceived performance when accessing globally distributed resources and the behavior of different HTTP versions. Complementing this, Kassem et al. [36] examined Starlink connectivity from the browser perspective, highlighting performance variations across different vantage points. Unlike geostationary satellites, which dominate current SatCom deployments at altitudes around 36000 km, LEO satellites orbit much closer to Earth, resulting in substantially lower propagation delays. This fundamental difference affects transport protocol behavior, the effectiveness of optimization mechanisms such as PEPs, and overall user-perceived quality of service, explaining why results from LEO systems cannot be directly compared to GEO-based measurements. All of these studies rely on active measurements, deploying test environments with SatCom devices to evaluate performance. In this work, we provide a different perspective by providing performance data on SatCom through passive measurements by observing traffic from about 10 k customers and complementing the results of previous work.

3. Measurement setup and methodology

In this section, we describe our measurement setup and the methodology we use to gather and analyze data from the actual deployment of a large international SatCom operator. We first provide an overview of the specific data-link technologies that can affect Internet access performance, but avoid detailing the complexity of the physical layer of SatCom transmissions.

3.1. The SatCom network

As in traditional SatCom networks, the operator has deployed a satellite infrastructure consisting of satellites in geostationary orbit, and a ground infrastructure. Referring to Fig. 1, subscribers employ a dedicated equipment, i.e., the Customer-Premises Equipment (CPE), to connect their devices (PC, smartphone, etc.) to the SatCom network. The CPE consists of a dish antenna and a router/modem that manages the satellite links and access protocols on the one side, while offering WiFi

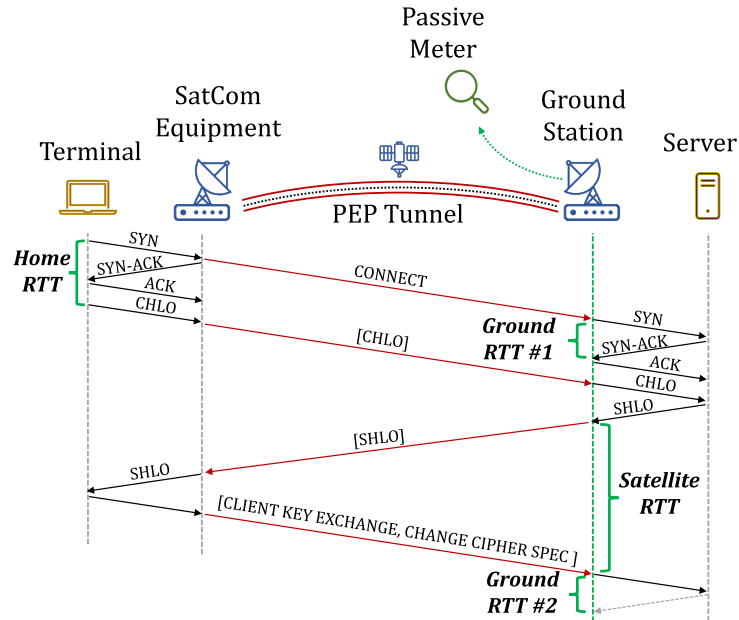


Fig. 1. Methodology for the estimation of the satellite segment RTT.

and Ethernet connectivity on the other side.¹ The satellite acts as relay for subscriber traffic, which traverses 35 786 km twice to reach the ground station, accumulating from 240 ms to 280 ms, depending on the location on Earth of the subscriber. The ground station terminates the satellite segment and forwards the traffic to the Internet.² Notice that this forces all traffic to enter the Internet from the location where the ground station is. In our measurement setup, we monitor the traffic managed by one satellite in geostationary orbit. This satellite offers service in Europe and Africa, from Ireland to South Africa. At the time of data collection, the satellite operator operates a single ground station in Europe, through which all traffic passes to reach all Internet services.

The satellite is equipped with multiple directional antennas, each managing a transmission beam that points to a specific region on Earth. This allows the reuse of frequencies to increase overall capacity while optimizing the use of spectrum reserved for satellite communications. Each beam acts as a separate and independent physical channel, providing aggregate capacity on the order of Gb/s, the actual capacity being configurable. Two separate beams (and frequencies) cover each area, one for the uplink (from users' CPE to the satellite) and one for the downlink (from the satellite to users' CPE). A separate beam pair also connects the satellite and the ground station.

On the shared uplink channel, the transmission of packets involves a complicated MAC protocol: a slotted-Aloha protocol allows the CPE to access the shared reservation channel the first time it needs to transmit. Then, a Time Division Multiple Access (TDMA) scheduling protocol run by the satellite allocates time-slots to each active CPE to avoid collisions and to fairly share capacity among the active users at each TDMA frame. The satellite then forwards the packets to the ground station via a dedicated high-capacity beam.

On the downlink channel, the ground station transmits the packets directly to the satellite, which then forwards them to the destina-

tion CPE by selecting the correct frequency and beam. In this case, the packets are broadcast to all receivers, which filters those destined for their CPE MAC address and discard the packets destined for other CPEs. In addition to the TDMA and MAC schemes, Forward Error Correction (FEC) and Automatic Repeat Request (ARQ) mechanisms provide a reliable data-link service. All in all, these proprietary algorithms provide a reliable, almost error-free, bi-directional point-to-point link between each CPE and the ground station. By combining these MAC, scheduling, FEC and ARQ protocols, further random delays are added to the communication between the CPE and the ground station.

To mitigate the potential performance degradation caused by this high latency, the SatCom operator relies heavily on a Performance Enhancing Proxy (PEP) to improve TCP performance on the satellite segment. A PEP is a network component that improves end-to-end performance by transparently manipulating TCP connections. Defined in RFC 3135 [37], PEP works as follows in our case. In the lower part of Fig. 1, the subscriber CPE acts as a transparent TCP proxy for the end user's TCP traffic. It terminates all TCP connections initiated by applications on end-user devices and forwards TCP payload to the ground station via a bidirectional reliable tunnel over UDP. In detail: When a subscriber's device initiates a new TCP connection via a SYN packet, the CPE impersonates the destination server and immediately completes the TCP three-way handshake, allowing the client application to send the initial data with no delay. Note that the PEP does not handle UDP traffic or any protocol on top of it – e.g., QUIC. Any UDP traffic is sent through the satellite link without any specific performance optimization, buffering or modification.

Here, the CPE acts as an L4 proxy. It buffers the TCP data stream and forwards it to the ground station via the bidirectional UDP tunnel at the allowed uplink rate. The ground station again works as an L4 proxy. When it receives a *Connect* request from the CPE, it establishes a new TCP connection to the actual destination server. It then forwards the data to/from the CPE via the downlink/uplink satellite tunnel. In this way, the TCP congestion control algorithm is effectively decoupled, allowing the ground station proxy to retrieve the data from the origin server at the backbone-path rate and the CPE to forward the data to the end user's device as quickly as possible. Note that the download rate of the ground station from the origin server is still regulated by the download rate of the end device because the buffer capacity of PEP in

¹ A customer may represent a single individual, a household, a company's office, or a community-based WiFi internet access solution.

² The total round trip time (RTT) of any communication is in the order of 550 ms since the packets must go through the satellite link on both the forward and backward paths.

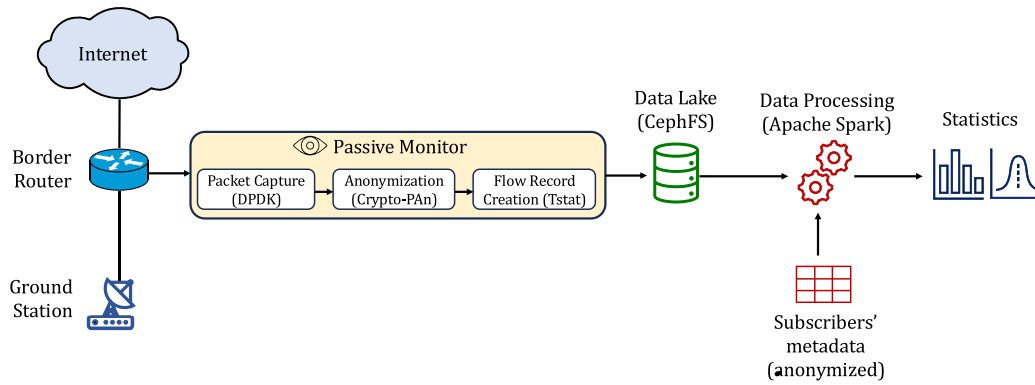


Fig. 2. Traffic capture and data processing flow.

the ground station is limited. Note that user traffic using UDP (e.g., DNS, QUIC) cannot benefit from PEP acceleration and therefore UDP packets are forwarded as is.³

At last, the ground station also acts as a DNS resolver, and supports Quality of Service (QoS) schedulers to prioritize and shape traffic depending on the application. To this end, the SatCom operator uses L3/L4 and domain name-specific rules to prioritize interactive traffic and shape video streaming flows. The traffic shaper of the ground station enforces commercial maximum capacity per subscriber of up to 5 Mb/s in the uplink, and 10, 20, 30, 100 Mb/s in the downlink based on the subscriber's contract.

The SatCom operator provides a dual-stack configuration, supporting both IPv4 and IPv6. Each subscriber is assigned a single (private) IPv4 and the ground station operates as a Network Address Translation (NAT) box for IPv4. The traffic thus undergoes a second Carrier-Grade NAT before reaching the Internet. Each subscriber is also assigned an IPv6 /64 or /60 prefix. Each user device auto-configures an IPv6 address, which it then uses to contact any destination. Notice that, in IPv6, NAT is not possible. Any connection (TCP or UDP) must be initiated by an end-user client, and no server can run on the customer's premises.⁴

3.2. Passive measurements

We instrument the SatCom operator's network to collect passive measurements of all subscriber traffic. The setup and the relevant steps of data capturing and processing are detailed in Fig. 2, along with the main technologies involved.

To this end, we deploy a passive probe at the operator's ground station. Here we collect all traffic after the operation of the PEP, which handles traffic in the satellite segment. We observe all packets exchanged by each customer – that we uniquely identify by their SatCom CPE IP address. Using two router span ports, we mirror both downlink and uplink traffic to a high-end measurement server.⁵ It is equipped with two Intel(R) Xeon(R) Gold 6140 CPUs @2.30GHz and 376GB of memory, running Ubuntu Server 22.04. The server also holds two Intel X710 network cards. Each card supports a data rate of up to 10 Gb/s per port and uses long-range optic transceivers to receive traffic from the router. The server runs Tstat [38], a custom flow monitoring software that generates rich per-flow summaries in real time from the processed data packets. To handle the high rate, our software resorts to the Data Plane Development Kit (DPDK) library for packet capture [39], which enables accelerated packet processing by bypassing the kernel-space drivers and protocol stack, and guarantees that all packets are processed in real time

without information loss. Indeed, using the traditional NIC drivers entails that packets are received and processed by the machine's kernel before being delivered to the monitor application. This would impair the performance, as the kernel cannot sustain such a high traffic rate. Using the classic 5-tuple, Tstat identifies and tracks the evolution of TCP and UDP flows. For each flow, it extracts hundreds of statistics for both flow directions. The metrics we mainly rely on are: the i) flow size and duration, ii) the timing information of the first 10 packets, iii) the server and client IP address, iv) the TCP RTT between data and ACK segments, and v) the name of the contacted server as retrieved by the *Server Name Indication* (SNI) of TLS, HTTP or DNS protocol.

The measurement of the TCP RTT deserves a careful explanation because the presence of PEP makes the measurement of RTT particularly troublesome. Indeed, the PEP causes the total RTT to be divided into three components, as shown in the bottom part of Fig. 1: (i) *home RTT* – between the user device and the user's SatCom CPE; (ii) *satellite RTT* – between the CPE and the ground station, where TCP segments are forwarded over the satellite PEP tunnel; and (iii) *ground RTT* – between the PEP terminator at the ground station and the destination server. Our vantage point is co-located with the ground station and therefore observes traffic exchanged from the PEP proxy to the Internet. For the ground RTT measurement (iii), Tstat uses the TCP connection initiated by the ground station PEP. For each TCP segment sent, it measures the time to the corresponding ACK, calculating the minimum, maximum, average, and standard deviation of all RTT samples in a TCP flow. To measure the RTT (ii) of the Satellite segment, we need an additional ingenuity. Specifically, we leverage the TLS handshake, initiated by the Client Hello message (CHLO), to measure the time from the Server Hello message (SHLO) to the next Client Key Exchange message/Change Cipher Spec message. This time also includes the Home RTT (i), which we can consider negligible compared to the satellite. In this way, and only for TLS flows completing the TLS negotiation, we can safely estimate the delay caused by the satellite at least once per flow. For both TCP and UDP, the software runs a Deep Packet Inspection (DPI) module that identifies the most popular protocols and extracts various information from headers. In particular, it annotates each flow with the server *domain name* as extracted from the *Host* header in case of plain-text HTTP, or from the SNI field in the case of TLS or QUIC flows.⁶ For QUIC flows that still use the initial "Google QUIC" version, the presence of the User Agent (UA) information in the TLS handshake returns coarse-grained details about the service, client device and application generating the traffic. In the latest version of QUIC, this data is encrypted. For DNS traffic, the software logs each requested domains and obtained responses, including the DNS server IP address the client used to resolve the name.

³ The PEP can only act as a L4 proxy without violating the authenticity provided by TLS.

⁴ The SatCom provider offers hosting of servers in the data center for customers interested in running services.

⁵ A similar setup would be possible using two optical splitters.

⁶ We use the term *domain* meaning Fully Qualified Domain Name.

3.3. Ethical aspects

Passive monitoring involves capturing and processing traffic generated by human beings; thus, we need to take proper actions to protect as much as possible the individual's privacy. Indeed, IP address is considered Personally Identifiable Information (PII) and it can be used to identify and track individuals. The characteristics of traffic, such as the list of visited websites, can be considered Sensitive Personal Information (SPI), as they can reveal personal aspects and habits of an individual.

For this work, we take all possible countermeasures to properly handle our measurements. First, the setup, management, and data collection were physically managed uniquely by the operator personnel, who control the data collection process. Second, we configured the data collection to limit as much as possible the exposed information. We process packets in real time and save only strictly required information in flow logs. In detail, we do not store any information present in headers that can be associated with a single user. Customers' IP addresses are anonymized in real time using the CryptoPan algorithm [40], which preserves the subnet structure of the original IP addresses.

The only metadata we collect are the server IP addresses, the domain names customers visited and the QUIC UA field, which we process only to extract aggregated statistics for the most popular services. Regarding the UA in QUIC, we limit our analysis to general information, such as the application and Operating System (OS), while discarding detailed data like application version, device model, or OS version to avoid overly specific information.

Third, we only have access to the anonymized logs that we store in our secure big data cluster, which is not reachable from the public Internet and has strictly controlled physical access. The operating system and software are kept up to date to avoid possible vulnerabilities, and strict user access policies limit access to the data only to authorized users.

The operator SatCom Data Protection Officer (DPO) has approved the above process and we have verified with our institutional review board that the data we collect is exempt from their approval.

4. Dataset processing and overview

In this extended version of our study, we focus on data collected during June-July 2023, resulting in 634 TB and a total of 5.6 billion flows. All analyses presented here, including new and updated results, are based exclusively on this dataset. On a daily basis, flow summaries are transferred from the measurement server at the SatCom provider premises to a big data cluster. Files are stored using the CephFS distributed file system, and they are post-processed using Apache Spark with custom-designed analytics to compute various statistics and distributions.

4.1. Data enrichment and aggregation

4.1.1. Service identification

In the first step of processing, we enrich the data by adding information about the customer's country (obtained by mapping the encrypted customer subnet to the corresponding country with the support of the SatCom operator) and about the service offered by the server. We focus here on six classes of services: Video Streaming, Social Networks, Audio Streaming, Chat, Work-related applications, and Search Engines. We rely on custom regular expressions that map popular server names to services. In detail, for each service class, we enumerate the top and local players by manually inspecting the list of most popular domains by volume and popularity. For each service, we enumerate the list of fully qualified domains and second-level domains used to serve its content.⁷ In some cases, we use regular expressions to generalize the set of domains. This is the typical case of CDN server names, which often

include numbers or country codes in the domain. We report the full list of domains and regular expressions in the [Appendix A](#). We uniquely use the domain to classify the service, as we do not capture the full HTTP URL, which is typically encrypted within the TLS session. For TLS flows, we obtain the domain from the SNI field of the Client Hello messages. As a result, we are sometimes unable to distinguish between sub-services from the same provider (e.g., Google Search and Maps both share the SNI *.google.com). As an alternative to manually curated lists, we could rely on online ranking and analytics tools (e.g., Alexa, Cisco Umbrella, or Similar Web). However, these services are known to list only the main domain of a given service, while not listing the domains of third-party services such as content-delivery providers and support services used by the first-party service. Given the small number of services we are interested in, we choose to create the lists manually.

The second step is to create aggregated views of the data to obtain traffic breakdowns by protocols, server domains, time (with 1 h granularity), country of the customer, and contacted service. This aggregation step facilitates subsequent data processing by reducing the amount of data to be processed by several orders of magnitude, enabling real-time data exploration.

4.1.2. Device identification

Following, we perform a second data enrichment step to identify the class of user devices. Our goal is to study the adoption of different classes and brands of devices in the monitored population. To this end, we consider uniquely IPv6 traffic, as the absence of the NAT allows us to run a classification at the device level.⁸ Indeed, all devices behind a subscription (thus using the same CPE) are assigned the same IPv6 prefix, but a unique per-device Interface Identifier (IID). This, combined with the absence of NAT, creates a unique identifier for each device on the network. We classify user devices into one of four categories: Linux, Windows, Apple, and Android.

To classify a device, we collect all the flows generated by the corresponding IPv6 address and perform a twofold process. First, we use the UA QUIC field if present in at least a single record as it describes the application and operating system of the device. We find the UA field only in a minority of devices still using the outdated Google QUIC version instead of the IETF standard one. Thus, for the rest of the traffic, we use a second method based on specific domains strongly associated with specific device types. These domains are typically used for services like operating system updates, network connectivity check (e.g., the Network Connectivity Status Indicator, NCSI in Windows⁹), or other built-in applications. We build sets of regular expressions matching domains contacted by those services and associate a device to a class in case we find a match. In the appendix, we report the full list of regular expressions that we match against domains.

Notice that we did not perform a systematic validation of our device identification approach; instead, we verified its correctness only on the limited set of devices available in our research labs, and we acknowledge this as a limitation of our study. Regarding the User Agent field in Google QUIC, it allowed us to spot a number of Android devices still using that outdated QUIC version. Despite the low number of identified devices, the User Agent is a valuable source of information as it leads to certain/deterministic classification. This classification allows us to compare the composition of device types across countries, while we do not perform any additional analytics at the per-device level.

4.2. Dataset overview

We first present an initial overview of the dataset by presenting the breakdown by geographic area and protocol. [Table 1](#) summarizes the

⁸ For IPv4, the presence of the NAT prevents us from distinguishing flows originated by different devices behind the same subscription.

⁹ <https://learn.microsoft.com/it-it/windows-server/networking/ncsi/ncsi-overview>

⁷ We handle the case of two-label top-level domains – e.g., co.uk.

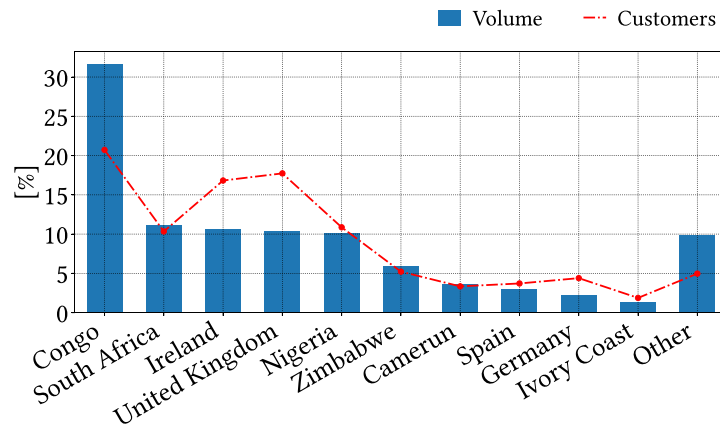


Fig. 3. Per country breakdown of traffic volume and user base.

Table 1
TCP/UDP traffic breakdown by protocols.

Protocol	Volume share
TCP/HTTPS	59.9%
TCP/HTTP	10.1%
Other TCP	3.5%
UDP/QUIC	22.9%
UDP/RTP	0.6%
UDP/DNS	< 0.1%
Other UDP	3.0%

Table 2
Distribution of data volume for L3 and L4 protocols, separately by continent.

	Europe		Africa	
	TCP	UDP	TCP	UDP
IPv4	60.3%	12%	69.5%	24.9%
IPv6	17.9%	9.8%	3.7%	1.9%

latter. As expected, web traffic accounts for most of the traffic, with HTTPS and QUIC accounting for 59.9% and 22.9% of the total volume, respectively. HTTP still accounts for a 10.1% share. Notably, compared to our original study, the presence of HTTPS and QUIC has increased. This is consistent with other studies that showed the convergence of Internet protocols towards encrypted web protocols [5,22,41]. Interestingly, despite the high latency due to the satellite link, we observe a non-negligible amount of video or voice traffic using Real-Time Transport Protocol (RTP).

Looking at traffic by country, we observe a large imbalance in the number of customers and thus traffic, as shown in Fig. 3. The blue bars indicate the share of data traffic for the top-10 countries. The red line shows the share of customers. The countries are sorted by decreasing total volume. Interestingly, the two figures are not completely proportional and show that customers in African countries consume much more traffic on average than customers in European countries. For example, Congolese customers are 21% of overall customers, but they generate 33% of volume (each generates about 640 MB per day). Britons are about 17% of customers, but generate only 11% of volume (each generates only 160 MB per day). This suggests that African customers may share Internet connections with multiple end-users, while European customers may resort to SatCom access only when forced to do so. In Section 6 we will explore this direction in more detail.

We now break down traffic volume by IP version and by L4 protocol (TCP or UDP). This allows us to measure the adoption of IPv6 across




Fig. 4. IPv6 usage by volume.

different continents. In Table 2, we report the percentage of volume data by IP version, separately for TCP and UDP protocols. The results show a disparity in IPv6 adoption between the continents. Europe exhibits a large share, 27.7% if we consider both TCP and UDP traffic. Conversely, in Africa, only 5.6% of the volume is IPv6. This highlights a significant disparity between the two continents, with Africa lagging well behind.¹⁰ In the following, we will dig into this disparity, analyzing how popular services adopt different IP versions in the two continents. Regarding L4 protocols, the share of TCP is similar across the two continents, with 78.2% and 73.2% for Europe and Africa, respectively.

Fig. 4 reports the percentage of traffic volume for the top 10 countries, with Europe on the right and Africa on the left. The y-axis represents the byte-wise percentage of IPv6, separately for download (blue bars) and upload (red bars). In European countries, the average IPv6 download volume accounts ~ 31.5% of total volume, with upload volume reaching a peak of 46% in the United Kingdom. Ireland and Germany follow with 29.9% and 38.6%, respectively, while the share of IPv6 is stable for uploaded data. Spain shows no IPv6 traffic at all, as the operator's reseller in Spain does not support IPv6. In Africa, the IPv6 volume share is significantly lower. On average 6.8% of download volume is IPv6, with an even lower percentage for uploads - except for Congo and Ivory Coast. In Ivory Coast, the IPv6 in upload reaches 26%, which is five times the average of the other African

¹⁰ See for example <https://blog.cloudflare.com/ipv6-from-dns-pov/> or <https://www.google.com/intl/en/ipv6/statistics.html> for up-to-date statistics on IPv6 adoption.

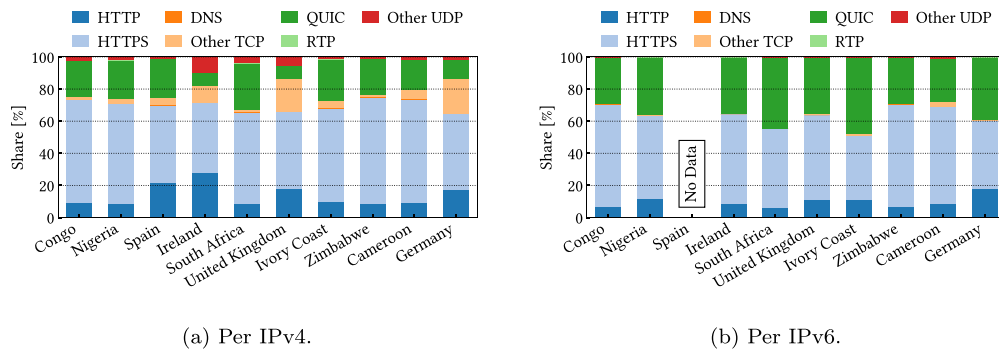


Fig. 5. Protocol share per country.

countries (5.2%). This is caused by a single subscriber responsible for 62% of uploaded volume in the country, with 95% of this volume going to <https://www.googleapi.com>, likely a cloud-based backup service. This difference is mainly due to the use of old devices in African countries, which do not support IPv6.

To complete the traffic overview, we break down traffic volume by L7 protocol. We limit the analysis to the top-10 countries, splitting the results into two Figures: Fig. 5a for IPv4 and Fig. 5b for IPv6 traffic.

The data exhibit some considerable differences across countries. In Fig. 5a (for IPv4 traffic) web protocols (HTTP and HTTPS) account for the majority of TCP traffic, while a substantial portion in European countries consists of non-web protocols. The case of Germany is extreme: 35% of all TCP traffic consists of non-web protocols. Manual inspection suggests that this is due to a higher use of Virtual Private Network (VPN) solutions (unknown protocol, non-standard port, long-lived flows without parallelism). Note also that the percentage of unencrypted HTTP traffic is higher in European than in African countries. This is due to popular Microsoft and Sky services that use HTTP to distribute software updates and video content. In contrast, Congo, Nigeria, and Ivory Coast show a very similar protocol breakdown. This may reflect the different customer base in Europe and Africa, with some business customers in the former (confirmed by the SatCom operator) using non-web-based protocols for VPN and internal services.

In Fig. 5b we break down IPv6 traffic. We observe a greater use of web-based protocols in all countries, with less than 1% not being web-based. UDP over IPv6 is used almost exclusively for QUIC traffic, as IPv6's modern infrastructure is preferred by large providers adopting QUIC. Additionally, there is a more widespread use of TLS for secure communication over IPv6, confirming the preference for more secure and modern communication with the latest IP version technology. The smaller amount of DNS traffic observed over IPv6 can be explained by the fact that the provider's recursive resolvers are reachable only via IPv4. As a result, even devices and networks that support IPv6 send their DNS queries over IPv4. The IPv6 DNS traffic we do observe is primarily generated by clients relying on public resolvers that offer IPv6 connectivity.

Given this initial overview, we limit our analysis below to the top 3 countries in Europe and the top 3 countries in Africa to compare usage, performance, and services used.

5. How much customers consume

In this section, we discuss the temporal patterns and volume of traffic generated by customers in different countries. We also discuss the users' device type and their habits in terms of subscription sharing.

5.1. Daily pattern

We start our analysis from the traditional hourly traffic pattern, which we report in Fig. 6. The y-axis reports the percentage of traffic

volume at a specific hour, normalized over the maximum value for the given country. For each time bin, we report the average value seen at that time during the whole time period, summing the upload and download traffic. We use the UTC time zone and countries in different time zones appear shifted.

We immediately observe how African (dashed lines) and European (solid lines) countries exhibit very different traffic patterns. In Europe, the traffic peak happens during evening prime time between 18:00-20:00-UTC. Conversely, during the day the traffic volume settles to lower values, down by 50% in the morning and as low as 20% at night. Conversely, in African countries, we observe a much higher traffic consumption during the morning too. For Congo (dashed red line), the absolute peak is at 9:00-UTC (10:00 local time). In Nigeria and South Africa, the morning peak reaches 90% of evening peak time. Notice also at night the low-peak that is almost as high as 40% of peak-time. This suggests that customers' use of the SatCom access differs between Europe and Africa, hinting at a classic leisure usage for the former.

5.2. Subscribers' traffic volume

We now focus on the volume of traffic generated per customer, per day. In Fig. 7, we report the empirical Complementary Cumulative Distribution Function (CCDF), using log-log scales. We measure both the total number of flows and the total byte-wise volume generated in each day by a given customer. We consider both TCP and UDP, and distinguish upload and download volume. We account for both IPv4 and IPv6 traffic generated by the same customer. Note that a single subscription generates many samples in the distribution, one per day.

Start from the number of daily flows per customer shown in Fig. 7a, and focus on the main body of the CCDF (the leftmost part of the figure). We clearly observe that a significant fraction of subscribers generates less than a few thousand flows in 24 h. In fact, the curves show a gradual slope change between 1k and 10k flows – i.e., overall only 85% of customers generate more than 1k flows in a given day. We thus opt to exclude from our analyses customers that are inactive on a given day, to allow a fair comparison of users' traffic and habits. From now on, we define *active customers* as those that generate at least 250 flows per day.

Focus now on the tail of the curve. Here we can clearly see that African countries present a heavier tail than Europeans, with, for example, Congo having 1% of customers generating more than 400M flows in a day. This is due to the presence of some WiFi access points that share SatCom access in community internet solutions or internet caf  s. In fact, most Africans in Congo and more in Nigeria and South Africa than in Europe, lack home Internet access. Therefore, people have to go to public places – Internet cafes, libraries, workplaces, etc. – to access the Internet. This situation was observed by the scientific community as early as 2013 [42] and more recently in 2019 [43]. The multiplexing of several end-users behind a single customer CPE IP address results in an inflated number of per-customer daily flows. In Section 5.3, we quantify this aspect by leveraging the absence of NAT for IPv6 traffic.




Fig. 6. Daily trends per country.

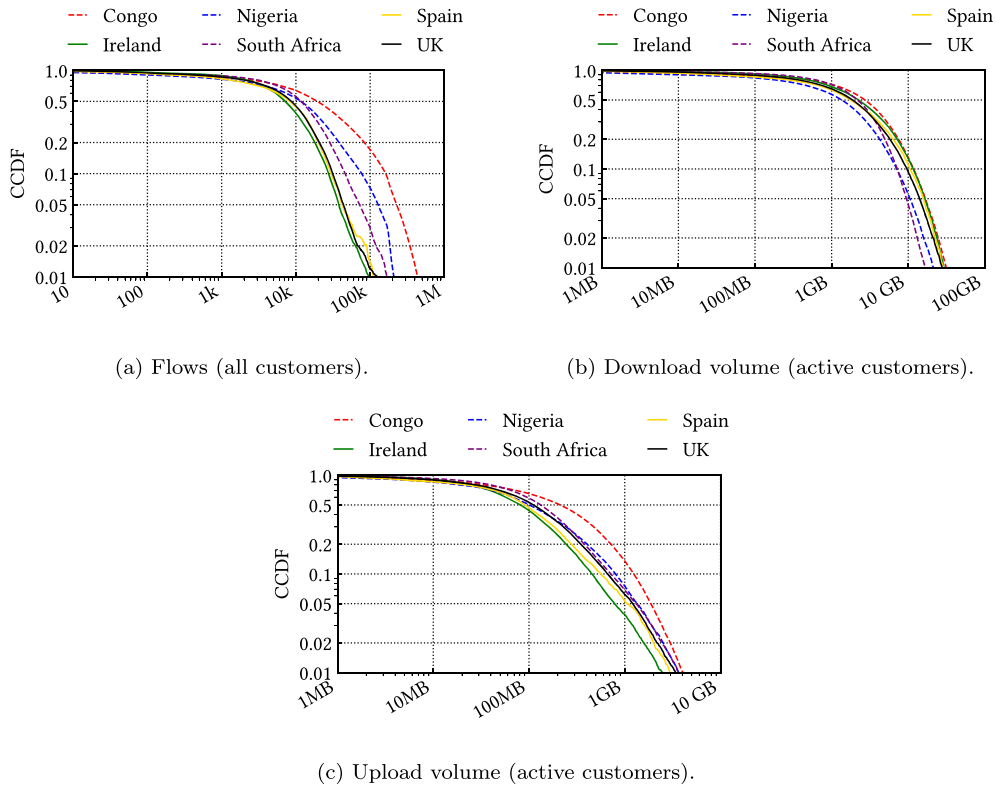


Fig. 7. Distribution of daily volume. Notice the log scale on both axes.

Next, we characterize the amount of daily downloaded and uploaded amount of traffic by each active customer. We show results in Fig. 7b and c, respectively, for downlink and uplink. First, we observe that there is no striking difference between European and African countries. The median downloaded data volume in Ireland is 2.2 GB, while Nigeria's is significantly lower, at 1.4 GB. However, Congo stands out with the highest median value observed (2.7 GB), exceeding even European countries. For instance, in Congo, the percentage of heavy hitters (those customers downloading more than 10 GB in a day) is twice as much (13%) than in South Africa (4%) and Nigeria (5%), following the behavior of European countries. When compared to the traffic in the previous month, we observe that the tail of the distribution for European countries is expanding – meaning that European heavy hitters nowadays do consume more traffic.

Interestingly, the difference between European and African countries is more pronounced in upload volume than in download volume. As shown in Fig. 7c African countries consistently exhibit a higher me-

dian upload value than European countries. Moreover, Congo, Nigeria and South Africa have 13%, 8% and 7% heavy hitters (those customers uploading more than 1 GB of data in a day), respectively, compared to less than 4–6% in the U.K., Spain and Ireland. As we will see in the next section, customers who tend to upload a lot of content tend to exhibit a large usage of instant messaging applications, likely sharing images and videos from their mobile app.

Overall, we find that SatCom customers generate similar traffic volumes as FTTH and ADSL customers. Compared to some recent work [5], the average download (upload) volume per FTTH customer was on the order of 1 GB (100 MB) per day in 2017. We thus observe a significant increase in the volume of traffic exchanged by customers despite the limited possibilities offered by SatCom access.

We finally remark that, in the above measurements, we account together IPv4 and IPv6 traffic by the same customer. However, not all customers necessarily hold IPv6-capable devices and, even if so, only a portion of the customer's traffic is carried over IPv6. In Europe a

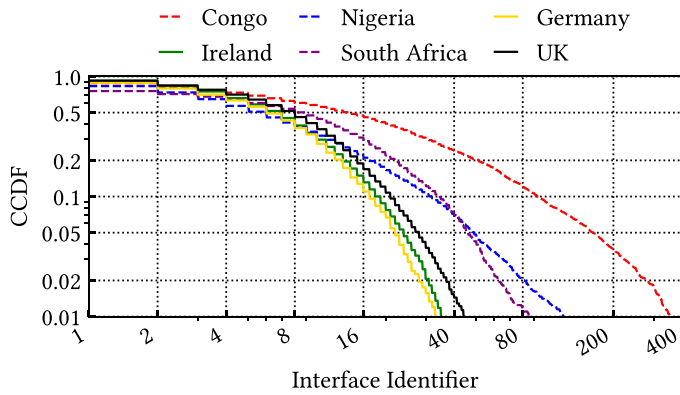


Fig. 8. Distribution of the number of IID per IPv6 prefix.

substantial portion of customers consistently utilize IPv6 for a large share of their traffic; many (40%) achieve over 50% of their daily volume on IPv6, while a minority of users (20%) still rely exclusively on IPv4. In Africa, however, IPv6 adoption remains limited, with only a small fraction of customers (15%) exceeding 50% of traffic volume on IPv6, and the vast majority (80%) relying exclusively on IPv4. In a nutshell, when a customer is provided with IPv6 connectivity and an IPv6-aware device, he/she generates approximately half of its traffic over IPv6.

5.3. Number of devices per CPE

Now, we study customers' habits in terms of subscription sharing. Indeed, each subscription can serve a household, a commercial activity, a small business, or even an entire village located in inaccessible areas. To this end, we leverage IPv6, where the absence of a NAT allows us to count the number of devices behind each CPE. Indeed, each device corresponds to a unique IPv6 Interface Identifiers (IID) within the IPv6 subscriber prefix. By quantifying the number of IIDs observed per Network Identifier (NI) per day, we have an insight into the number of users behind a single CPE. In Fig. 8, we report the CCDF of the unique IIDs per NI in a log-log scale. Each point represents the number of IID per subscription per day on the x -axis. The analysis counts a device if it generates more than 50 flows of at least 50 KB in download. Notice that counting devices cannot give a perfect estimate of the number of devices per subscription due to two aspects. First, we may obtain an overestimate, as some operating systems modify their IID periodically for privacy reasons. However, there is no way to match the same device using two or more different IIDs. Second, we could get an underestimate, as not all devices are IPv6 aware, and we would thus not spot them by looking at IPv6 traffic.

In European countries, the number of devices per subscription is relatively uniform compared to African countries. In Europe, the median number is 7 in Ireland and Germany, and 8 in the UK. Only 20% of German and Irish subscribers have more than 14 devices, while extreme outliers (<1%) have more than 35 in Ireland and 32 in the UK. In African countries, we find different values. In Nigeria and South Africa, the median value is respectively 5 and 9 devices per subscriber, similar to European countries patterns, but with a softer descent in the distribution tail. Congo has higher median values – 14 devices – but a heavy tail. Indeed, in Congo, 25% of subscribers have more than 40 devices – for comparison, this happens in less than 1% of cases in the UK. We observe peaks of over 360 devices connected to a single CPE. Similar peaks (<1%) are present in Nigeria and South Africa with 95 and 140 devices per CPE, respectively. These findings highlight the diverse patterns in SatCom usage, with a clear distinction between continents. European countries tend to have more consistent and lower values of devices per subscription, while African countries, especially in the case of Congo and South Africa, show high variability and larger peak values, hinting

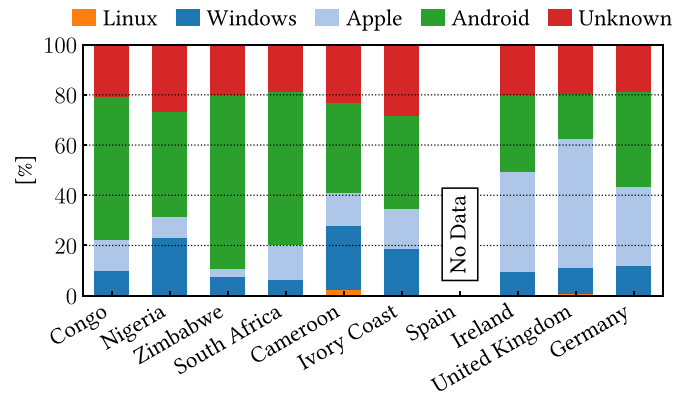


Fig. 9. Device classification results.

that SatCom subscriptions are likely shared between entire villages or businesses located in remote areas.

5.4. Type of devices

We now offer a coarse overview of the device type in the different countries. As detailed in Section 4.1, we use the User Agent field in the Google QUIC header and the domain names contacted by customers to classify them as Windows, Linux, Apple, or Android devices. We execute the match using IPv6 traffic only, since we want to classify the user device individually. Therefore we neglect IPv4-only devices.

The UA field of the Google QUIC protocol allows us to identify only 2.3% of devices. Among these, 95% resulting in Android devices, with 47% of the flows related to YouTube application and 11% to Google Mobile Services (GMS), which provide APIs for Android devices. We use domain names to classify those IIDs (i.e., devices using IPv6) that are active. Additionally, we discard a negligible portion of devices for which we find ambiguous matches – i.e., would be mapped to two or more device types. Cross-verification of the UA-based approach and domain-based labeling process demonstrate a 100% agreement between the two methods.

In Fig. 9, we show the percentage of classified devices by country and by the four types, with unclassified devices represented in red. On average, 78.2% of devices are classified. In Africa, classification rates are lower, with results from 81.2% in Congo to 71.2% in Ivory Coast. This disparity is attributed to the greater volume of IPv6 flows per single device in European countries. The device distribution that emerges from Fig. 9 highlights different regional patterns. In African countries, Android devices dominate, accounting for 69% and 60.8% of IIDs in Zimbabwe and South Africa, respectively, though with some variability, as Cameroon shows a reduced share of 35.5%. Windows and Apple devices hold the second position, with values varying by country, e.g., Apple devices represent 8.2% in Nigeria and 15.9% in Ivory Coast. Linux devices are almost absent in our results with only some presence of 2% in Cameroon. In Europe, Apple devices dominate the landscape, reaching a 51% share in the UK, while Android represents only the 18%. These results reveal a pronounced divergence in device distribution between African and European countries, especially regarding Android and Apple market shares.

6. What customers consume

We now examine the habits of SatCom subscribers in terms of services they access. As described in Section 3, we identify the services by examining the domain of TLS, HTTP, and QUIC flows. For each country, we extract the list of popular services, manually create regular expressions to identify them from the domain, and assign a category among Audio streaming, Chat, Search engine, Social, Video streaming, and Work.

Google	94.28	86.36	97.11	88.97	98.06	82.53	90.23
Whatsapp	92.93	73.83	94.34	67.74	89.68	58.20	81.35
Snapchat	61.13	49.39	28.19	51.26	21.94	41.26	47.03
Wechat	12.61	4.05	1.45	0.56	0.65	0.54	5.20
Telegram	3.32	3.44	1.93	0.28	0.65	0.54	2.65
Instagram	68.08	56.63	67.47	54.89	71.61	43.15	60.87
Tiktok	74.91	52.46	65.30	57.68	58.71	54.30	62.85
Netflix	18.88	27.89	56.27	66.48	60.00	63.04	38.79
Primevideo	6.83	8.85	15.18	34.78	40.65	42.74	17.33
Sky	19.07	6.39	10.12	35.75	9.68	41.80	19.93
Spotify	50.80	39.93	52.17	49.86	66.45	42.07	47.63
Dropbox	9.66	6.51	16.51	8.80	12.90	13.58	10.90
	Congo	Nigeria	South Africa	Ireland	Spain	UK	Average

Fig. 10. Heatmap of the service popularity in different countries.

6.1. Service popularity

In Fig. 10, we first give, for each category, the percentage of customers accessing different services on a daily basis, separated by country. We focus on a subset of the services for which we can write regular expressions that match domains that we know the user intentionally visited. For example, we do not report on Social Media Networks or widely used services (e.g., YouTube) because these services often appear as third-party services in web pages (e.g., embedded social buttons, videos, tracking services).

Among Chat services, WhatsApp is the most widely used service - comparable to Google, which is the most popular service, as expected. This is consistent with our earlier findings [5]. Interestingly, Snapchat ranks second, while in Congo more than 12.61% of customers also use WeChat for communication. This suggests the presence of Chinese-related communities. Telegram has yet to gain momentum.

Instagram and TikTok have similar penetration in all countries. TikTok is, on average, slightly more popular than Instagram, with the leading platform varying across countries depending on local usage trends.

In the paid-video category, it is notable that they are more popular in Europe than in Africa, with only South Africa achieving similar penetration. This is likely due to both economic and cultural effects, as well as the investments these platforms make in each country. For completeness, as previously said, Sky uses HTTP rather than HTTPS for serving the video content, and its popularity in Europe leads to the increase in HTTP traffic observed in Fig. 5.

Our findings corroborate recent research on African Internet usage, which shows the strong presence of social and chat services. For example, a recent report from the Pew Research Center [44] shows that chat and social networking are much more popular than paid services in Africa. In fact, paid video streaming services are not yet very popular in Africa, with Netflix having an estimated 2 million users on the continent, according to the 2022 annual report from Digital TV consulting [45]. However, the same report predicts rapid growth, which is also confirmed by a report from Conviva, a major video distribution company [46]. We note that South Africa is peculiar among African countries and that the strong penetration of streaming in South Africa is already well known.¹¹

6.2. Service volume

Next, we focus on the volume of traffic generated when accessing the different service categories per country. Here we consider all services,

assuming that social buttons, tracking cookies, etc., consume little volume compared to customers using the actual service. Fig. 11 uses box-plots to show the distribution of daily volume per customer accessing each category. The box extends from the lowest to the highest quartile, with a line at the median. Whiskers that extend from the box show the 5th and 95th percentiles.

Audio streaming services consume the least amount of traffic in Africa and slightly more in Europe, where some customers consume more than 50 MB in some days. Chat application usage is much more heterogeneous. While customers in Europe consume a median of less than 10 MB per day for Chat services, Spain is a small exception, with a median slightly above 10 MB. In African countries, this value surprisingly increases by more than three orders of magnitude in African countries. Customers in Congo have a daily median of 250 MB, with the top-5% of the heaviest customers consuming more than 2 GB on some days. These are likely those community WiFi Access Point (AP) that share SatCom access with multiple end users. The same effect is observed in the Social Media category, with a daily median of 300 MB in Congo, but only 30 MB in European countries. In contrast, the differences are smaller in the Video streaming category. However, the share of video traffic comes from different services: free YouTube in Africa, and paid video streaming services in Europe.

6.3. Service protocols

We now dissect how the top services are accessed using different protocols, aiming at quantifying the emergence of IPv6 and QUIC. Indeed, the adoption of new protocols depends not only on the device and network but, clearly, also on the server-side support. In Fig. 12, we break down traffic volume for 7 top services by IP version and transport protocol. For this study, we consider only flows exceeding 100 KB in download size. Blue represents TCP while Green represents QUIC, with different shades per IP version (IPv6 darker than IPv4).

Although TCP is predominant in most cases, we observe that QUIC traffic already constitutes the majority of Google and YouTube traffic, for both Europe and Africa. This is no surprise as QUIC was originally developed by Google: Indeed, it accounts for 61.4% of traffic in Africa and 52.8% of Europe. Meta services (Facebook, Instagram and WhatsApp) also support QUIC, but to a lesser extent with varying proportions across the different platforms. For Instagram, QUIC represents respectively 21.3% and 24.5% in Africa and Europe, while for WhatsApp and Facebook, it does not reach the 10% of their volume. For all services, QUIC is more present in Europe than in Africa, with the curious exception of Google. Curiously, TikTok and Netflix do not use QUIC at all; it is hard to find the reasons behind the choice of these two over-the-top content providers, but this reflects a conservative approach

¹¹ <https://www.finder.com/za/streaming-statistics>

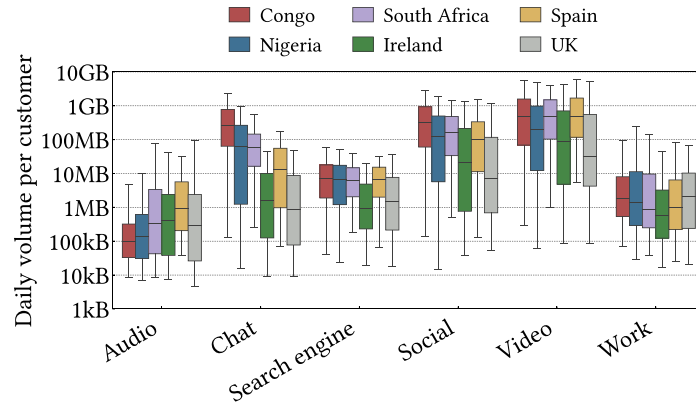


Fig. 11. Boxplot of the daily volume consumption per customer when accessing different service categories.

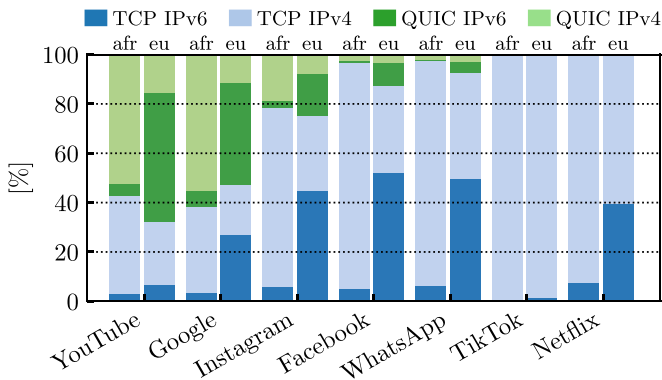


Fig. 12. IP version share per service and continent.

concerning novel protocols. Notice that TikTok is the only service not using IPv6 at all.

Finally, we take a look into IPv6, which all services but TikTok adopt. Remains evident the difference in IPv6 adoption between continents, where European services seem to be more into the IPv6 standard. Indeed, in Europe, more than half of traffic to Google and Meta is carried over IPv6. This means that, for over-the-top services, IPv6 already represents the majority of traffic. The situation is different in Africa, where structural delays in IPv6 adaptation can be attributed to the lower share of devices that support IPv6 across the continent.

7. Which performance consumers get

We now discuss the performance SatCom customers get. We focus on classical Quality of Service (QoS) indicators, namely RTT and throughput, and finally drill down on DNS performance. We then compare TCP and QUIC performance, as the presence of the PEP makes them behave differently in this SatCom scenario.

7.1. Satellite RTT analysis

Here, we examine how the satellite access link and SatCom network architecture affect end-to-end RTT. We consider the satellite RTT and the ground RTT separately, as we defined in Section 3.

Focus first on the satellite RTT shown in Fig. 13. We report the measurements at night – when we expect low congestion on the satellite link – and at peak time. The dashed lines indicate the median, while dotted lines are used for the 25th and 75th percentiles. As expected, the minimum satellite RTT is above 550 ms. However, the distributions show very large variability with RTT that can be higher than 2 s and varies widely in each country. This variability is due to several factors:




Fig. 13. Distributions and quartiles of the satellite RTT per country for night time (from 2:00 to 5:00 local time) and peak time (from 13:00 to 20:00 local time).

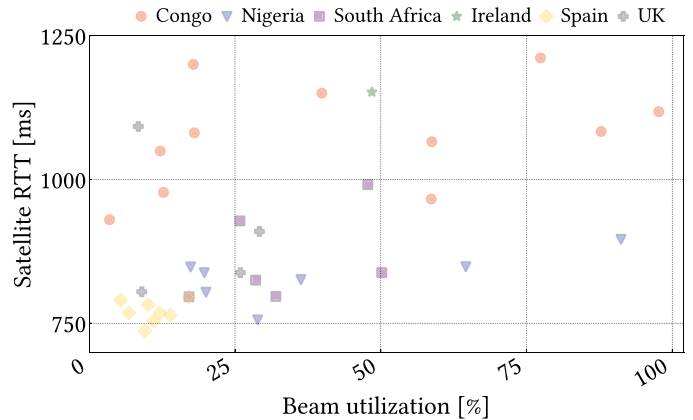


Fig. 14. Median satellite RTT per beam.

Queuing delay at various forwarding elements; Processing and transmission delay for limited-performance terminals; Packet losses and TCP retransmission; but the main reason is the SatCom access technology. Specifically, Nigeria has the best RTT at night in general. This is due to Nigeria favorable position, where the satellite is closer to the zenith. Follow Spain, South Africa and U.K. suffer from a larger zenith angle and thus a larger RTT.

In contrast, Congo and Ireland suffer from a much higher and more variable RTT. For Congo, the main cause could be the congestion on the satellite beams covering the country or the saturation of the bandwidth portion allocated by the SatCom operator to the Virtual Network Operator (VNO) offering Internet services to end users in Congo. In such

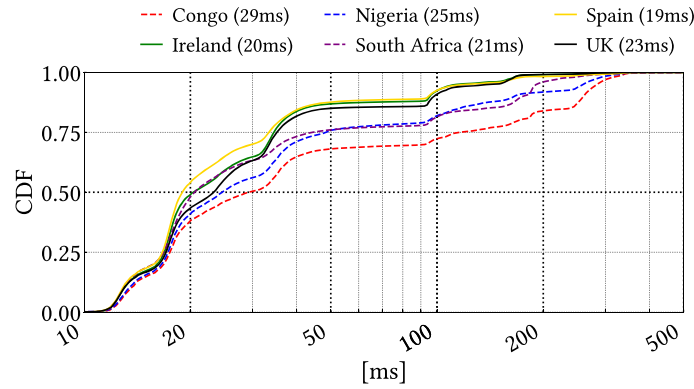


Fig. 15. Ground segment RTT computed as the average RTT in each TCP flow. Legend details the median.

bandwidth-constrained scenarios, the MAC protocol and the PEP scheduler may delay the transmission of packets by several frames, affecting the satellite RTT. For example, note that about 20% of RTT samples are longer than 2 s. The high RTT values already occur during periods of low traffic and worsen during periods of high traffic, with RTT values increasing significantly - compare median values.

To give more details, Fig. 14 shows the median satellite RTT for each beam and relates it to the beam utilization¹². We consider the peak time interval and mark each beam with the corresponding country it serves. Both Congo and Ireland suffer from high delay almost regardless of beam utilization. Nigeria, Spain and U.K. exhibit in general lower per-beam RTT. We hypothesize that some congestion occurs, due to beam capacity or PEP processing capabilities. This, in turn, slows down the forwarding of packets, especially during the initial phase of the connection setup. This clearly shows the overall complexity of the SatCom access technology, which has implications for end-user quality of experience.

7.2. Ground RTT analysis

We now focus on the ground RTT shown in Fig. 15. This RTT considers only the part of the path between the SatCom ground station and the server in the Internet.

In general, the ground RTT is much more deterministic than the satellite RTT. Here, clear bumps reflect the proximity of the servers on the Internet to the SatCom ground station. Focus on the European countries. The closest group has an RTT of about 12 ms and serves about 20% of the traffic. These are CDN nodes of well-known players with widespread infrastructures with which the SatCom provider has direct peering agreements. A second group of servers is around 15 – 17 ms, and a third around 35 ms. All of these servers are located in Europe and serve more than 80% of the traffic for European customers. Continuing this analysis, we observe another group of servers at about 95 ms (180 ms). Most of these servers are cloud servers located in the East (West) coast of the U.S., which suffer from the latency of crossing the Atlantic Ocean (and the U.S.). This reflects the typical RTT on the common Internet paths [5–7].

Now look at the RTT for African countries. Surprisingly, they have a higher RTT than their European counterparts. Since all traffic must be routed through the same ground station in Europe, African countries experience additional ground RTT when the final server is located in the original African country, e.g., when the end-user accesses a local service that is not served by global CDNs. In other words, the location of the ground station forces all traffic to be routed through it. This creates the rightmost bumps, where RTT on the ground increases to 300 – 400 ms.

By manually examining the services offered by these servers, we can confirm that they are likely popular services in the country of origin. Again, we also observe a significant proportion of Chinese services that are particularly popular in Congo (note the last bump in the ground RTT). This is related to the presence of Chinese companies in the country.

7.3. DNS performance

Given the importance of the server IP address location to reduce latency, DNS resolution plays an important role. For this, we drill down on DNS resolver choice and performance. We consider DNS/UDP traffic, for which we observe the original end-user device request and resolver response. First, we look for popular DNS resolvers, quantify their resolution latency, and next we observe the impact on server choice.

In total, we observe 4 195 different resolvers, some of them only sporadically. Interestingly, we found that customers use well-known open resolvers instead of operator resolver, and strangely choose custom, unusual, and geographically distant resolvers. In Fig. 16, we break down the top-8 resolvers in terms of volume, separated by country. For a given country, each column shows the percentage of DNS traffic for the different resolvers. In the rightmost column, we report the median response time observed at the ground station.

We note that the operator DNS (first row) is quite used only in European countries. In Ireland, Spain, and the UK, it accounts for 44%, 29%, and 38% of the DNS volume, respectively. With a mean resolution time of only 3.98 ms, it offers the best performance. As expected, Google DNS is popular everywhere. In Africa it resolves 86% of requests in Congo and more than 50% in the other African countries. Other popular open resolvers, namely CloudFlare DNS and OpenDNS, have a different popularity, usually below 10%. The resolution time for all of them is on the order of 20 ms.

Nigeria is a peculiar case. We find that 12% of traffic goes to a local Nigerian operator resolver. For this resolver, the ground RTT artificially inflates the resolution time to about 120 ms since packets have to travel from the ground station in Europe to Nigeria and back.

Interestingly, we observe two Chinese DNS resolvers (Baidu DNS and 114 DNS) in African countries, confirming the assumption that there is a significant Chinese community that use homeland resolvers. For Baidu, the resolution time is terrible, with a median response time higher than 350 ms that must be added to the satellite RTT to reach the actual end-user device.

In summary, in most cases SatCom customers do not adopt the operator DNS and resort instead to open resolvers. Due to the particular routing in the SatCom network, we observe cases of resolvers that suffer very high RTT, yet they are widely used in Africa. This greatly impacts the DNS response time and introduces an additional 100 – 300 ms delay on top of the satellite RTT. This has a clear negative impact on user experience.

¹² We normalize the results to the maximum utilization observed across all beams to avoid disclosing the actual per-beam utilization.

Operator	Congo	Nigeria	South Africa	Ireland	Spain	UK	Median Response Time [ms]
Operator-EU	0.87	9.10	1.87	43.75	28.95	38.10	3.98
Google	85.68	50.69	63.47	38.49	61.27	34.67	21.98
CloudFlare	3.02	2.54	10.36	2.03	2.05	6.04	19.97
Nigerian	0.00	11.84	6.32	0.00	0.00	0.00	119.98
Open DNS	1.22	4.00	0.65	0.49	0.72	6.97	17.99
Level3	0.45	7.63	0.09	0.00	0.00	0.49	23.99
Baidu	0.68	0.32	0.22	0.12	0.11	0.05	355.97
114DNS	2.97	3.43	1.64	0.05	0.03	0.01	109.98
Other	5.11	10.46	15.38	15.07	6.87	13.67	29.97

Fig. 16. Adoption and median response time of DNS resolvers.

7.4. Implications on server selection policies of CDNs and DNS resolvers

The superposition of i) routing constraints through the same ground station in Europe, ii) an intercontinental service presence that includes African and European countries, iii) the adoption of different DNS resolvers by customers, creates a very tangled picture that complicates the server selection policies of different CDN and DNS resolver operators. To examine these implications, we observe whether there are differences in ground RTT to the same service when using different DNS resolvers.

Table 3 shows some examples. We report the average ground RTT for some sample domains and some of the most popular resolvers for Nigeria and U.K. A more complete version of this table can be found in Appendix C. For U.K. (and European countries in general), the DNS resolver has little impact on performance. This is expected since i) the ground station is located in Europe and ii) customers tend to access European services that are well served by CDNs in Europe.

However, this is not the case for African countries such as Nigeria. For example, the server IP address resolved to serve the `captive.apple.com` service results in 19.1 ms if resolved via the Operator-EU DNS for U.K. customers. It results instead in 110.4 ms if resolved by the 114DNS for customers in Africa. Interestingly, even the Google DNS resolver returns two different CDN nodes for U.K. (26.0 ms) and Nigerian (38.4 ms) customers. These resolvers provide more distant CDN IP addresses because they are likely confused by the originating customer request geo-position which conflicts with the routing through Europe. Not shown for brevity, we even observe that some DNS resolvers point to some CDN server in the original African country of the customer. This clearly inflates the ground RTT by several hundreds of milliseconds. Some domains, e.g., `nflxvideo.net`, are less affected by these phenomena. This may be because resolvers and CDNs have accurate information, or because they do not rely on DNS resolution to determine the closest CDN node, e.g., because they use Anycast-based CDN solutions (which are not affected by the DNS resolution issue).

A possible solution to the DNS inconsistency problem is to either force the use of the SatCom operator's resolver or work with the Open Resolver providers to correctly instruct the server selection policies to return the closest server to the ground station instead of the original location of the end user's terminals.

7.5. Throughput analysis

We finally briefly discuss the download throughput of SatCom customers. This poses significant challenges for a number of reasons. First, flow records allow us to measure the average per-flow bitrate, which depends on several factors – the channel capacity, congestion control in place, and type of data transfer (e.g., a bulk download of a file, or video streaming). Thus, the throughput does not necessarily indicate the available bandwidth at the customer. Moreover, the SatCom environment is highly peculiar: Recall that we observe the TCP data flow from the Internet server to the ground station PEP. Thus, download throughput is regulated by the PEP.

Table 3

Average ground segment RTT per country and DNS resolver.

	UK		Nigeria		
	Op-EU	Google	Op-EU	Google	114DNS
<code>captive.apple.com</code>	19.1 ms	26.0 ms	23.1 ms	38.4 ms	110.4 ms
<code>play.googleapis.com</code>	16.3 ms	17.7 ms	38.7 ms	36.0 ms	114.2 ms
<code>*.nflxvideo.net</code>	–	25.5 ms	33.6 ms	28.8 ms	20.1 ms


We first compare the average throughput of TCP connections, by calculating the gross ratio between bytes downloaded and the duration of the flow (calculated from the first to the last TCP segment with data sent). To obtain a reliable measurement, we only consider flows large enough for the throughput to reach stable values. For this purpose, we only consider flows carrying at least 10 MB of data, for which we limit the representations to 1 million samples from a three-day interval. Even in these cases, not all flows can be considered valid bulk download samples (e.g., persistent HTTP flows or rate-limited video streaming flows), and competing traffic may limit throughput. This figure can only be considered a rough estimate of the actual performance a customer gets.

Fig. 17a shows the CCDF for download throughput separately by country. The operator offers several commercial plans with different maximum throughput. This is reflected in the knees of the curves in the figure. In Europe, where 30 Mb/s, 50 Mbit/s, and 100 Mbit/s plans are popular, we find that these customers can saturate their capacity with a single TCP flow. Overall, European countries have similar download throughput, with Ireland achieving slightly lower values due to its particular physical channel characteristics (see Section 7.1). For brevity, we limited our analysis to video streaming flows, and separated off-peak and on-peak times. In both cases, we could not find any signs of congestion.


In the African countries, the picture is quite different. First, the operator sells plans with a capacity of 10 Mb/s and 30 Mbit/s. Only few customers can saturate this capacity, with Nigerian customers tending to achieve slightly higher throughput. This is likely due to lower traffic on the satellite link. However, the higher utilization of the link, the less optimal server selection and routing, the presence of community WiFi APs, and likely the less powerful end-user terminals limit the maximum download throughput that customers can achieve.

This is confirmed by Fig. 17b, which shows the distribution of download speed for each country, with nighttime hours separated from peak hours. Again, European customers have higher throughput than African customers. In all countries, throughput is lower during peak hours than at night, as shown in the body of the distribution and the lower percentiles and medians. The change is more pronounced in Congo and South Africa.

We conclude by comparing the performance of TCP and QUIC, the latter being a novel and rapidly emerging protocol that, however, complicates the operation of the Performance Enhancing Proxy (PEP). The results are presented for two major over-the-top content providers:



(a) Per country.



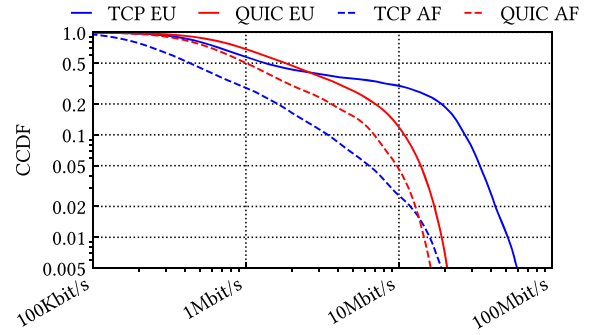
(b) Per time of the day.

Fig. 17. Download speeds per customer.


Google (including YouTube and Google services) and Meta (with Instagram, Facebook, and WhatsApp). We select these services due to their global presence and the high number of connections. To ensure a fair comparison, we present the results for each content provider separately in Fig. 18a and b, for Google and Meta, respectively. The plots depict the CCDF of the average per-flow bitrate on a log-log scale, with TCP represented in blue and QUIC in red. Dashed lines denote data from Africa, while solid lines correspond to data from Europe.

For Google, QUIC achieves superior median throughput on both continents: 1 Mbit/s in Africa and 2 Mbit/s in Europe, nearly double the TCP performance. This advantage can likely be attributed to QUIC's rapid connection establishment and faster handshake process (up to 0-RTT in cases of session resumption). However, QUIC struggles to achieve high throughput. As illustrated by the tails of the red lines in Fig. 18a, the throughput for QUIC appears capped at 20 Mbit/s for both Europe and Africa. This limitation is likely due to the absence of PEP support for QUIC traffic. In this case, the client and server rely solely on their built-in congestion control algorithm – CUBIC being the default for QUIC. Consequently, communication over a high bandwidth-delay channel, such as SatCom, is constrained by the intrinsic limitations of sliding window protocols when not explicitly optimized for such scenarios.

Conversely, with TCP, in Europe, 20% of flows surpass a throughput of 20 Mbit/s, potentially due to the PEP's enhancements which improve TCP performance in high-delay environments. Furthermore, 2% of TCP flows achieve throughput exceeding 30 Mbit/s, a level not observed for QUIC. In Africa, these trends are less pronounced due to higher network traffic, the larger number of devices sharing a single subscription and the widespread use of subscription plans with limited speed. Nevertheless, the dashed blue curve surpasses the dashed red curve at the



(a) Per Google services.



(b) Per Meta services.

Fig. 18. Download speeds per protocols.

tail, suggesting that, even in Africa, the highest throughput flows are predominantly achieved using TCP.

The results for Meta (Fig. 18b) are even more striking. In Europe, TCP achieves a median throughput of 11 Mbit/s - an order of magnitude higher than QUIC. Additionally, 2% of TCP flows exceed 50 Mbit/s, while QUIC rarely reaches 10 Mbit/s. In Africa, QUIC achieves a better median throughput of 0.7 Mbit/s compared to TCP's 0.6 Mbit/s. However, TCP demonstrates a longer tail, with the top 2% of flows exceeding 11 Mbit/s, compared to 6 – 10 Mbit/s for QUIC.

In summary, QUIC performs comparably to TCP when low bandwidth is sufficient or available, despite lacking PEP support. However, its throughput seems limited, likely due to the substantial delays characteristic of the SatCom channel. Conversely, TCP, assisted by the PEP, is capable of achieving significantly higher throughput under favorable conditions.

8. Conclusion

In this work, we provided a thorough characterization of the SatCom network through passive measurements. By instrumenting the ground station of a major SatCom operator, where traffic from tens of thousands of customers is aggregated, we had the opportunity to characterize the traffic from different angles, providing a new perspective on a mature but complex technology.

Our results shed light on SatCom usage habits and traffic patterns, allowing us to compare Internet usage in different countries. We find that SatCom customers in Europe are more likely to be active during evening peak hours, while in Africa we observe a traffic peak in the morning. This suggests a classic leisure usage by European customers, while in Africa (at least partially) a business usage. Our measurements also show that African subscriptions are used by multiple people. Moreover, they consume many more chat services, which is hardly consistent

with sole or domestic use. Customers can reach the nominal throughput of their plan and often access popular high-definition video streaming platforms. SatCom traffic reaches the Internet via a single ground station in Europe, which adds delays in accessing destinations in other continents. Nevertheless, we found that a non-negligible proportion of African customers use DNS resolvers in other continents. Finally, we find that the new QUIC protocol struggles to achieve top throughput, as the network cannot provide any support due to full-header encryption.

This work also offered valuable insights into the overall mixture of Internet traffic. Over more than two years of continuous measurements, we identified some interesting shifts in traffic patterns. Although providing a comprehensive longitudinal analysis is beyond the scope of this paper, our observations highlight trends worth mentioning. In particular, we observed that the share of QUIC traffic increased from 20% (January 2022) to 27% (May 2024), while no significant trends emerged for the other protocols. At the same time, TikTok experienced a boom in traffic volume, rising from 2% (January 2022) to 6% (May 2024), whereas other bandwidth-hungry services, such as YouTube and Netflix, remained at stable levels.

CRedit authorship contribution statement

Gabriele Merlach: Writing – review & editing, Writing – original draft; **Daniel Perdices:** Writing – review & editing, Writing – original draft; **Gianluca Perna:** Writing – review & editing, Writing – original draft; **Martino Trevisan:** Writing – review & editing, Writing – original draft; **Danilo Giordano:** Writing – review & editing, Writing – original draft; **Marco Mellia:** Writing – review & editing, Writing – original draft.

Data availability

The authors do not have permission to share data.

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The research leading to these results has been funded by the SmartData@PoliTO Center for Big Data Technologies. We also acknowledge funding support from the Comunidad de Madrid under the project RAMONES-CM (TEC-2024/COM-504) and the project LINO (SI4/PJI/2024-00221) through the direct grant agreement for the promotion of research and technology transfer at Universidad Autónoma de Madrid and from the Spanish State Research Agency under the project MOAI (PID2024-158424OB-C22). Additional support was provided by the projects 20228FT78M “DREAM” and 2022M2Z728 “COMPACT” under the Italian Ministry of University and Research 2022 PRIN program.

Appendix A. Services

[Table A.1](#) reports the regular expressions defined to identify services and categories from domain names.

Table A.1
Regular expressions used to identify services and service category.

Service	Regex	Category
Spotify	[.spotify.com\$, .scdn.co\$, .scdn.com\$]	Audio
Youtube	[.googlevideo.com\$, .ytimg.com\$, .youtube.com\$, .gvt1.com\$, .gvt2.com\$, .youtube-nocookie.com\$]	Video
Netflix	[.netflix., .nflxext., .nflximg., .nflxvideo., .nflxso.]	Video
Sky	[.sky.com\$]	Video
Primevideo	[amazonvideo.com\$, primevideo.com\$, pv-cdn.net\$, atv-ps.amazon.com\$, atv-ext.amazon.com\$, atv-ext-eu.amazon.com\$, atv-ext-fe.amazon.com\$, atv-ps-eu.amazon.com\$, atv-ps-fe.amazon.com\$]	Video
Facebook	[.facebook.com\$, .fbcdn.net\$, .facebook.net\$, .fbcdn, .fbstatic, .fbexternal, .fbxs.com\$, .fb.com\$]	Social
Twitter	[.twitter., .twimg., ^twitter.com\$, twitter.com.edgesuite.net, twitter-any.s3.amazonaws.com, twitter-blog.s3.amazonaws.com]	Social
LinkedIn	[.linkedin.com\$, .licdn.com\$, .lnkd.in\$]	Social
Instagram	[.instagram.com\$, .cdninstagram.com\$, ^igcdn]	Social
Tiktok	[tiktok.com\$, tiktokcdn, tiktokv.com\$]	Social
Google	[^www.google., ^google.]	Search engine
Bing	[.bing.com\$]	Search engine
Yahoo	[.yahoo.com\$, .yahoo.net\$, .ymg.com\$]	Search engine
Duckduck	[.duckduckgo.]	Search engine
Whatsapp	[.whatsapp.com\$, .whatsapp.net\$]	Chat
Telegram	[.telegram.org\$, ^telegram.org\$]	Chat
Snapchat	[.snapchat.com\$, feelinsonice.appspot.com\$, feelinsonice-hrd.appspot.com\$, feelinsonice.l.google.com\$]	Chat
Skype	[.skypeassets.com\$, .skype.com\$, .skype.net\$]	Chat
Wechat	[wechat.com\$, weixin.qq.com\$, wxs.qq.com\$]	Chat
Office365	[.sharepoint.com\$, .office.net\$, .onenote.com\$, .office365.com\$, .office.com\$, teams.microsoft, teams.office, lync, skype, live.com\$]	Work
Gsuite	[.googledrive.com\$, .drive.google.com\$, .docs.google.com\$, .sheets.google.com\$, .slides.google.com\$, .talk.google.com\$, .takeout.google.com\$]	Work
Dropbox	[dropbox, db.tt\$]	Work

Appendix B. Devices

Table B.1 reports the regular expression of the defined domains used to identify devices from domain names.

Appendix C. Ground segment RTT

For the sake of completeness, Tables C.1 and C.2 provide the average ground segment RTT for several countries and DNS resolvers.

Table B.1
Regular expressions used to identify devices category.

Label	domains
Apple	[.device-config.pcms.apple.co\$, .icloud.com\$, .contacts.icloud.com\$, .apple-mapkit.com\$, .apple-dns.cn\$, .ocsp2.apple.com\$, .stats.gc.apple.com\$, .oscdn.apple.com\$, .cdn-apple.com\$, .aaplimg.com\$, .itunes.apple.com\$, .push.apple.com\$, .metrics-config.icloud.com\$, .swscan.apple.com\$, .swcdn.apple.com\$, .swdist.apple.com\$, .appldnld.apple.com\$, .configuration.apple.com\$, .gdmf.apple.com\$, .ssl.ls.apple.com\$, .weather-data.apple.com\$, .fmfmobile.icloud.com\$, .swdownload.apple.com\$, .amp-api.podcasts.apple.com\$, .device-config.pcms.apple.com\$, .gs.apple.com\$, .bookkeeper.itunes.apple.com\$, .init.push.apple.com\$, .mesu.apple.com\$, .iphone-ld.apple.com\$, .courier.push.apple.com\$, .config-chr.health.apple.com\$, .apple.news\$, .metrics-config.icloud.com\$, .guizzoni.apple.com\$, .kittyhawk-ssl.ls.apple.com\$, .xp.apple.com\$, .mzstatic.com\$, .ig.apple.com\$, .config-chr.health.apple.com\$]
Android	[.android.clients.google.com\$, .android.googleapis.com\$, .connectivitycheck.android.com\$, .mi.com\$, .xiaomi.net\$, .micloud.xiaomi.net\$, .dbankcloud.eu\$, .dbankcloud.cn\$, .samsungvisioncloud.com\$, .samsungdiroute.net\$, .huawei.com\$, .harmonyos.com\$, .mi-img.com\$, .samsungdm.com\$, .api.sec.miui.com\$, .weather.oppomobile.com\$, .android.pool.ntp.org\$, .helpnewsrepublic1.ksmobile.com\$, .api.accuweather.com\$, .samsungapps.com\$, .dbankcdn.cn\$, .data.mistat.xiaomi.com\$, .samsapps.cust.ildns.net\$, .cmdts.ksmobile.com\$, .dbankedge.net\$, .android.apis.google.com\$, .samsungdec.com\$, .samsungdive.com\$, .findmymobile.samsung.com\$, .ms.cmc.com\$, .sbixby.com\$, .qepodownload.mediatek.com\$, .mdp-appconf-sg.heytaapl.com\$, .configcenter.iteldrive.com\$]
Windows	[.windowsupdate.com\$, .update.microsoft.com\$, .dl.delivery.mp.microsoft.com\$, .msftncsi.com\$, .msftconnecttest.com\$, .delivery.mp.microsoft.com\$, .maps.windows.com\$, .storeimages.microsoft.com\$, .dns.msftncsi.com\$, .msftspeechmodelsprod.azureedge.net\$, .telemetry.microsoft.com\$, .tlu.dl.delivery.mp.microsoft.com\$, .definitionupdates.microsoft.com\$, .edgeassetservice.azureedge.net\$, .arming.msedge.net\$, .g.ceipmsn.com\$, .ipv6.msftconnecttest.com\$, .watson.telemetry.microsoft.com\$, .arc.msn.com\$, .modern.watson.data.microsoft.com.akadns.net\$, .prod.do.dsp.mp.microsoft.com\$, .winget.ec.azureedge.net\$, .gameplayapi.intel.com\$, .microsoft-com.akamaized.net\$, .settings-win.data.microsoft.com\$, .client.wns.windows.com\$, .time.windows.com\$, .forefrontdl.microsoft.com\$, .ocsp.msocs.com\$, .storecatalogrevolution.storequality.microsoft.com\$, .activity.windows.com\$, .msn-com.akamaized.net\$, .sadownload.mcafee.com\$, .shasta-rrs.norton.com\$, .mediaredirect.microsoft.com\$]
Linux	[.ubuntu.com\$, .canonical.com\$, .fedoraproject.org\$, .debian.org\$, .debian.net\$, .archlinux.org\$, .kali.org\$, .linuxmint.com\$, .snapcraft.io\$, .redhat.com\$, .redhat.io\$, .clearos.com\$, .fedorapeople.org\$, .mirrorlist.centos.org\$, .centos.org\$, .freebsd.org\$, .turnkeylinux.org\$, .pop-os.org\$, .vyos.net\$, .vyos.io\$, .pop.system76.com\$, .security.ubuntu.com\$, .centos.mirror.garr.it\$, .centos.melbourneitmirror.net\$, .getfedora.org\$, .ubuntu.pool.ntp.org\$, .tizen.org\$, .launchpad.net\$, .fedorainfracloud.org\$, .clearos.uberglobalmirror.com\$, .joss.org\$, .snapcraftcontent.com\$, .centos.ca-west.mirror.fullhost.io\$, .archive.ubuntu.com\$, .launchpadcontent.net\$]

Table C.1
Ground segment RTT for most popular second-level domains for Congo and South Africa per DNS resolver.

Name Resolver	Congo										South Africa					
	Congo										South Africa					
	114	Aliyun	Baidu	Cloudflare	Google	Level3	Norton	OpenDNS	Operator-EU	114	Aliyun	Baidu	Cloudflare	Google	OpenDNS	Operator-EU
apple.com	118.3 ms	-	-	91.3 ms	169.8 ms	51.0 ms	39.9 ms	129.4 ms	75.4 ms	111.3 ms	-	-	62.6 ms	190.6 ms	187.3 ms	69.0 ms
doubleclick.net	150.6 ms	45.9 ms	292.1 ms	130.2 ms	62.7 ms	29.3 ms	46.4 ms	64.4 ms	45.4 ms	114.5 ms	27.5 ms	287.4 ms	73.9 ms	64.9 ms	21.1 ms	70.3 ms
facebook.com	68.1 ms	78.9 ms	59.8 ms	47.2 ms	53.4 ms	56.5 ms	23.1 ms	28.6 ms	40.2 ms	67.6 ms	132.9 ms	279.5 ms	42.6 ms	39.7 ms	29.3 ms	33.6 ms
fbcdn.net	40.6 ms	27.8 ms	31.9 ms	32.1 ms	29.7 ms	46.3 ms	190.8 ms	32.2 ms	41.9 ms	27.4 ms	35.1 ms	22.6 ms	31.4 ms	30.1 ms	27.7 ms	34.3 ms
google.com	154.0 ms	47.5 ms	186.0 ms	43.1 ms	79.8 ms	36.5 ms	47.5 ms	28.1 ms	77.9 ms	104.6 ms	61.2 ms	201.3 ms	106.3 ms	52.9 ms	98.2 ms	55.3 ms
googleapis.com	165.5 ms	42.4 ms	323.3 ms	66.6 ms	77.5 ms	54.7 ms	32.3 ms	31.3 ms	74.9 ms	111.0 ms	32.7 ms	275.3 ms	66.5 ms	64.2 ms	28.7 ms	79.5 ms
googlevideo.com	145.5 ms	34.3 ms	69.8 ms	95.0 ms	87.2 ms	63.4 ms	63.4 ms	25.5 ms	81.8 ms	24.2 ms	30.9 ms	86.9 ms	35.8 ms	38.5 ms	84.1 ms	48.4 ms
gstatic.com	146.8 ms	31.2 ms	286.1 ms	96.6 ms	67.9 ms	30.8 ms	39.1 ms	22.9 ms	80.2 ms	107.6 ms	27.5 ms	288.5 ms	67.8 ms	90.3 ms	824.3 ms	94.7 ms
microsoft.com	118.0 ms	-	281.0 ms	59.1 ms	101.1 ms	96.0 ms	77.7 ms	97.4 ms	77.0 ms	106.1 ms	-	-	46.6 ms	107.3 ms	107.1 ms	72.6 ms
netease.com	251.2 ms	193.2 ms	240.3 ms	247.9 ms	252.0 ms	209.4 ms	234.8 ms	-	172.1 ms	-	-	-	-	-	-	265.4 ms
qq.com	266.7 ms	254.3 ms	269.3 ms	269.2 ms	251.5 ms	255.4 ms	233.6 ms	254.2 ms	243.3 ms	-	-	253.8 ms	115.1 ms	253.8 ms	-	243.9 ms
scoopertnews	120.2 ms	50.0 ms	-	97.8 ms	132.9 ms	29.5 ms	36.1 ms	23.1 ms	36.5 ms	-	-	-	1026.0 ms	61.0 ms	-	19.9 ms
shallyr.com	82.0 ms	40.9 ms	43.3 ms	95.2 ms	68.1 ms	83.9 ms	66.7 ms	37.4 ms	69.1 ms	-	-	-	42.7 ms	41.7 ms	-	39.0 ms
tiktokcdn.com	215.2 ms	-	207.3 ms	124.8 ms	93.4 ms	29.6 ms	100.7 ms	69.1 ms	53.7 ms	137.3 ms	-	-	56.5 ms	173.6 ms	129.5 ms	73.3 ms
tiktokv.com	202.9 ms	24.9 ms	280.6 ms	94.5 ms	190.3 ms	50.5 ms	59.1 ms	114.7 ms	54.6 ms	89.3 ms	-	-	41.6 ms	193.0 ms	186.3 ms	48.5 ms
umeng.com	307.9 ms	106.7 ms	302.8 ms	300.8 ms	242.3 ms	252.2 ms	262.2 ms	82.1 ms	239.4 ms	346.1 ms	-	-	267.6 ms	140.9 ms	-	104.0 ms
whatsapp.net	70.9 ms	54.2 ms	85.2 ms	59.7 ms	51.8 ms	47.9 ms	47.3 ms	23.5 ms	44.6 ms	67.8 ms	134.9 ms	76.4 ms	38.6 ms	42.3 ms	20.6 ms	29.9 ms
yximgs.com	206.6 ms	184.7 ms	269.4 ms	174.6 ms	162.7 ms	-	-	174.2 ms	153.9 ms	-	-	-	182.8 ms	29.9 ms	-	89.2 ms

Table C.2
Ground segment RTT for most popular second-level domains for Nigeria and United Kingdom per DNS resolver.

Name Resolver	Nigeria										United Kingdom					
	Nigeria										United Kingdom					
	114	Aliyun	Baidu	Cloudflare	Google	Level3	OpenDNS	Operator-EU	Yandex	114	Aliyun	Baidu	Cloudflare	Google	OpenDNS	Operator-EU
apple.com	116.9 ms	264.9 ms	-	35.7 ms	76.6 ms	51.8 ms	64.1 ms	65.0 ms	56.0 ms	-	-	-	40.6 ms	70.7 ms	54.1 ms	45.1 ms
doubleclick.net	153.2 ms	31.8 ms	287.8 ms	28.9 ms	91.6 ms	29.3 ms	21.6 ms	41.2 ms	43.3 ms	-	-	284.4 ms	24.0 ms	23.3 ms	23.9 ms	26.9 ms
facebook.com	65.2 ms	41.7 ms	67.4 ms	48.1 ms	44.9 ms	44.0 ms	32.3 ms	42.2 ms	75.8 ms	-	-	33.9 ms	24.8 ms	20.5 ms	19.8 ms	25.7 ms
fbcdn.net	43.7 ms	39.2 ms	25.4 ms	40.2 ms	35.4 ms	24.2 ms	32.7 ms	45.0 ms	71.9 ms	-	19.1 ms	21.7 ms	24.4 ms	23.3 ms	26.3 ms	21.0 ms
google.com	147.9 ms	31.2 ms	162.2 ms	25.8 ms	75.6 ms	175.5 ms	27.6 ms	72.3 ms	151.9 ms	-	-	125.6 ms	27.0 ms	38.0 ms	20.7 ms	23.1 ms
googleapis.com	164.2 ms	34.6 ms	287.2 ms	25.0 ms	71.9 ms	29.7 ms	22.9 ms	95.0 ms	136.4 ms	-	45.3 ms	273.0 ms	23.2 ms	154.9 ms	22.3 ms	25.5 ms
googlevideo.com	106.4 ms	34.8 ms	64.2 ms	24.3 ms	55.2 ms	45.9 ms	28.3 ms	78.9 ms	113.2 ms	-	19.7 ms	293.0 ms	22.2 ms	24.1 ms	18.3 ms	23.7 ms
gstatic.com	147.5 ms	29.8 ms	299.8 ms	25.9 ms	101.2 ms	663.2 ms	23.6 ms	94.0 ms	40.2 ms	-	24.0 ms	-	24.0 ms	20.5 ms	19.2 ms	18.7 ms
microsoft.com	169.3 ms	-	-	-	93.2 ms	90.0 ms	79.7 ms	77.4 ms	46.9 ms	-	-	-	56.4 ms	56.7 ms	57.3 ms	36.6 ms
netease.com	245.4 ms	227.6 ms	286.0 ms	254.3 ms	249.2 ms	231.4 ms	201.5 ms	179.7 ms	273.8 ms	-	-	-	-	-	-	-
qq.com	245.2 ms	327.4 ms	282.2 ms	268.5 ms	244.7 ms	94.4 ms	226.5 ms	238.9 ms	90.6 ms	-	-	-	-	-	-	-
scoopertnews	101.5 ms	-	-	25.2 ms	25.9 ms	30.6 ms	21.1 ms	40.2 ms	140.2 ms	-	-	-	-	-	-	-
shallyr.com	83.3 ms	-	-	42.0 ms	97.7 ms	41.3 ms	42.9 ms	91.4 ms	93.7 ms	-	-	-	-	-	-	-
tiktokcdn.com	299.5 ms	-	-	30.7 ms	84.3 ms	35.0 ms	87.7 ms	73.2 ms	47.4 ms	-	257.9 ms	243.5 ms	-	30.5 ms	31.1 ms	-
tiktokv.com	177.5 ms	-	-	40.4 ms	111.4 ms	43.4 ms	101.3 ms	56.2 ms	56.6 ms	-	-	273.2 ms	-	31.7 ms	28.2 ms	-
umeng.com	305.1 ms	43.2 ms	301.5 ms	328.8 ms	216.3 ms	296.2 ms	201.6 ms	216.1 ms	216.1 ms	-	-	-	-	-	-	-
whatsapp.net	63.7 ms	65.8 ms	119.4 ms	23.6 ms	59.7 ms	58.7 ms	56.4 ms	51.3 ms	49.0 ms	-	-	-	23.3 ms	52.6 ms	14.0 ms	26.2 ms
yximgs.com	226.8 ms	36.8 ms	263.5 ms	-	157.9 ms	-	-	180.9 ms	-	-	-	-	-	-	-	-

References

- [1] J.L. García-Dorado, A. Finamore, M. Mellia, M. Meo, M. Munafo, Characterization of isp traffic: trends, user habits, and access technology impact, *IEEE Trans. Netw. Serv. Manage.* 9 (2) (2012) 142–155.
- [2] M. Sargent, M. Allman, Performance within a fiber-to-the-home network, *ACM SIGCOMM Comput. Commun. Rev.* 44 (3) (2014) (22–30).
- [3] D. Xu, A. Zhou, X. Zhang, G. Wang, X. Liu, C. An, Y. Shi, L. Liu, H. Ma, Understanding operational 5G: a first measurement study on its coverage, performance and energy consumption, in: *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication*, 2020, pp. 479–494.
- [4] A. Narayanan, E. Ramadan, J. Carpenter, Q. Liu, Y. Liu, F. Qian, Z.-L. Zhang, A first look at commercial 5G performance on smartphones, in: *Proceedings of the Web Conference 2020*, 2020, pp. 894–905.
- [5] M. Trevisan, D. Giordano, I. Drago, M.M. Munafo, M. Mellia, Five years at the edge: watching internet from the ISP network, *IEEE/ACM Trans. Netw.* 28 (2) (2020) 561–574.
- [6] T. Koch, E. Katz-Bassett, J. Heidemann, M. Calder, C. Ardi, K. Li, Anycast in context: a tale of two systems, in: *Proceedings of the 2021 ACM SIGCOMM 2021 Conference, SIGCOMM '21*, Association for Computing Machinery, New York, NY, USA, 2021, p. 398–417.
- [7] how much, Can edge computing change network latency? in: *2021 IFIP Networking Conference (IFIP Networking) (2021)* pp. 1–9. <https://doi.org/10.23919/IFIPNetworking52078.2021.9472847>
- [8] Y. Hu, V.O.K. Li, Satellite-based internet: a tutorial, *IEEE Commun. Mag.* 39 (3) (2001) 154–162. <https://doi.org/10.1109/35.910603>
- [9] O. Kodheli, E. Lagunas, N. Maturo, S.K. Sharma, B. Shankar, J.F.M. Montoya, J.C.M. Duncan, D. Spano, S. Chatzinotas, S. Kisseleff, et al., Satellite communications in the new space era: a survey and future challenges, *IEEE Commun. Surv. Tutorials* 23 (1) (2020) 70–109.
- [10] J. Border, M. Kojo, J. Griner, G. Montenegro, Z. Shelby, Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations, RFC 3135, (2001) (1–45).
- [11] J. Zhu, S. Roy, J.H. Kim, Performance modelling of TCP enhancements in terrestrial-satellite hybrid networks, *IEEE/ACM Trans. Netw.* 14 (4) (2006) 753–766. <https://doi.org/10.1109/TNET.2006.880167>
- [12] A. Pirovano, F. Garcia, A new survey on improving TCP performances over geostationary satellite link, *Netw. Commun. Technol.* 2 (1) (2013) 1.
- [13] F. Michel, M. Trevisan, D. Giordano, O. Bonaventure, A first look at starlink performance, in: *Proceedings of the 2022 Internet Measurement Conference*, 2022. <https://doi.org/10.1145/3517745.3561416>
- [14] M. Trevisan, A.S. Khatouni, D. Giordano, ERRANT: realistic emulation of radio access networks, *Comput. Netw.* 176 (2020) 107289.
- [15] D. Perdices, G. Perna, M. Trevisan, D. Giordano, M. Mellia, When satellite is all you have: watching the internet from 550 ms, in: *Proceedings of the 22nd ACM Internet Measurement Conference*, 2022, pp. 137–150.
- [16] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, S.C. Diot, Packet-level traffic measurements from the sprint IP backbone, *IEEE Netw.* 17 (6) (2003) 6–16.
- [17] M. Fomenkov, K. Keys, D. Moore, K.C. Claffy, Longitudinal study of internet traffic in 1998–2003, in: *Proceedings of the Winter International Symposium on Information and Communication Technologies*, 2004, pp. 1–6.
- [18] G. Maier, A. Feldmann, V. Paxson, M. Allman, On dominant characteristics of residential broadband internet traffic, in: *Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement*, 2009, pp. 90–102.
- [19] Z.S. Bischof, F.E. Bustamante, R. Stanojevic, Need, want, can afford: broadband markets and the behavior of users, in: *Proceedings of the 2014 Conference on Internet Measurement Conference*, 2014, pp. 73–86.
- [20] P. Borgnat, G. Dewaele, K. Fukuda, P. Abry, K. Cho, Seven years and one day: sketching the evolution of internet traffic, in: *IEEE INFOCOM 2009*, IEEE, 2009, pp. 711–719.
- [21] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, F. Jahanian, Internet inter-domain traffic, *ACM SIGCOMM Comput. Commun. Rev.* 40 (4) (2010) 75–86.
- [22] A. Feldmann, O. Gasser, F. Lichtblau, E. Pujol, I. Poese, C. Dietzel, D. Wagner, M. Wichtlhuber, J. Tapiador, N. Vallina-Rodriguez, O. Hohlfeld, G. Smaragdakis, The lockdown effect: implications of the COVID-19 pandemic on internet traffic, in: *Proceedings of the ACM Internet Measurement Conference, IMC '20*, Association for Computing Machinery, New York, NY, USA, 2020, p. 1–18. <https://doi.org/10.1145/3419394.3423658>
- [23] T. Favale, F. Soro, M. Trevisan, I. Drago, M. Mellia, Campus traffic and e-Learning during COVID-19 pandemic, *Comput. Netw.* 176 (2020) 107290.
- [24] T. Böttger, G. Ibrahim, B. Vallis, How the internet reacted to Covid-19: a perspective from Facebook's edge network, in: *Proceedings of the ACM Internet Measurement Conference*, 2020, pp. 34–41.
- [25] D.L. Johnson, V. Pejovic, E.M. Belding, G. Van Stam, Traffic characterization and internet usage in rural Africa, in: *Proceedings of the 20th International Conference Companion on World Wide Web*, 2011, pp. 493–502.
- [26] D. Minoli, Innovations in Satellite Communications and Satellite Technology: The Industry Implications of DVB-S2X, High Throughput Satellites, Ultra HD, M2M, and IP, John Wiley & Sons, 2015.
- [27] J. Deuschmann, T. Heyn, C. Rohde, K.-S. Hielscher, R. German, Broadband internet access via satellite: state-of-the-art and future directions, in: *Broadband Coverage in Germany; 15th ITG-Symposium*, 2021.
- [28] M. Tropea, P. Fazio, Evaluation of TCP versions over GEO satellite links, in: *2013 International Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS)*, IEEE, 2013, pp. 86–90.
- [29] F. Peng, Á.S. Cardona, K. Shafiee, V.C.M. Leung, TCP performance evaluation over geo and leo satellite links between performance enhancement proxies, in: *2012 IEEE Vehicular Technology Conference (VTC Fall)*, 2012.
- [30] M. Muhammad, M. Berioli, T. De Cola, A simulation study of network-coding-enhanced PEP for TCP flows in GEO satellite networks, in: *2014 IEEE International Conference on Communications (ICC)*, IEEE, 2014, pp. 3588–3593.
- [31] J. Iyengar, M. Thomson, QUIC: A UDP-Based Multiplexed and Secure Transport, 2021, (RFC 9000). <https://doi.org/10.17487/RFC9000>
- [32] A. Abdelsalam, M. Luglio, M. Quadri, C. Roseti, F. Zampognaro, QUIC-proxy based architecture for satellite communication to enhance a 5G scenario, in: *2019 International Symposium on Networks, Computers and Communications (ISNCC)*, IEEE, 2019, pp. 1–6.
- [33] L. Thomas, E. Dubois, N. Kuhn, E. Lochin, Google QUIC performance over a public SATCOM access, *Int. J. Satell. Commun. Netw.* 37 (6) (2019) 601–611.
- [34] N. Kuhn, F. Michel, L. Thomas, E. Dubois, E. Lochin, QUIC: opportunities and threats in SATCOM, in: *2020 10th Advanced Satellite Multimedia Systems Conference and the 16th Signal Processing for Space Communications Workshop (ASMS/SPSC)*, IEEE, 2020, pp. 1–7.
- [35] J. Deuschmann, K.-S. Hielscher, R. German, Satellite internet performance measurements, in: *2019 International Conference on Networked Systems (NetSys)*, IEEE, 2019, pp. 1–4.
- [36] M.K. Mohamed, A. Raman, D. Perino, N. Sastry, A browser-side view of starlink connectivity, in: *Proceedings of the 2022 Internet Measurement Conference*, 2022. <https://doi.org/10.1145/3517745.3561457>
- [37] J. Griner, J. Border, M. Kojo, Z.D. Shelby, G. Montenegro, Performance enhancing proxies intended to mitigate link-related degradations, 2001, (RFC 3135). <https://doi.org/10.17487/RFC3135>
- [38] M. Trevisan, A. Finamore, M. Mellia, M. Munafo, D. Rossi, Traffic analysis with off-the-shelf hardware: challenges and lessons learned, *IEEE Commun. Mag.* 55 (3) (2017) 163–169.
- [39] D. Intel, Data plane development kit, 2014.
- [40] J. Fan, J. Xu, M.H. Ammar, Crypto-pan: cryptography-based prefix-preserving anonymization, *Comput. Netw.* 46 (2) (2004) 253–272.
- [41] L. Schumann, T.V. Doan, T. Shreedhar, R. Mok, V. Bajpai, Impact of evolving protocols and COVID-19 on internet traffic shares, *arXiv preprint arXiv:2201.00142* (2022).
- [42] C. Stork, E. Calandro, A. Gillwald, Internet going mobile: internet access and use in 11 African countries, *info* (2013).
- [43] S.M. Mojapelo, The internet access and use in public libraries in limpopo province, South Africa, *Public Library Quarterly* 39 (3) (2020) 265–282.
- [44] L. Silver, C. Johnson, Internet Connectivity Seen as Having Positive Impact on Life in Sub-Saharan Africa, Pew Research Center, (2018).
- [45] D.T.V. Research, Africa SVOD forecasts, 2022, (<https://digitaltvresearch.com/product/africa-svod-forecasts/>).
- [46] Conviva, Conviva's state of streaming, 2021, (https://pages.conviva.com/rs/138-XJA-134/images/RPT_Conviva_State_of_Streaming_Q3_2021.pdf).