

A Characterization of the COVID-19 Pandemic Impact on a Mobile Network Operator Traffic

Andra Lutu
Telefonica Research

Diego Perino
Telefonica Research

Marcelo Bagnulo
Universidad Carlos III de Madrid

Enrique Frias-Martinez
Telefonica Research

Javad Khangosstar
Telefonica UK

ABSTRACT

During early 2020, the SARS-CoV-2 virus rapidly spread worldwide, forcing many governments to impose strict lockdown measures to tackle the pandemic. This significantly changed people's mobility and habits, subsequently impacting how they use telecommunication networks. In this paper, we investigate the effects of the COVID-19 emergency on a UK Mobile Network Operator (MNO). We quantify the changes in users' mobility and investigate how this impacted the cellular network usage and performance. Our analysis spans from the entire country to specific regions, and geodemographic area clusters. We also provide a detailed analysis for London. Our findings bring insights at different geotemporal granularity on the status of the cellular network, from the decrease in data traffic volume in the cellular network and lower load on the radio network, counterposed to a surge in the conversational voice traffic volume.

ACM Reference Format:

Andra Lutu, Diego Perino, Marcelo Bagnulo, Enrique Frias-Martinez, and Javad Khangosstar. 2020. A Characterization of the COVID-19 Pandemic Impact on a Mobile Network Operator Traffic. In *ACM Internet Measurement Conference (IMC '20)*, October 27–29, 2020, Virtual Event, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3419394.3423655>

1 INTRODUCTION

After the emergence of SARS-CoV-2 in the province of Wuhan (China) in December 2019, the virus rapidly spread to neighboring countries and to the rest of the world. It was declared by the World Health Organization a Public Health

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IMC '20, October 27–29, 2020, Virtual Event, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8138-3/20/10...\$15.00

<https://doi.org/10.1145/3419394.3423655>

Emergency on January 30th, 2020 and a pandemic on March 11th, 2020 (week 11 of 2020). As a result, different countries have implemented a variety of interventions to contain the virus. These included forced or recommended confinement, intended to reduce transmission by reducing contact rates between individuals[24]. These policies resulted in a dramatic change in human mobility, which in turn affected the traffic patterns and operations in telecommunication networks.

In this paper, we focus on the cellular network of O2 UK, and evaluate how the changes in people's mobility impacted this Mobile Network Operator (MNO) traffic patterns. The coronavirus outbreak reached the UK on January 31st 2020, when the first two (imported) cases with the respiratory disease COVID-19 were confirmed in York. On March 16th 2020 (week 12 of 2020), the government recommended all citizens to work from home, and on March 20th (also week 12), it implemented the closure of sporting events, schools, restaurants, bars and gyms.¹ On the 23rd of March (week 13), the government imposed a lockdown on the whole population, banning all non-essential travel and contact with people outside the home. London was particularly affected by the outbreak, with 27,000 positive cases at the end of May[19].

We collect and analyze statistics on the operational status of the MNO network, and our main findings are as follow.

By analyzing cellular network signalling information regarding users' device mobility activity, we observe a decrease of 50% in mobility (according to different mobility metrics) in the UK during the lockdown period. We find no correlation between this reduction in mobility and the number of confirmed COVID-19 cases, showing that only the enforced government order was effective in significantly reducing mobility. We observe this reduction is more significant in densely populated urban areas than in rural areas. We further note regional differences in how people relax the mobility restrictions, with an increase in mobility in London and West

¹We indicate the weeks of the year 2020 for the different relevant dates because our analysis in the following sections refers to them in the different graphs. Subsequent weeks also refer to the year 2020, but we omit to mention the year for brevity.

Yorkshire in weeks 18-19. For London, specifically, we observe that approximately 10% of the residents temporarily relocated during the lockdown.

We find that these mobility changes have immediate implications in traffic patterns of the cellular network. The downlink data traffic volume aggregated for all bearers (including conversational voice) decreased for all UK by up to 25% during the lockdown period. This correlates with the reduction in mobility we observe country-wide, which results in people likely relying more on the broadband residential Internet access to run download intensive applications such as video streaming. We also note a decrease in the radio cell load, with a reduction of approximately 15% across the UK after the stay-at-home order, which further corroborates the drop in cellular connectivity usage. This effect is consistent with the reported surge in traffic for residential ISPs. This decrease rewound the traffic load on the MNO infrastructure by one year, to levels similar to those of March 2019.

The total uplink data traffic volume, on the other hand, experienced little changes (between -7% and +1,5%) during lockdown. This is mainly due to the increase of 4G voice traffic (i.e., VoLTE) across the UK that peaked at 150% after lockdown compared to the national medial value before the pandemic, thus compensating the decrease in data traffic in the uplink. At the same time, we observe an increase of more than 100% in the downlink packet loss error rate for voice traffic on week 10 and 11. This was caused by excess of congestion in the interconnection infrastructure MNOs use to exchange voice traffic, whose capacity was exceeded during the steep surge of voice traffic. The error rate reverted its previous levels during the following weeks thanks to rapid response of the network operations.

Finally, we observe mobility changes have different impact on network usage in geodemographic area clusters. In densely populated urban areas, we observe a significantly higher decrease of mobile network usage (i.e., downlink and uplink traffic volumes, radio load and active users) than in rural areas. By looking into the case of London, we observe that this is likely due to geodemographics of the central districts (e.g., Eastern-Central(EC) and Western-Central(WC)), which include many seasonal residents (e.g., tourists), business and commercial areas.

The rest of the paper is organized as follows. In Section 2 we describe the measurement infrastructure, as well as the data feeds and the metrics used. In Section 3, we describe the evolution in mobility observed throughout the lockdown imposed in March and April. Then, in Section 4, we describe the changes observed in different parameters representative of the MNO’s network performance in the UK and specifically for the case of London. Next, we present the related work and conclude the paper.

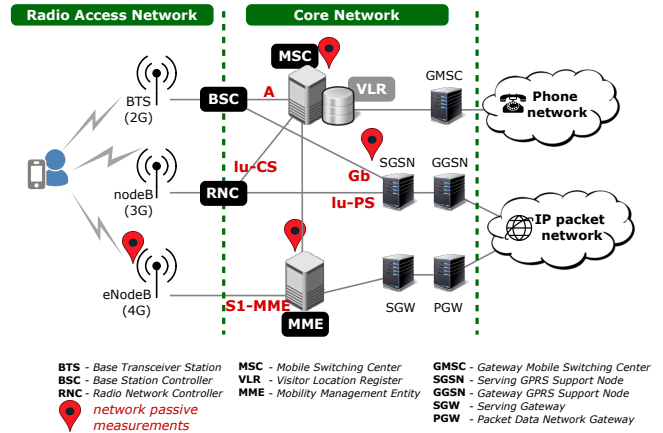


Figure 1: High-level architecture of the measurement infrastructure integrated in the cellular network.

2 DATASET

In this section, we describe the measurement infrastructure we leverage for collecting network data from a large commercial MNO in the UK (with more than 25% market share in the UK in 2019) . We detail the dataset we built and the metrics we use to capture the activity of smartphone devices.

2.1 Measurement Infrastructure

The cellular network we study supports 2G, 3G and 4G mobile communication technologies. In Figure 1, we illustrate a high-level schema of the MNO architecture. Such a network can be simplified to consist of three main domains: (i) the cellular device (in our case, the smartphone used as primary device by end-users), (ii) the Radio Access Network (RAN) and (iii) the Core Network (CN).

Our passive measurement approach relies on commercial solutions the MNO integrates within its infrastructure. The red pins in Figure 1 mark the network elements that we monitor, namely the Mobility Management Entity (MME), the Message Sequence Chart (MSC), the Serving GPRS Support Node (SGSN)/Serving Gateway (SGW), and the Cell Sites. We collect control plane information for both voice and data traffic from the total population of devices connected to the MNO’s radio network, as well as Key Performance Indicators (KPI) of cell sites.

Cell Sites. Cell sites (also called cell towers) are the sites where antennas and equipment of the RAN are placed. Every cell site hosts one or multiple antennas for one or more technologies (i.e., 2G, 3G, 4G), and includes multiple cells and sectors. For every cell site we have detailed information including location, radio technologies available, number of cells and radio sectors. We collect KPI for every radio sectors (e.g., radio load, average user throughput, traffic volume) that we aggregate at postcode level or larger granularity.

Radio Interfaces. We capture and process logs reporting on activities on the lu-PS (for 3G) and Gb (for 2G) interfaces, which carry events related to data packet transmissions and mobility management. Similarly, for the LTE networks, we capture the logs at the MME nodes on the S1 interface, reporting on mobility management events and bearer management, and user plane S1-UP interface for data and voice (over data) events. We also capture and process logs that report on events on the lu-CS (for 3G) and A (for 2G) interfaces, for 2G and 3G voice events. Note that the A interfaces also carries mobility management information. For the complete detailed specifications, we direct the reader to [1].

2.2 Data Feeds

From our measurement infrastructure, we capture various data feeds from the mobile network that we describe next. Note that these feeds are aggregated at postcode level or larger granularity.

General Signaling Dataset. As described in the previous section, we capture the activity of the users in the control plane for the different Radio Access Technologies (RATs) supported by the cellular provider. Specifically, for every RAT, the signalling dataset we collect includes the (control plane) signaling messages related to events triggered by the MNO’s subscribers, including Attach, Authentication, Session establishment, Dedicated bearer establishment and deletion, Tracking Area Update (TAU), ECM-IDLE mode transition, Service request, Handover and Detach. Each event we capture carries the anonymized user ID, Subscriber Identity Module (SIM) Mobile Country Code (MCC) and Mobile Network Code (MNC), Type Allocation Code (TAC) (the first 8 digits of the device IMEI, which are statically allocated to device vendors), the radio sector ID handling the communication, timestamp, and event result code (success / failure). Further, we aggregate this information at postcode level or larger granularity.

Devices Catalog. We consider a commercial database provided by Global System for Mobile communications (GSM) Association (GSMA). This catalog maps the device TAC to a set of device properties such as device manufacturer, brand and model name, operating system, radio bands supported, etc. With this information we are able to distinguish between smartphones (likely used as primary devices by the mobile users) and Machine-to-Machine (M2M) devices.

Radio Network Topology. To account for potential structural changes in the radio access network (e.g., new site deployments), we rely on a daily snapshot of the network topology. This includes metadata (location and configuration) and the status (active/inactive) of each cell tower.

Radio Network Performance. We rely on a commercial solution the MNO deploys to collect the radio network performance dataset. This dataset includes various KPI, including average cell throughout, average user throughput, average percentage of resources occupied, average number of users, total volume of data traffic uplink/downlink and total volume of conversational voice traffic. We collect this data hourly, and aggregate at postcode level or larger granularity.

UK Administrative and Geo-demographic Datasets. We use the National Statistics Postcode Lookup (NSPL) dataset for the UK as at February 2020 to group the postcode areas into Upper Tier Local Authority (UTLA). The NSPL is produced by ONS Geography, providing geographic support to the Office for National Statistics (ONS) and geographic services used by other organisations. Furthermore, we use the latest available Area Classification for Output Areas (2011 OAC) released in 2011, which represents a widely used public domain census-only geodemographic classifications in the UK [16]. The 2011 OAC dataset summarizes the social and physical structure of postcode areas using data from the 2011 UK Census, and is updated every 10 years when census is performed (the next one will be available in 2021).

2.3 Mobility Statistics

Since we are interested in analyzing mobility of people, we focus on their primary devices. We use the TAC database to filter only the devices that are smartphones (i.e., we drop M2M devices such as smart sensors). We are also able to separate the native users of the MNO, and drop the international inbound roamers from further analysis. Using the signalling data-set described above, we then associate each (anonymized) user to a radio tower throughout the time they are connected to the MNO’s network. Based on the radio network topology, we further attach to each radio tower its geographic location (postal code and approximate coordinates). With this, we then generate aggregated mobility statistics over six disjoint 4-hour bins of the day (e.g., 04:00AM - 08:00AM, 08:00AM - 12:00PM, 12:00PM - 04:00PM), and also over the entire day (i.e., 24 hours time window).

For each user, we determine the total duration of time they spend connected to every cell tower and select the top 20 towers. This allows us to identify all relevant places, as previous studies have shown than more than three quarters of people have between 3 to 6 important places, and in general no more than 8 [18, 21]. After doing this initial filtering, we obtain information regarding roughly 22 million native users aggregated at postcode level or larger granularity (e.g., UTLA or geodemographic cluster).

Mobility Metrics. From the aggregated mobility statistics, we focus on two metrics: entropy and radius of gyration. The

combination of both metrics gives a wide view of changes in mobility: while entropy measures the repeatability of movements, radius of gyration is an indication of the distance travelled. The two metrics are independent, one could have a high entropy with a reduced gyration, implying someone that moves in a reduced physical space almost randomly; or, on the contrary, have a low entropy with a large value of gyration, implying someone that moves over a large area but repeating the trajectories done. These metrics are computed over a day for each individual and aggregated to obtain an average value per day. Even if we compute these metrics per user at cell tower level, we aggregated them at postcode or larger granularity.

Entropy is a measure of the randomness of the movements of an individual, and as such, a metric for the predictability of movements [28]. From the variety of ways to calculate entropy in mobility [28], we implemented a temporal-uncorrelated entropy, that characterizes the heterogeneity of visitation patterns. Formally:

$$e = - \sum_{j=1}^N (p(j) \log(p(j))) \quad (1)$$

with $p(j)$ is the fraction of the time spent in the j^{th} visited cell tower (being a proxy for the probability for the user to be in that cell tower).

Radius of gyration is a key characteristic to model travelled distance [18], and measures how far from the center of mass the mass is located [2]. It is defined as the root mean squared distance between the set of cell towers and its center of masses. Formally:

$$g = \sqrt{\frac{1}{N} \sum_{j=1}^N (t_j \mathbf{l}_j - \mathbf{l}_{cm})^2} \quad (2)$$

where \mathbf{l}_j represents the location of the j^{th} visited cell tower, t_j represents the time spent in the j^{th} visited cell tower and \mathbf{l}_{cm} represents the location of the center of mass of the user's trajectory, calculated as $\mathbf{l}_{cm} = \frac{1}{N} \sum_{j=1}^N \{t_j \mathbf{l}_j\}$ and N the total number of towers visited.

Home Detection. For our analysis, locating the home postcode of the end-users is important when capturing their mobility patterns. Home Detection algorithms are a specific kind of a wider group of algorithms used to identify meaningful places from mobility information. The main idea consists in using some criteria to define time slots for home, work and other activities and then use the mobility information to identify these places [3, 15, 21, 26]. We estimate home location for each user at postcode granularity. For this, we use the cell tower to which the user connects more time during nighttime hours (12:00 PM through 8:00 AM) for at

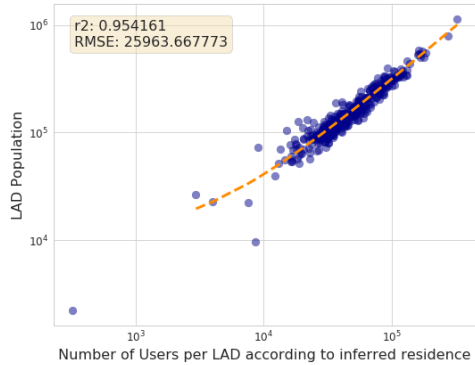


Figure 2: Comparison between inferred residential Local Authority District (LAD) population and the actual LAD population from census data.

least 14 days (not necessarily consecutive) during February 2020. With that filtering, we were able to determine the home postal code for approximately 16 million users.

An inherent limitation with our inference is that the estimation of the home location distribution is influenced by the market share of the MNO, and how it reflects the general population. In order to validate its reliability, we assigned all subjects to a Local Authority District LAD [13] according to our home estimations, and compared with values of population estimation from the Office for National Statistics (see Figure 2). The result shows a linear relationship ($r^2 = .955$), thus validating the representativity of the dataset. The values obtained are in line with the literature [27].

2.4 Network Performance Statistics

Using the general signalling dataset, we evaluate the average time the users spend connected to the different RAT cells. We find that 4G is the most popular RAT, with users spending on average 75% of the time per day connected to 4G cells. Thus, for the network performance statistics, we focus on 4G cells as they have the highest load out of the three RATs.

Based on the Radio Network Performance data feed, we generate network performance statistics at the 4G radio cell level. For each cell, we separate the following hourly performance metrics: the Uplink (UL) and the downlink (DL) data volume (the sum of all data transferred on all cell bearers corresponding to QoS Class Identifier (QCI) from 1 to 8 in each direction, UL and DL), average number of active DL users (users with active data transmission in the DL buffer), average radio load (as Transmission Time Interval (TTI) utilization, representing the number of active User Equipments (UEs) the LTE scheduler assigns per TTI), average user DL throughput (as the average throughput over all users active in the cell in one hour, considering all bearers corresponding to QCI from 1 to 8), and time (number of seconds)

with active data per cell. We also extract hourly metrics per cell specifically for conversational voice (separating only the bearers corresponding to QCI value 1), namely: voice traffic volume (total traffic with QCI equal to 1), average number of simultaneous voice active users, and the UL and DL average packet loss error rates.

For all the hourly metrics, we further aggregate them per day and extract the (hourly) median value per cell. This allows to capture one single value per metric per day, enabling further analysis with the daily mobility metrics. For each of the radio cell, we attach the location metadata information from the Radio Network Topology data feed. We further merge this (at the postcode level) with the UK Administrative and Geodemographic Datasets to append extra information such as the geodemographic cluster for each radio cell.

3 MOBILITY

In this section we present how the evolution of the pandemic and the social distancing measures impacted mobility by analyzing the change of the metrics detailed in Section 2.3. We capture the mobility metrics of users for 10 weeks (from week 10 to week 19 of 2020), which includes time before the SARS-CoV-2 pandemic was declared in the UK on March 11th 2020 (week 11), as well as during the government imposed measures to tackle the emergency. The total number of end-users whose data we aggregate for this study is approximately 16 million, and, unless otherwise specified, for all metrics we report for every day the percentage of change in the average daily value compared to average weekly value in week 9 (23 February - 1 March 2020).

3.1 National Mobility

We start our analysis by investigating the nation-wide time series for radius of gyration and the mobility entropy in Fig 3.

The average gyration evolution (Figure 3a) shows the reduction in the total area that users cover in their daily routines in reference to the average value over week 9². We note that people started implementing social distancing recommendations even before lockdown was enforced, with a decrease of 20% in gyration in week 12. With the government imposing the nation-wide lockdown in week 13, we also observe a steep decrease in gyration, with a drop of 50% towards the end of week 13 compared to the usual value from week 9. Mobility entropy per user follows a similar trend (Figure 3b). All metrics show a steep decrease in weeks 13-14, following the "stay-at-home" being enforced. In the following weeks, we note a slight relaxation, with mobility marginally increasing. It is worth noticing that, the reduction of entropy is smaller than the reduction of gyration. This

²Note that during the weekdays of week 9, the gyration is larger while during the weekend is smaller, yielding the aforementioned average.

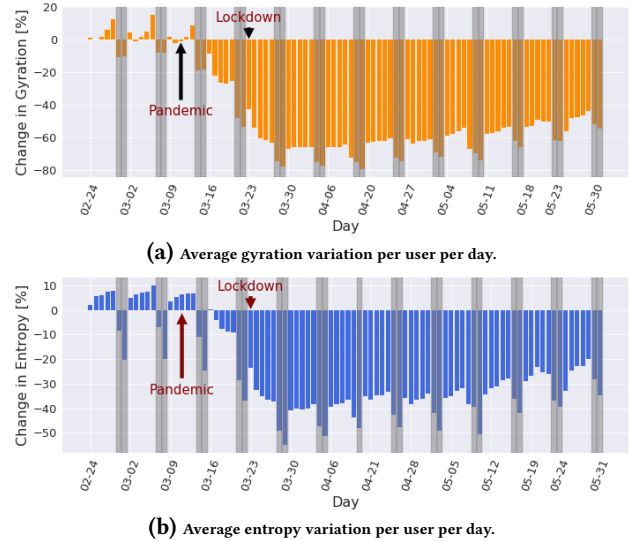


Figure 3: Percentage of change in the average value per user for radius of gyration and entropy, compared to their average value in week 9. Shaded bars correspond to weekends.




Figure 4: Entropy variation (from week 9 to week 18) vs cumulative number SARS-CoV-2 infections per day. Each point represents a different day. Colors encode the different days of the week (yellow shows week-end).

indicates that people, besides moving significantly less, tend to move close to their home location.

Figure 4 captures the correlation between the average mobility entropy per user and the nation-wide cumulative number of lab-confirmed SARS-CoV-2 cases, as reported by Public Health England[19]. Each point in the scatterplot represent a different day; we capture the interval between February 23rd until May 4th, 2020. We note that the mobility reduction is not impacted by the number of reported cases, i.e., there is not a correlation between number of cases

and mobility, but rather mobility is impacted by public announcement and lockdown measures. The decrease in the entropy starts only after the pandemic is declared (vertical red line in Figure 4, coinciding with 1,000 confirmed cases), and becomes significant after the lockdown.

Takeaway: Mobility metrics (gyration and entropy) show a steep decrease in people’s mobility in weeks 13-14, following the "stay-at-home" order being enforced. We find no correlation between this reduction in mobility and the number of confirmed COVID-19 cases, showing that only the enforced government order was effective in significantly reducing mobility. We also notice mobility slightly increases from week 15 despite the lockdown still being enforced.

3.2 Regional Mobility

We now focus our analysis on five different regions in order to observe potential geo-spatial difference in the mobility pattern changes. We select the regions that are best represented in our dataset with more than 500,000 users, namely Inner London (700k users), Outer London (1,1 million users), Greater Manchester (700k users), West Midland (600k users) and West Yorkshire (500k users). For each region, we capture the variation of the two mobility metrics (Figure 5) in reference to the nation-wide average value of the metric. The evolution of the metrics shows clearly the impact of the stay-at-home measures in every region, with a sharp decrease in weeks 13-14 in the values of all metrics. We note that for London (both Inner and Outer London), reference values for gyration are below national average (20% below the average for each corresponding week– see Figure 5a), while the reference value for mobility entropy is higher (20% above the average, see Figure 5b). This shows that within London, in general, people move more randomly and with less predictable mobility pattern, but cover smaller areas than the national average. This analysis also bring to our attention the regional differences in how people relax the mobility restrictions, with an increase in mobility in London and West Yorkshire in weeks 18-19. This is not the case for the regions of Greater Manchester and West Midlands, where mobility is consistently low after week 13. Finally, metrics distributions have little variance in all regions, and all percentiles are close to the median, following similar trends.

Takeaway: The impact of the lockdown is consistent over different regions in the UK, showing that people respected the lockdown, regardless where they live. We do, however, find regional differences in how people relax the mobility restrictions, with an increase in mobility in London and West Yorkshire in weeks 18-19. This is not the case for the regions of Greater Manchester and West Midlands, where mobility is consistently low after week 13.




Figure 5: Variation in the average gyration and entropy per region, compared to the national average during week 9.

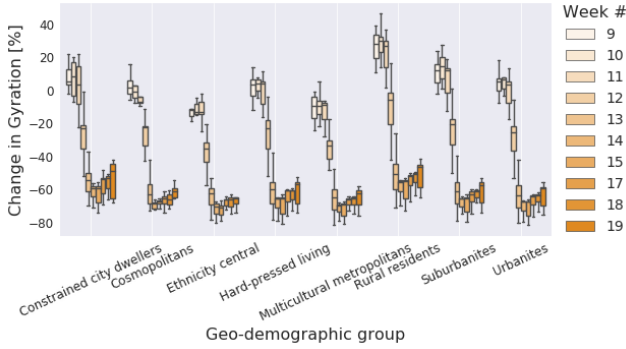
3.3 Geodemographic Mobility

To capture the changes in users’ mobility in correlation with different social and demographic characteristics of their residence area, we study the variation of mobility metrics across different geo-spatial clusters as defined by the UK Office for National Statistics (ONS)[14]. This consists of eight groups meant to be illustrative of the characteristics of areas in terms of their demographic structure, household composition, housing, socio-economic characteristics and employment patterns. Each of the eight categories provides the most generic descriptions of the corresponding population group in the UK (Table 1).

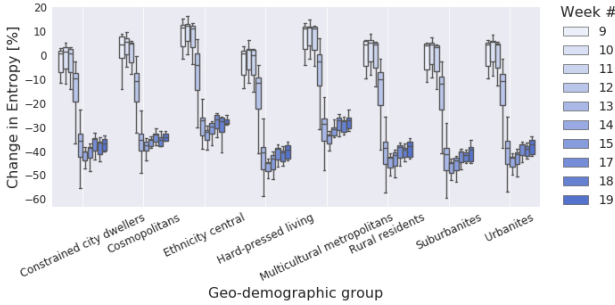
We break down the mobility metrics variations per week in each of the above-mentioned area groups (Figure 6). We find that mobility in rural areas is normally higher than the nation average and people usually cover wider areas in their daily movement (weeks 9-11 in Figure 6a). Contrariwise, in highly populated urban areas (such as Cosmopolitans or Ethnicity central), the population covers smaller areas (weeks 9-11 in Figure 6a), but the predictability of their mobility patterns is lower (higher entropy in weeks 9-11 in Figure 6b). All the different socio-economic groups present the same significant drop in mobility, with a transition period in week 12 and a steep drop from week 13. Gyration decreases across all groups by more than 50% of the national average in week

Table 1: Geodemographic clusters (2011 OAC).

Name	Definition
Rural Residents	Rural areas, low density, older and educated population
Cosmopolitans	Densely populated urban areas, high ethnic integration, young adults and students
Ethnicity Central	Denser central areas of London, non-white ethnic groups, young adults
Multicultural Metropolitans	Urban areas in transition between centres and suburbia, high ethnic mix
Urbanites	Urban areas mainly in southern England, average ethnic mix, low unemployment
Suburbanites	Population above retirement age and parents with school age children, low unemployment
Constrained City Dwellers	Densely populated areas, single/divorced population, higher level of unemployment
Hard-pressed Living	Urban surroundings (northern England/southern Wales), higher rates of unemployment



(a) Gyration.



(b) Mobility entropy.

Figure 6: Variation in average gyration and entropy per geodemographic cluster, compared to the national average in week 9.

9. As it was observed in the previous section, the ethnicity central group (which would basically correspond to Inner London), the reduction of gyration is the highest of all the groups but their reduction in entropy is the smallest, implying that they have reduced their area of mobility the most, but within that area their movement is more random. This is probably related to the high density of commercial services in central London. Finally, we observe that metrics’ distribution has little variance in all geodemographic clusters.

Takeaway: All users in different geodemographic clusters present the same significant drop in mobility, with a transition period in week 12 and a steep drop from week 13.

3.4 Temporary Relocation from London

In this section, we measure how many people from London decided to move elsewhere during (part of) the analyzed time frame. To evaluate the temporary relocation of Inner London residents to secondary locations, we generate its mobility matrix at a county level (see Figure 7). For each Inner London resident (obtained using the method described in 2.3), we check the top 20 locations (at county level) that they visit during each day (cf. Section 2.3). If none of the visited locations during a day matches their home county (Inner London in this case) we are able to identify relocations, i.e. people with residence in London that are elsewhere. Figure 7 shows the variation in the number of Inner London residents who are present in the different counties (we capture a different county per row) per day, compared to the average number we identified in week 9, before the lockdown.

By looking at the Inner London line we observe a permanent 10% decrease in the number of Inner London residents who actually are present in their area of residence from week 13 onward (after the lockdown is imposed). Contributions to the decrease in the number of Inner London residents who remain in their inferred homes during the lockdown period may include, for example, students who left campuses in London after schools closing on the 19th of March, or long-term tourist leaving Inner London center areas, or London residents who decided to spend the lockdown in their second residences. In particular, we observe an increase in the number of people from London who relocated to the Hampshire area during most of the duration of the lockdown.

Further, we note that mobility patterns of Inner London residents changed significantly after the stay-at-home order was imposed. Specifically, before the stay-at-home recommendations, we observe an increase of Londoners spending the weekend in other counties in the UK. This pattern disappears starting from weeks 11 and 12, concurrent with the recommendation of social distancing. We capture a large variation in the number of people travelling from Inner London to outside areas such as East Sussex on the 21st-22nd of March, just prior to the stay-at-home order. We also observe




Figure 7: Mobility matrix of users with residence in Inner London towards other counties in the UK. Each cell shows the variation per day (on the x axis) in the number of resident from Inner London active in the county marked in the y axis compared to the median value over week 9. We include in our mobility matrix the top 10 counties in terms of receiving inbound residents from Inner London according to the average in week 9.

an additional increase of Londoners going to Hampshire for the weekend by the end of April (also in Kent but less so), consistent with the overall (slight) increase observed at the end of the period in the overall mobility metrics.

Takeaway: We find a sustained 10% decrease in the number of Inner London residents who actually are present in their area of residence from week 13 onward (after the lockdown is imposed). We capture a large variation in the number of people travelling from Inner London to outside areas such as East Sussex on the 21st-22nd of March, just prior to the stay-at-home order being enforced.

4 MOBILE NETWORK PERFORMANCE

In this section, we analyze mobile network performance indicators to investigate how changes in people’s mobility induce shifts in the usual traffic patterns we observe in the radio network, and further impact the mobile network operations.

4.1 Data Traffic Evolution

We investigate the dynamics of the network performance metrics we defined in Section 2.4 across the UK, aggregating all bearers per cell, including voice and data traffic (i.e., we aggregate all bearers for QCI 1 to 8). Our analysis spans a period of 10 weeks, from 23rd of February until the 10th of May. Unless otherwise stated, for all metrics we present the delta variation percentage of each metric from its median value over week 9.

Figure 8: *Downlink Data Volume* captures the variation of the downlink data volume during the weeks we mark on the x-axis. Throughout the UK, we note an increase of 8% in the average data volume per cell in week 10. This is followed by a steady decrease, with the lowest registered volume in week 17 (-24% compared to week 9). Our observation is consistent with people’s mobility decrease after the stay-at-home order (Section 3). We conjecture that after lockdown people relied less on the cellular network for data connectivity (e.g., using home WiFi connectivity instead), thus contributing to the surge of traffic reported by residential ISPs.³

³See <https://www.cnn.com/2020/03/27/coronavirus-can-the-internet-handle-unprecedented-surge-in-traffic.html> for instance.



Figure 8: MNO performance characterization, using metrics for all data traffic. Each plot corresponds to a metric, as we indicate in the title of the plot. We show the median values for the delta variation percentage for each metric over one week (each point on the x-axis corresponds to a week #), using as a baseline for comparison the median value in week 9. Each line corresponds to different geo-regions within the UK, including "UK - all regions", which accounts for the entire country.

Indeed, we observe in Figure 8: *Uplink Data Volume* a significant reduction of active downlink users per hour per cell compared to week 9 throughout the entire UK (all regions), with a minimum number of active users in week 19 (-28.6% users compared to week 9). Note that, despite we observe a slight recovery of mobility in week 14 (cf. Section 3), this is not sufficient to have an impact on the downlink data volume, the number of active users, or the radio load per cell, which do not increase.

We find a small reduction in the downlink user throughput in the 10 weeks we analyze (Figure 8: *User Downlink*

Throughput)), with the lowest value being a 10% drop compared to week 9. This follows the same trend as the decrease in downlink traffic volume. This is unexpected, as in presence of less traffic (cf. Figure 8: *Downlink Data Volume*) and overall lower active downlink users per cell (cf. Figure 8: *Downlink Active Users*), the per user throughput should increase. We also observe that congestion in the radio network is not the cause of this decrease because the radio load also decreases (Figure 8: *Cell Resource Utilization* shows a 15.1% radio load decrease in week 16 compared to week 9). We conjecture the per-user throughput is application limited and not network limited during the weeks of the analysis. This is in line with the choice of many content providers to reduce content quality during the pandemic⁴ and by the shift of more data-intensive applications to wifi access.

We also investigate the variation of the total uplink data volume per cell in the UK (Figure 8: *Uplink Data Volume*). We find overall modest changes for this metric, with, less than 5% decrease in the median daily volumes in week 18. This suggests that applications making intensive use of the downlink (e.g., video streaming) suffered a significantly higher traffic reduction than applications with symmetric uplink/downlink usage (e.g., audio/video conferences, or voice traffic). In particular, as we present in the next section, we observe that voice traffic (corresponding to QCI=1) has significantly increased during the lockdown period.

Altogether, because the downlink data volume is one order of magnitude larger than the uplink data volume, the volume of data traffic carried by the network experienced an overall reduction of roughly 20%, reverting to 2019 traffic volume (when the MNO had less customers and applications were less bandwidth hungry).

Finally, we note that metrics' distribution across cells does not significantly change across weeks with respect to the reference week 9. A wide distribution is common in large operational environments, where cells support different RATs types across the country (e.g. LTE, LTE-A), different number of customers each cell serves, different radio coverage, etc. The only exception we note is in the 90th percentile in the downlink active user per hour per cell (not picture for readability reasons), which slightly reduces during the lockdown phase, thus showing the nation-wide impact on the users connectivity habits.

Takeaway: We find an overall reduction (-24% in week 17 compared to week 9) in downlink data traffic volume throughout the UK during lockdown. Unintuitively, we also note a decrease in the average user throughput. Though a decrease in radio cell resource utilization and in overall traffic

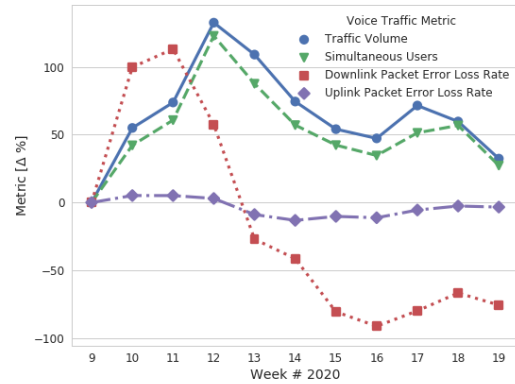


Figure 9: Analysis of 4G voice traffic (corresponding to QCI=1) patterns in the UK. We show the median values for the delta variation percentage, which we generate compared to the national average during week 9 for all the UK, over a period of 10 weeks (each point corresponds to a week # we mark on the x-axis), from the end of February until mid-May. Each line corresponds to different metric.

may have hinted to the opposite, we find the throughput to drop by at most 10%, likely due to application throttling. We also observe uplink traffic volume has little variation.

4.2 Voice Traffic

We now analyze the changes in voice traffic patterns, to further drill down on the potential dynamics that impact the total volumes of data traffic we observe. We filter 4G traffic marked with QCI equal to 1 (conversational traffic), and analyze the performance metrics we previously explained in Section 2.4. Figure 9 shows the evolution of these metric over the period of 10 weeks we analyze in March - May 2020, for all the UK compared to the average of week 9.

We note that in week 12 there is a spike of 140% in the median value of voice traffic volume (Figure 9: *Traffic Volume*) and a significant increase of its top 90 percentile value. This corresponds to a predicted seven years of growth in voice traffic in the operator's network, which the MNO had to accommodate in the space of few days (week 12 in Figure 9: *Traffic Volume*). Though striking, we find that the increase in voice traffic is not enough to compensate the decrease in other traffic classes, resulting in the overall decrease in traffic presented in the previous section. The surge in voice traffic volume is consistent with a spike in the average number of simultaneous voice users per cell (Figure 9: *Simultaneous Users*). This aligns with the time when the pandemic was declared (11th of March), and the lockdown restrictions started in the UK (23rd of March, week 13). This is an anticipated result of the lockdown and the subsequent changes in people's mobility, since more are likely to carry on voice conversations in these circumstances.

⁴<https://www.theverge.com/2020/3/20/21187930/youtube-reduces-streaming-quality-european-union-coronavirus-bandwidth-internet-traffic>

We report in Figure 9: *Uplink Packet Error Loss Rate* and Figure 9: *Downlink Packet Error Loss Rate* the daily average uplink and downlink packet loss error rate per cell respectively. While the uplink packet loss decreases during the pandemic period, we observe a significant increase in the downlink packet loss during weeks 10-11-12, after which it decreases below normal values. This is again unexpected, as the radio network is less congested and there is less downlink traffic (see previous section). The issue during these weeks was, in fact, that the unexpected surge in the amount of voice traffic exceeded the capacity of the inter-MNO interconnection infrastructure. The rapid response of the network operators and service providers quickly restored the DL error below the normal values.

Takeaway: We observe a spike of 140% in the median value of voice traffic volume around lockdown. Although the MNO was dimensioned to accommodate the traffic, this unexpected surge exceeded the capacity of the interconnection infrastructure between MNOs. The rapid response of the network operators and service providers restored the network status back to normal quickly after week 13.

4.3 High Density Areas

To capture regional traffic dynamics, we now zoom into the network performance of five counties with high density of users in our dataset. Specifically, we analyse the network performance in Outer London, Inner London, Greater Manchester, West Midlands and West Yorkshire.

We analyze four performance metrics, namely the daily uplink/downlink data traffic volume per cell, the number of active downlink users per cell, and the downlink average user throughput. We complement this analysis by also showing the total number of users connected (both active and idle), and the cell resource utilization. We report on the delta variation percentages of each metric per week with respect to the nation-wide average for each metric observed in week 9 (Figure 8, plot lines marked with the region names we mention above). We capture in our analysis 10 weeks, from the 23rd of February until the 10th of May 2020. We observe that while the overall trends in the five counties are similar to those over the entire UK (Figure 8, plot lines marked *UK*), the intensity of the different trends varies significantly from one region to another.

Regarding the downlink data traffic volume (Figure 8: *Downlink Data Volume*), we observe in all regions a mild increase in week 10 (ranging between 9% and 17% depending on the region), followed by a decrease in the following weeks reaching the minimum in weeks 17 and 18. However, the decrease is much larger in Inner London (with a 41% decrease) than in the other regions, which decrease by about

20%. Further, in outer London, there is the smallest decrease, with only 15%.

We also see in Figure 8: *Downlink Active Users* that during lockdown the number of active downlink users per cell significantly decreases in Inner London (e.g., -40% in week 15), while it is almost constant and even slightly higher in Outer London. These different trends between commercial/business and residential areas are also visible for other regions as well (i.e., West Midlands/Greater Manchester more business compared to Yorkshire more residential).

An additional contributing factor to the decrease of traffic and users in Inner London is the temporary relocation of people living in Inner London to other counties outside the city. In Section 3 we saw that approximately 10% of inferred residents were not signalling against radio cells covering London after lockdown, thus hinting a potential relocation and partially explaining the reduction in the data traffic volume. We make a detailed analysis of the different districts within Inner London in Section 5.

Regarding the uplink data volume (Figure 8: *Uplink Data Volume*), we observe different trends across regions. The uplink traffic volume grows for all regions in week 10. This growth is followed by a decrease in Inner London and West Midlands, and by a slight increase or steadier behavior in Outer London and West Yorkshire. The highest contrast is again in Inner/Outer London. In week 14, in Inner London, the UL data traffic volume decreased -22%; while in the Outer London region it increased by 17%. In the other regions, the effects are similar, but with a lower intensity.

We conjecture these differences correlate with the different geodemographic dynamics within these regions. For example, Inner London includes many business, touristic, commercial and recreational areas, which remained largely empty during lockdown, potentially explaining the steep decrease in the downlink data volume in the area. To capture these correlations, in the following section we further study the network performance on areas with the same geodemographic profiles (Table 1). Thus, we change from the current perspective of network performance within administrative boundaries, to capturing performance metrics across clusters of areas with the different geodemographic profiles.

Takeaway: We observe that while the overall trends in five counties are similar to those in the entire UK, their intensity varies significantly from one region to another. We conjecture these differences correlate with the different geodemographic dynamics within these regions.

4.4 Geodemographic Clusters

In this section, we incorporate into our analysis of network performance during COVID-19 demographic dynamics. Geodemographic classifications provide summary indicators of the

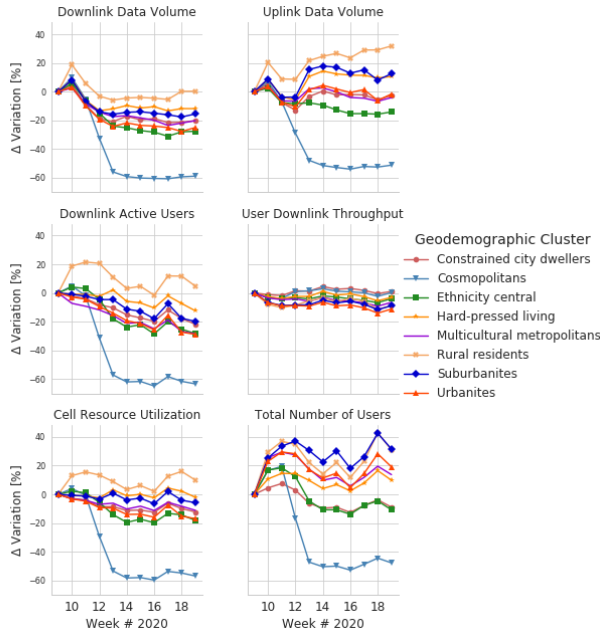


Figure 10: Focused analysis of network performance metrics on geodemographic clusters (OAC 2011 dataset). We show the median variation per cluster (see Table 1), where each line corresponds to a different cluster.

social, economic, demographic, and built characteristics of small areas. To this end, we use the geodemographic clusters (Table 1) that the ONS labeled to each postcode.⁵ We group the postcode areas in the corresponding geodemographic clusters and extract the network performance statistics (Section 2.4) per cluster. This analysis allows us to understand potential network performance differences between groups of areas with distinct mobility patterns across the entire UK (e.g., rural residents areas vs. cosmopolitan areas, see Figure 6), as we already observed different behavior for different areas in the previous section.

We find that the behaviour for most areas, with the exception of "Rural residents" and "Cosmopolitan" areas, exhibits similar trends than the ones observed in UK nation wide, with a decrease in downlink traffic and a stable uplink traffic volume. In areas within the "Rural residents" cluster, the average volume of downlink data traffic per cell remains largely stable after the lockdown period.

Figure 10 shows that the average volume of downlink data traffic per cell in areas with majority of "Cosmopolitan" residents decreased dramatically after week 13, coinciding with the stay-at-home order. The areas within the "Cosmopolitans" cluster correspond to densely populated urban areas,

⁵National Statistics Postcode Lookup (NSPL) for the United Kingdom as at February 2020. <https://geoportal.statistics.gov.uk/datasets/national-statistics-postcode-lookup-february-2020>

with high ethnic integration, where young adults and students reside (see Table 1). These areas include the neuralgic centers for business, commercial, recreational, educational activities throughout the country. The pattern we observe is different from the rest of the areas. This is likely to be caused by a dramatic reduction of people in these places, since are areas with a lower number of residents but with a lot of people visiting them due to work and recreational purposes. This is consistent with the observations we made in the previous section for Inner London, where approximately 45% of postcode areas cluster within the "Cosmopolitans" geodemographic group, while 50% of postcode cluster in the "Ethnicity Central" group.

To further verify this correlation, in Figure 10 we show the total number of users connected to the network, and observe a sharp decrease of up to -50% in the "Cosmopolitan" areas. We calculate the correlation between the total number of users (both idle and active users) and the downlink data volume for the geodemographic clusters. We find a high correlation for the "Cosmopolitan" (+0.973) and for "Ethnicity central" (+0.816) areas, low correlation for "Rural residents" areas (0.299), and a negative correlation for "Suburbanites" (-0.466), where the increase of users resulted in a decrease of downlink data volume (Figure 8). In these areas, the drop in downlink data volume is likely impacted by other factors (e.g., offloading to WiFi, changes in applications used).

Takeaway: Areas within "Rural residents" and "Cosmopolitan" geodemographic clusters deviate from the trends we observed in UK nation wide. In the former, the average volume of downlink data traffic per cell remains largely stable after the lockdown. Conversely, in the latter we observe a sharp decrease of up to -50% in the total number of users connected after lockdown, which led to a dramatic decrease in the downlink data traffic volume.

5 LONDON-CENTRIC ANALYSIS

In this section, we focus our analysis on the London area (i.e., the Inner London county). London is particularly interesting because of the large concentration of users in a single area, high mobility, diversity in user's profiles, and also because, unfortunately, the pandemic was particularly severe in the city. Our goal is to dissect how the changes in the mobility patterns of the population residing in this region (approximately 700,000 users in our dataset) resulted in changes in their habits of using the mobile network service.

5.1 Network Performance

Figure 11 shows the network performance evolution based on data volume metrics, compared to median values over week 9 within London. We break down the analysis for each of the Postcode Districts within London, and we only depict

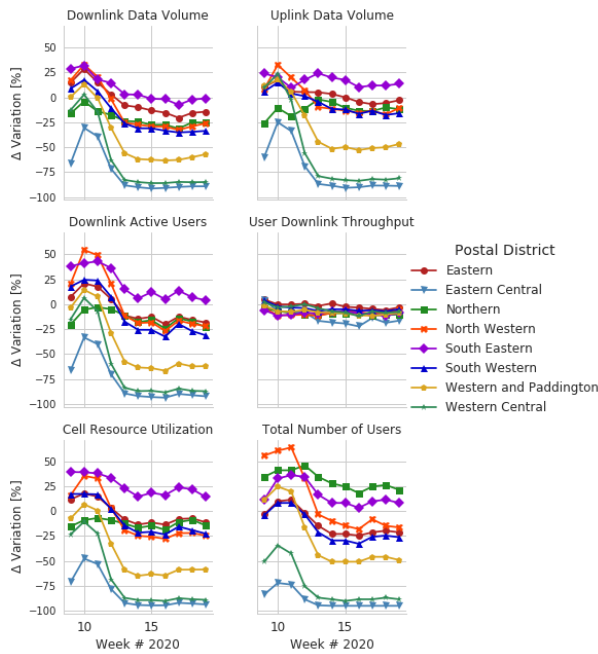


Figure 11: Network performance metrics for Inner London. Each line corresponds to a Postal District within London. We show the median variation percentage (on the y axis) per postal district, over a period of 10 weeks (on the x axis).




Figure 12: Network performance metrics per geodemographic clusters. We show the median variation percentage (on the y axis) per cluster (see Table 1), where each line corresponds to a different cluster, over a period of 10 weeks (on the x axis).

the median values over each week (for readability reasons). We note a general decrease in the total volume of downlink traffic (Figure 11: *Downlink Data Volume*), as well as in the volume of uplink traffic (Figure 11: *Uplink Data Volume*). The decrease is particularly large for Eastern Central (EC) and Western Central (WC). In the WC district, the uplink and the downlink traffic volumes exhibited a decrease of over 80% between weeks 14 and 19, while the EC district experienced a similar trend, with a decrease of more than 70% compared to week 9. These two districts cover the central area of London, where the density of population is also lower (i.e., approximately 30,000 residents in EC compared to 400,000 residents in SW) and includes many seasonal residents (e.g., tourists), business and commercial areas. The decreasing volume of data is likely due to a major reduction of people in the area, which is consistent with the decrease we observed in the average number of radio users per sector (see Figure 11: *Total Number of Users*). We note a similar trend both in the percentage of cell resources occupancy (Figure 11: *Cell Resource Utilization*) and in the average number of users per cell (Figure 11: *Downlink Active Users*).

We find that the trends in the Northern (N) district detach from the rest of the districts. In particular, we note an increase in the number of downlink users (varying between 10% and 23% in weeks 10 to 14), while the volume of downlink traffic keeps stable (unlike the other postcodes, where we note a decrease). This hints towards potential hot spots in the mobile network moving within London from central areas (EC, WC) to the north (N).

Takeaway: Performance in the central postal districts of London (EC and WC) differ from the rest, with decreases of over 70% in the downlink data traffic. This is due to a major reduction of people in the area, which is consistent with the decrease we observed in the average number of total radio users in those sectors.

5.2 Geodemographic Clusters

We further dissect the performance in London in relation to the geodemographic clusters. We find that only three clusters map to the area of London, each with different patterns. Figure 12 shows the evolution of four performance metric over the period of analysis (median values of delta variation percentage per week compared to corresponding median values in week 9 in London). We find again a sharper decrease in the areas marked as "Cosmopolitans", largely matching the behaviour we observed in the EC and WC districts, due to the severe reduction of users in these areas. We find that all areas follow the same trends for the user downlink throughput, which is consistent with the UK-wide observation (Figure 8: *User Downlink Throughput*). However,

for the uplink/downlink traffic volume, we note distinct patterns for the three clusters, showing distinct usage patterns across the different areas. We observe a sharp decrease in the traffic volumes in the radio cells in Cosmopolitans cluster areas (decrease of more than 50% in week 13 in the uplink and the downlink data volume), while cells in the "Multicultural clusters" cluster areas show an increase the mobile traffic volume (e.g., 40% increase in the uplink data volume). These correlate with the patterns we observe for the number of active downlink users in the same areas (i.e., a more than 20% increase in week 13).

Takeaway: We find again a sharper decrease in the uplink/downlink data volume in the areas that cluster as "Cosmopolitans", largely matching the behaviour we observed in the EC and WC districts, due to the severe reduction of users in these areas. Cells in the "Multicultural clusters" cluster areas show an increase the mobile traffic volume largely due to an increase in the number of active users.

6 RELATED WORK

Even though the spread of SARS-CoV-2 started only a few months ago, there are already in the literature different studies that present efforts for understanding human mobility during the interventions. Apart from cell phone traces, as used in our study and in others [22], there are other datasets that other studies investigate. Google has available anonymized and aggregated private counts of visits to places in different categories [30][20][23]; also, Facebook through its Disease Prevention maps has provided aggregated data of spatial distribution of individuals [25][17]; Cuebiq has provided GDPR-compliant information to model mobility in Boston [4] and Italy [10]; and Baidu has provided in-flow and out-flow indexes for Chinese cities [8]. In general, all the previous studies highlight the drastic changes in mobility and evaluate the impact of the interventions in the spread of the virus. In the particular case of the UK, TomTom [29] observed a reduction between 60% and 80% in the number of trips during lockdown, which is consistent with the reduction we observe from the MNO. Also, [6], by combining cell phone traces and Facebook data to identify number of journeys, concluded that mobility began to decrease around one week before lockdown was enforced, but that the sharpest drop was after the closure, similarly to our findings using different metrics.

The overall topic of the impact of COVID-19 in network traffic has also been subject to numerous reports. For example, Comcast, which provides mobile connectivity using wifi while at home but leverages on Verizon LTE while away, has observed an increase of 39% in wifi and a decrease of 17% of mobile LTE traffic. Regarding the increase of the different types of applications, there was a 215-285% increase in VoIP

and videoconferencing traffic, 30-40% in VPN traffic and a 20-40% in streaming and web video consumption.[9]. In the context of network traffic and distance learning, of critical relevance during COVID-19, the work by Favale et al. [11] found that incoming traffic drastically decreased, while outgoing traffic more than doubled to support online learning. By combining datasets from one ISP, three IXPs, and one metropolitan educational network, Feldmann et. al. [12] show that the fixed internet infrastructure was able to sustain the 15-20% increase in traffic that happened rapidly during a short window of one week.

The Cellular Telecommunications Industry Association (CTIA) provides daily reports on the variations of voice and data minutes observed by ATT, T-Mobile, U.S. Cellular and Verizon. With respect to the period between the 23 of February and the 16 of March, it was observed a surge of the upper range of voice minutes used of up to +24% in the last week of march (week 13) and an increase of about +20% in the upper range of data minutes[5]. While the surge reported in voice traffic is within the ranges observed in this paper, we do not observe the surge in data traffic reported by the main MNOs in the U.S. This may be due to different prevalence of flat fee contracts in the U.S. operators and the U.K. MNO we study. Similarly, BT in the U.K., on March the 20th 2020 [7], reported a decrease of 5% in mobile data traffic which is consistent with the decrease described in this paper for week 12. It was also observed a surge in voice traffic, and as a result BT encouraged users to use landline services, especially for long conference calls.

7 CONCLUSIONS

The lockdown measures implemented by the UK government to limit the spread of SARS-CoV-2 virus drastically changed user mobility patterns and network traffic. The results of this paper present an analysis of the changes in mobility and their impact on the cellular network traffic. Regarding mobility, we observed an overall decrease of 50%, with non-uniform changes across different geographical areas and social backgrounds, which confirm previous findings using other datasets and metrics [29][6]. These variations translated into a surge of voice traffic (up to 150%) accompanied by an overall decrease of download traffic (20%), especially in densely populated urban areas (-60%), and an increase of uplink traffic in suburbanities (10%). Despite significant pattern changes, the MNO was able to provide service maintaining quality standards: the radio load was below common values and per user throughput was likely application limited. We identified one issue in voice traffic packet loss due to excess of congestion in the interconnection infrastructure MNOs use to exchange voice traffic, which was fixed by network operation teams.

A ETHICAL CONSIDERATIONS

The data collection and retention at network middle-boxes and elements are in accordance with the terms and conditions of the MNO and the local regulations. All datasets used in this work are covered by NDAs prohibiting any re-sharing with 3rd parties even for research purposes. Further, raw data has been reviewed and validated by the operator with respect to GPDR compliance (e.g., no identifier can be associated to person), and data processing only extracts aggregated user information at postcode level. No personal and/or contract information was available for this study and none of the authors of this paper participated in the extraction and/or encryption of the raw data.

B ACKNOWLEDGEMENTS

We thank the IMC anonymous reviewers and our shepherd for their helpful comments and guidance. The work of Andra Lutu was supported by the EC H2020 Marie Curie Individual Fellowship 841315 (DICE).

REFERENCES

- [1] 3GPP. [n. d.]. 3rd Generation Partnership Project (3GPP). ([n. d.]). Retrieved May 14, 2020 from <https://www.3gpp.org>
- [2] MA Abramowicz, JC Miller, and Z Stuchlík. 1993. Concept of radius of gyration in general relativity. *Physical Review D* 47, 4 (1993), 1440.
- [3] Rein Ahas, Siiri Silm, Olle Järvi, Erki Saluveer, and Margus Tiru. 2010. Using mobile positioning data to model locations meaningful to users of mobile phones. *Journal of urban technology* 17, 1 (2010), 3–27.
- [4] Alberto Aleta, David Martin-Corral, Ana Pastore y Piontti, Marco Ajelli, Maria Litvinova, Matteo Chinazzi, Natalie E Dean, M Elizabeth Halloran, Ira M Longini, Stefano Merler, et al. 2020. Modeling the impact of social distancing, testing, contact tracing and household quarantine on second-wave scenarios of the COVID-19 epidemic. *medRxiv* (2020).
- [5] Cellular Telecommunications Industry Association. 2020. The Wireless Industry responds to COVID-19. (2020). <https://www.ctia.org/homepage/covid-19#network-performance>
- [6] Kylie E C Ainslie Oliver Eales Constanze Ciavarella Sangeeta Bhatia Sarah Hayes Marc Baguelin Adhiratha Boonyasiri Nicholas F. Brazeau Gina Cuomo-Dannenburg Richard G FitzJohn Katy Gaythorpe William Green Natsuko Imai Thomas A Mellan Swapnil Mishra Pierre Nouvellet H Juliette T Unwin Robert Verity Michaela Vollmer Charles Whittaker Neil Ferguson Christl A. Donnelly Benjamin Jeffrey, Caroline E Walters and Steven Riley. 2020. Report 24 - Anonymised aggregated crowd level mobility data from mobile phones suggests initial compliance with COVID19 social distancing interventions was high geographically consistent across UK. (2020). <https://www.imperial.ac.uk/media/imperial-college/medicine/mrc-gida/2020-05-29-COVID19-Report-24.pdf>
- [7] BT.com. 2020. The facts about our network and Coronavirus. (2020). <https://newsroom.bt.com/the-facts-about-our-network-and-coronavirus/>
- [8] Matteo Chinazzi, Jessica T Davis, Marco Ajelli, Corrado Gioannini, Maria Litvinova, Stefano Merler, Ana Pastore y Piontti, Kunpeng Mu, Luca Rossi, Kaiyuan Sun, et al. 2020. The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak. *Science* 368, 6489 (2020), 395–400.
- [9] COMCAST. 2020. COVID-19 Network Update. (2020). <https://corporate.comcast.com/covid-19/network>
- [10] Laetitia Gauvin Filippo Privitera Ciro Cattuto Michele Tizzoni Emanuele Pepe, Paolo Bajardi. 2020. COVID-19 outbreak response: first assessment of mobility changes in Italy following lockdown. (2020). <https://covid19mm.github.io/in-progress/2020/03/13/first-report-assessment.html>
- [11] Thomas Favale, Francesca Soro, Martino Trevisan, Idilio Drago, and Marco Mellia. 2020. Campus Traffic and e-Learning during COVID-19 Pandemic. *Computer Networks* (2020), 107290.
- [12] Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poesse, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapidor, Narseo Vallina-Rodriguez, et al. 2020. The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic. *arXiv preprint arXiv:2008.10959* (2020).
- [13] Office for National Statistics. 2019. Local Authority Districts (April 2019) Names and Codes in the United Kingdom. (2019). <http://geoportal.statistics.gov.uk/datasets/local-authority-districts-april-2019-names-and-codes-in-the-united-kingdom>
- [14] Office for National Statistics. 2020. Pen portraits and radial plots. (2020). <https://www.ons.gov.uk/methodology/geography/geographicalproducts/areaclassifications/2011areaclassifications/penportraitsandradialplots>

- [15] Vanessa Frias-Martinez, Jesus Virseda, Alberto Rubio, and Enrique Frias-Martinez. 2010. Towards large scale technology impact analyses: Automatic residential localization from mobile phone-call data. In *Proceedings of the 4th ACM/IEEE international conference on information and communication technologies and development*. 1–10.
- [16] Christopher G Gale, Alexander D Singleton, Andrew G Bates, and Paul A Longley. 2016. Creating the 2011 area classification for output areas (2011 OAC). *Journal of Spatial Information Science* 2016, 12 (2016), 1–27.
- [17] Alessandro Galeazzi, Matteo Cinelli, Giovanni Bonaccorsi, Francesco Pierri, Ana Lucia Schmidt, Antonio Scala, Fabio Pammolli, and Walter Quattrociocchi. 2020. Human Mobility in Response to COVID-19 in France, Italy and UK. *arXiv preprint arXiv:2005.06341* (2020).
- [18] Marta C Gonzalez, Cesar A Hidalgo, and Albert-Laszlo Barabasi. 2008. Understanding individual human mobility patterns. *nature* 453, 7196 (2008), 779–782.
- [19] Health and GOV.UK Social Care. 2020. COVID-19: track coronavirus cases. (2020). <https://www.gov.uk/government/publications/covid-19-track-coronavirus-cases>
- [20] Google Inc. 2020. Google COVID-19 Community Mobility Report. (2020). <https://www.google.com/covid19/mobility>
- [21] Sibren Isaacman, Richard Becker, Ramón Cáceres, Stephen Kobourov, Margaret Martonosi, James Rowland, and Alexander Varshavsky. 2011. Identifying important places in people’s lives from cellular network data. In *International Conference on Pervasive Computing*. Springer, 133–151.
- [22] Jayson S Jia, Xin Lu, Yun Yuan, Ge Xu, Jianmin Jia, and Nicholas A Christakis. 2020. Population flow drives spatio-temporal distribution of COVID-19 in China. *Nature* (2020), 1–5.
- [23] Moritz UG Kraemer, Adam Sadilek, Qian Zhang, Nahema A Marchal, Gaurav Tuli, Emily L Cohn, Yulin Hswen, T Alex Perkins, David L Smith, Robert C Reiner, et al. 2020. Mapping global variation in human mobility. *Nature Human Behaviour* (2020), 1–11.
- [24] Moritz UG Kraemer, Chia-Hung Yang, Bernardo Gutierrez, Chieh-Hsi Wu, Brennan Klein, David M Pigott, Louis Du Plessis, Nuno R Faria, Ruoran Li, William P Hanage, et al. 2020. The effect of human mobility and control measures on the COVID-19 epidemic in China. *Science* 368, 6490 (2020), 493–497.
- [25] Paige Maas, Shankar Iyer, Andreas Gros, Wonhee Park, Laura McGorman, Chaya Nayak, and P Alex Dow. 2019. Facebook Disaster Maps: Aggregate Insights for Crisis Response and Recovery. In *Proceedings of the 16th International Conference on Information Systems for Crisis Response and Management (ISCRAM), Valencia, Spain*. 19–22.
- [26] Raul Montoliu and Daniel Gatica-Perez. 2010. Discovering human places of interest from multimodal mobile phone data. In *Proceedings of the 9th International Conference on Mobile and Ubiquitous Multimedia*. 1–10.
- [27] Santi Phithakkitnukoon, Zbigniew Smoreda, and Patrick Olivier. 2012. Socio-geography of human mobility: A study using longitudinal mobile phone data. *PLoS one* 7, 6 (2012).
- [28] Chaoming Song, Zehui Qu, Nicholas Blumm, and Albert-László Barabási. 2010. Limits of predictability in human mobility. *Science* 327, 5968 (2010), 1018–1021.
- [29] TomTom.com. 2020. The effect of the COVID-19 measures on mobility in the UK. (2020). <https://www.tomtom.com/covid-19/country/uk/>
- [30] Gregory A Wellenius, Swapnil Vispute, Valeria Espinosa, Alex Fabrikant, Thomas C Tsai, Jonathan Hennessy, Brian Williams, Krishna Gadepalli, Adam Boulange, Adam Pearce, et al. 2020. Impacts of state-level policies on social distancing in the united states using aggregated mobility data during the covid-19 pandemic. *arXiv preprint arXiv:2004.10172* (2020).