This is the author's accepted version of the article. The final version published by IEEE is A.F. Zanella, O.E. Martínez-Durive, S. Mishra, Z. Smoreda and M. Fiore, "Impact of Later-Stages COVID-19 Response Measures on Spatiotemporal Mobile Service Usage," IEEE INFOCOM 2022 - IEEE Conference on Computer Communications, 2022, pp. 970-979, doi: 10.1109/INFOCOM48880.2022.9796888

Impact of Later-Stages COVID-19 Response Measures on Spatiotemporal Mobile Service Usage

André Felipe Zanella*†, Orlando E. Martínez-Durive*†, Sachit Mishra*†, Zbigniew Smoreda‡ and Marco Fiore*
*IMDEA Networks Institute, Spain, †Universidad Carlos III de Madrid, Spain, ‡SENSE / Orange Innovation, France
{andre.zanella, orlando.martinez, sachit.mishra, marco.fiore}@imdea.org, zbigniew.smoreda@orange.com

***Abstract*—The COVID-19 pandemic has affected our lives and how we use network infrastructures in an unprecedented way. While early studies have started shedding light on the link between COVID-19 containment measures and mobile network traffic, we presently lack a clear understanding of the implications of the virus outbreak, and of our reaction to it, on the usage of mobile apps. We contribute to closing this gap, by investigating how the spatiotemporal usage of mobile services has evolved through different response measures enacted in France during a continued seven-month period in 2020 and 2021. Our work complements previous studies in several ways: (*i*) it delves into individual service dynamics, whereas previous studies have not gone beyond broad service categories; (*ii*) it encompasses different types of containment strategies, allowing to observe their diverse effects on mobile traffic; (*iii*) it covers both spatial and temporal behaviors, providing a comprehensive view on the phenomenon. These elements of novelty let us lay new insights on how the demands for hundreds of different mobile services are reacting to the new environment set forth by the pandemics.**

I. Introduction

The COVID-19 pandemic has affected lives worldwide in a way that is unprecedented in modern times. The response measures that governments have adopted to contain the virus have changed the lifestyle of billions. Under severely restrained mobility regulations, the telecommunication infrastructure has played a key role, allowing people to communicate, work, entertain and even carry out physical activity in the most normal way possible. As proven by early studies, this has determined significant changes in the use of networks [1], [2].

Mobile services and COVID-19. In this paper, we contribute to the body of knowledge about the impact of the COVID-19 emergence on network usage. More precisely, we focus on mobile services, or *apps*, and investigate how their consumption has evolved throughout periods characterized by different pandemic containment measures.

To this end, we analyse the demands generated by hundreds of apps in the whole territory of France. We observe the modification of such usages across seven months from 2020 to 2021, and correlate them with the heterogeneous restrictions enacted by the local government over that time frame. Our study builds on mobile data traffic information collected in the nationwide infrastructure of Orange, a major telecommunications operator with a 35% market share and over 20 million clients in France, and a 99% coverage of the population. The spatial scale and penetration level of the data lets us explore also the geographical dimension of mobile service usage changes.

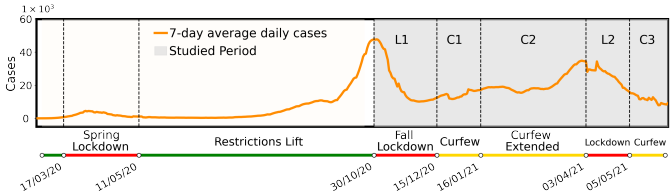


Fig. 1: Timeline of COVID-19 cases and responses in France.

As detailed in Section II, the perspectives we take in our work, combining individual mobile services, multiple response measures, and both temporal and spatial dimensions of the phenomenon, have not been explored by previous studies on the impact of COVID-19 on network traffic.

COVID-19 measures in France. The COVID-19 containment measures adopted in France resulted in three lockdowns, spaced out by periods with varied responses, as illustrated in Figure 1, along with the 7-day moving average of daily cases in France [3]. The first lockdown (March 17 – May 10, 2020) forced the majority of public places, including schools and restaurants, to close, social gatherings to be forbidden, and personal mobility to be limited to essential tasks. As in many countries in Europe, the subsequent period (May 11, 2020 – October 29, 2020) was characterized by a progressive lift on the restrictions, with a re-opening of public spaces under the requirement to preserve social distancing. The growth of COVID-19 cases later (October 17 – October 30, 2020) pushed authorities to enforce a 9 PM – 6 AM curfew, first in a few major cities and then to the majority of the country. A second nationwide lockdown followed (October 31 – December 14, 2020), with similar measures to the first one, except for primary and secondary schools staying open. Afterwards, non-essential services started to re-open, and travel restrictions were lifted, but an overnight curfew was maintained between 8 PM and 6 AM (December 15, 2020 – January 15, 2021). The curfew was later (January 16 – April 2, 2021) anticipated to 6 PM. A new increase in infections determined varied local measures, and ultimately a third national lockdown (April 3 – May 4, 2021); non-essential travel was again prohibited, and schools closed, but with lighter measures overall. This period was again followed (May 5 – May 30, 2021) by a progressive lift of restrictions. On a side note, vaccination started in France on December 27, 2020, but the incidence was not substantial during the observed period, with only 16.35% of the population fully vaccinated by May 30, 2021 [4].

©2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Contributions and key insights. Our study covers the period from October 2020 to May 2021, hence encompasses two lockdowns (L1 and L2 in the following), and three curfew periods (C1, C2, and C3). The rationale is that previous studies have already explored the impact of the first lockdown and subsequent period on mobile traffic; our aim is instead investigating if and how diverse response measures have affected mobile service usage at later stages of the epidemics.

By adopting such an approach, we provide new insights on the spatial and temporal dynamics of both total traffic and demands for specific services, which stem from the succession of more and less restrictive pandemic response strategies. Our key insights, later detailed throughout the paper, are as follows.

- We unveil the extremely rich and variegate range of reactions that individual mobile services have in front of restrictions of different severity; such diversity was hidden in previous analyses that have focused on traffic aggregates or service categories –as we show that apps on same categories can have very heterogeneous behaviors.
- We prove that there are services that are more affected than others by later-stages control measures in both time and space patterns, and detail representative cases.
- We reveal that the pre-COVID-19 weekly activity peaks of the total mobile data traffic have changed in a non-negligible manner one year into the pandemic, due to shifts in the consumption of a few high-volume services.
- We expose the de-urbanization of mobile service usage during lockdowns, and show how less restrictive measures such as curfews reverse that phenomenon.

II. Related work

The impact of the pandemic on internet traffic has been extensively studied in the past year, at different levels across the network. Most works have focused on Internet traffic at large. At Internet Service Providers (ISPs) located in Central Europe, traffic increased by 15% to 20% during the initial 2020 lockdown, with a much higher growth than in a typical year. Such dynamics can be ascribed to the restrictions mandated by governments, resulting also in dynamic changes in weekday patterns that started looking similar to those in weekends [6], [2]. Similar trends occurred in large ISPs in the United States, with an increase from 30% to 60% of peak traffic rates during the first quarter of 2020 [7]. Not only operators, but companies that provide online services also noted significant traffic changes; for instance, Facebook initially observed short periods of sharp increase in their edge network traffic, with a subsequent steady increase of load; they also reported user behavior variations, such as an increase of interest in live streaming services [8]. The sharp traffic changes due to new user behaviors have been also seen at smaller scales: a 90% reduction in downlink traffic was recorded in the university network traffic, due to the classes becoming remote; at the same time, uploads grew due to the much more frequent usage of locally-hosted online teaching platforms [9]. The pandemic also impacted latency in the Internet, with delays values 3 to 4 times higher than in 2019 [10].

When looking at the specific context of mobile networks, the overall trend is different: restrictions in the UK resulted in a decrease of 24% in downlink mobile data traffic over the whole country. The changes were sharper across cosmopolitan areas, which experienced a 50% mobile data traffic decrease, while rural areas were more stable after the lockdown [1]. A lot of attention was in fact drawn by mobility measurements based on mobile network metadata: notably, network metrics showed a steep decrease in the population mobility over the whole UK during the first two weeks after the initial movement restrictions, followed by a slight uniform increase afterwards; also, city-level analyses proved that people tended to move from the more dense metropolitan areas to the urban outskirts just before the measures took place [1]. At a national level, network mobility metrics showed a 65% decrease in displacements over France during the first nationwide lockdown, for both short- and long-range trips [5]. The phenomenon started one week before the lockdown was enforced, and, also in this case, fluxes could be observed leaving large conurbations towards more rural or touristic places. Restrictions during the first French lockdown also disrupted rush-hour commuting patterns.

Table I summarizes the related work to date, and highlights how our study complements and goes beyond those available to date. Specifically, the vast majority of previous analyses focused on the impact of the first lockdown on network usage, and none investigated later stages of the pandemic response, where the enacted measures became more varied, and the population increasingly accustomed to those. Also, most research has considered aggregate traffic volumes, typically in ISPs or Internet Exchange Points (IXPs), possibly disaggregated into few macroscopic service categories; very limited attention has been paid to individual mobile services, which are at the core of our study. In the light of these considerations, our work offers new insights on how the COVID-19 pandemic affected the spatiotemporal consumption of hundreds of mobile apps.

III. Measurement data

Our study hinges on extensive mobile network data, whose collection, preprocessing, and ethical aspects are detailed next.

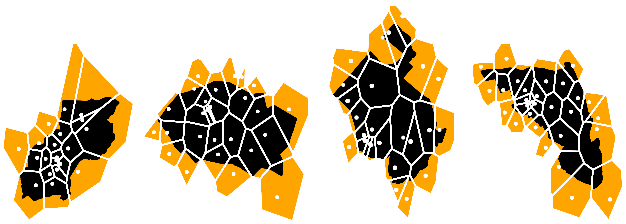
A. Data collection

The mobile data traffic information we use was collected in the production network of Orange servicing the metropolitan France territory. The operator employed passive measurement probes to monitor the Gi, SGi, and Gn interfaces that connect Gateway GPRS Support Nodes (GGSNs) and Packet Data Network Gateway (PGWs) to external Public Data Networks (PDNs), gathering data about the traffic generated by mobile subscribers using 2G, 3G and 4G connectivity. This basically represents the full mobile demand, as 5G generated less than 1% of the total mobile data traffic as of May 2021 [11].

The mobile service generating each IP session was identified via a combination of Deep Packet Inspection and proprietary traffic classifiers deployed by the operator. Each IP session was geo-referenced at the level of Base Transceiver Station (BTS), leveraging the User Location Information (ULI) contained in

Study	Temporal coverage	Spatial coverage	Network	Metrics
Lutu *et al.*, [1]	First wave	UK, at regional and city levels	Mobile network	Total data traffic, mobility metrics
Pullano *et al.*, [5]	First wave	France, at regional level	Mobile network	Mobility metrics
Feldmann *et al.*, [2], [6]	First/second waves	Europe, US, Madrid	ISP, IXP, mobile network	Total traffic and 5 service categories
Liu *et al.*, [7]	First wave	US, nationwide	ISP	Total traffic
Bottget *et al.*, [8]	First wave	Thirteen countries worldwide	Edge network	Facebook traffic and 4 categories
Favale *et al.*, [9]	First wave	Turin, Italy, at a campus level	LAN	Total traffic
Candela *et al.*, [10]	First wave	Five countries in Europe	ISP	Latency
Ours	Second/third waves	France, nationwide / communes	Mobile network	Total traffic and 285 individual apps

TABLE I: Summary of previous studies of the impact of the COVID-19 pandemic on telecommunication infrastructures.



(a) Roanne (b) Saint-Priest (c) Pantin (d) Noisy-le-Grand

Fig. 2: Examples of spatial interpolation of Voronoi cells for 4G BTSs (gold) on four random communes (black) of France.

Packet Data Protocol (PDP) Contexts and Evolved Packet System (EPS) Bearers over the GPRS Tunneling Protocol control plane (GTP-C). Based on this, the operator computed the hourly traffic demand for each mobile service at every BTS, by aggregating the uplink and downlink traffic of all users attached to the same BTS during one-hour intervals.

B. Data preprocessing

The data collected by the operator for this study encompasses all $370,189$ BTSs servicing in territory of Metropolitan France, for the period between October 30, 2020 and May 31, 2021. As a spatial granularity at BTS level is exceedingly high for a nationwide study, we further aggregated the hourly traffic demand of each mobile app over 35,180 *communes*, *i.e.*, local administrative units. To perform the spatial interpolation of BTS-level traffic to communes zoning, we proceed as illustrated in Figure 2. First, we tell apart BTSs into three sets according to their technology, *i.e.*, 2G, 3G, or 4G. Second, for each set of BTS separately, we use the coordinates of BTSs as bi-dimensional anchors to compute a Voronoi tessellation of space; we assume each Voronoi cell to represent the coverage area of the corresponding BTS, and the BTS traffic to be uniformly distributed over the cell [12]. Third, we assign to each commune a fraction of the mobile app traffic recorded at each BTS and hour, proportional to the percentage of the Voronoi cell surface intersecting the commune territory. Finally, we compute the hourly demand for a service in a given commune as the sum of the contributions from all BTSs.

The final data we use in the paper describes the hourly traffic generated in every commune of France by 285 mobile services. Those encompass a wide range of apps for video and audio streaming, gaming, messaging, social networking, or travel planning, just to name a few categories. The monitored apps include all most popular and traffic-exacting services: a non-comprehensive list of sample apps is in Table II.

C. Ethics considerations

The dataset supporting this study consists of hourly demands for individual mobile apps recorded in French communes. This representation entails a high spatiotemporal aggregation that merges the IP sessions of thousands of subscribers or more, hence granting that no individual data subject can be re-identified from the data. In fact, traffic demands over communes do not configure as personal data according to the General Data Protection Regulation (GDPR) [13], and our research does not involve risks for the mobile subscribers. In addition, the data collection and aggregation occurred in full compliance of article 89 of the GDPR, under the supervision of the Data Protection Officer (DPO) of the operator, and upon authorization by the French National Commission on Informatics and Liberty (CNIL). Specifically, the aggregation was performed in a secure platform at the operator premises, and the raw measurements were deleted immediately afterwards.

IV. Temporal analysis

In the first part of our analysis, we focus on the temporal dynamics of mobile service usage over the whole country of France, looking at both demand volumes and typical weekly patterns throughout the studied period.

We will use the following notation. Let $d_c^s(t)$ be the demand for service $s \in \mathcal{S}$ recorded in commune $c \in \mathcal{C}$ at time t; we then refer to $d_c(t) = \sum_s d_c^s(t)$ as the total demand generated by all services in c at t, and to $d^s(t) = \sum_c d_c^s(t)$ as the demand for service s at time t over the whole country. Similarly, the global demand at t is $d(t) = \sum_s \sum_c d_c^s(t)$. Let us also define six macroscopic time periods, each associated with different pandemic containment measures: T_{L1}, T_{C1}, T_{C2}, T_{L2}, T_{C3} denote the time span of the periods in the subscript. Also, T_{21} is the concatenation of all these periods.

A. Demand volume analysis

For the volume analysis, we reduce the temporal granularity of the data to days, since this is a better resolution to highlight behaviors over the whole seven-month time span. Hence, the time index t indicates the day over which the demand is aggregated. We first briefly explore trends in total mobile data traffic, and then delve into the dynamics of individual apps.

1) Total traffic: The daily evolution of the total traffic volume recorded in the nationwide Orange network between October 2020 and May 2021 is displayed in Figure 3. A standard score normalization is used in this plot, and in all subsequent time series, so as not to disclose the actual traffic loads, which are considered sensitive by the operator.

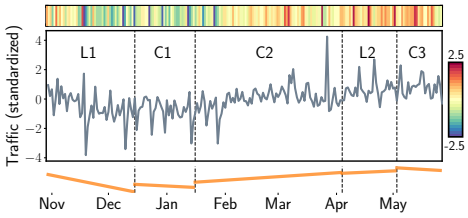


Fig. 3: Total traffic volume transiting in the Orange mobile network during the observed seven-month period, as a color scale (top), time series (middle), and linear interpolation over time periods with different responses (bottom).

Formally, the standard score of daily total traffic in T_{21} is

$$z(t) = \frac{d(t) - \frac{1}{|T_{21}|}\sum_{t\in T_{21}} d(t)}{\frac{1}{|T_{21}|}\sqrt{\sum_{t\in T_{21}}\left(d(t) - \frac{1}{|T_{21}|}\sum_{t\in T_{21}} d(t)\right)^2}}. \quad (1)$$

The figure clearly shows how L1 determined a continued drop in mobile data traffic; L2 had a softer effect, but still curbed the traffic volume. The linear interpolations at the bottom of Figure 3 emphasizes these trends. After both lockdowns, mobile traffic tended to recover fairly slowly, as C1 and C3 are both characterized by fairly constant loads in time. C2 is the only period where mobile traffic starts growing, which used to be the norm in pre-pandemic times [14]; the reduced growth at the end of C2 can be ascribed to the fact that several densely populated departments of France already started local lockdowns a couple of weeks before the nationwide one.

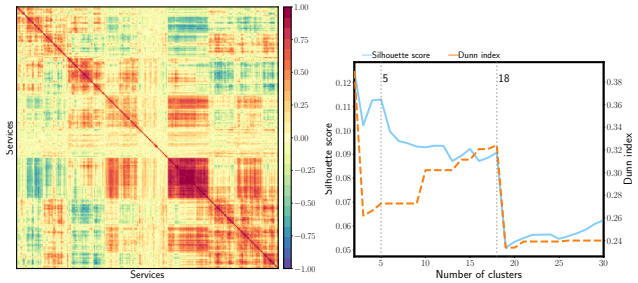
Insights. *Our analysis corroborates the findings of previous studies, confirming that the stringent restrictions enforced during lockdowns hinder the utilization of mobile networks [1], and expose that this occurs also at later stages of the pandemic and not only during responses to its first wave. A new element unveiled by our study is that milder measures such as curfews do not instead curb mobile traffic, but allow for a slow recovery towards conventional trends of previous years.*

2) Individual mobile services: An interesting question is whether the dynamics above homogeneously characterize all mobile services, or if the pandemic had varying impact on diverse apps. To answer the question, we describe each service as its normalized daily time series $z^s(t)$ over the seven-month observation period[1]; formally,

$$z^s(t) = \frac{d^s(t) - \frac{1}{|T_{21}|}\sum_{t\in T_{21}} d^s(t)}{\frac{1}{|T_{21}|}\sqrt{\sum_{t\in T_{21}}\left(d^s(t) - \frac{1}{|T_{21}|}\sum_{t\in T_{21}} d^s(t)\right)^2}}. \quad (2)$$

It is worth noting that this normalization makes the times series of different apps directly comparable, removing biases due to assorted popularity and service categories (which entail very diverse loads), and burstiness of usage (which causes heterogeneous variances). Therefore, we can compute

[1]Here, we filter out from the data and subsequent analysis of $z^s(t)$ all vacation periods, which, as we will later see, can severely affect apps usage. By doing so, we ensure that our results do not reflect (dis)similarities among services caused by the way they are consumed during holidays.



(a) (b)

Fig. 4: (a) Matrix of pairwise distances between traffic volume time series of apps. (b) Stopping rules versus cluster number.

Euclidean distances between the normalized time series of each couple of services, which results in the pairwise distance matrix in Figure 4a. The presence of clear clusters of apps in the matrix calls for further investigation. We thus run a hierarchical clustering based on the Ward algorithm [15], and use the Silhouette score [16] and Dunn index [17] as stopping rules to determine the best number of clusters in the matrix. Figure 4b shows that both rules indicate[2] 18 clusters as the optimal number. Moreover, the Silhouette score highlights 5 clusters as a good partitioning of services.

We provide a comprehensive representation of the traffic volume evolution of the clusters of individual mobile services in Figure 5. The 5 macro-clusters tell apart high-level behaviors: I contains apps with higher weekend traffic, II apps with decreasing traffic, III apps with noisy dynamics, IV apps with higher working hours usage, and V apps with increasing traffic.

The macroscopic behaviors are too coarse to be informative, hence we focus on the micro-clusters. The full list, including a brief description of their main characterizing features, as representative examples of apps, is in Table II. The table, jointly with Figure 5, reveals the substantial complexity and variety of the temporal evolution of single mobile services under the hood of the simpler view offered by total traffic. Many micro-clusters show trends that are resilient to the COVID-19 containment measures enforced throughout our period of observation: some, like B, H or N, show steady patterns across all periods; others, like C, are characterized by a growing popularity; and others, like G, suffer from a fairly consistent loss of users. Instead, our interest is on dynamics that can be ascribed, at least in part, to the epidemics response.

In that sense, the micro-clusters A, D, E and F all show a significant increase in traffic volumes during the lockdowns. Sample time series of specific services in those clusters are in Figure 6. Although there are discrepancies that make the Ward algorithm cluster the apps separately, we can observe in all cases sustained higher traffic during L1 (in E, such as Pinterest in Figure 6c, and in F, such as Xiaomi Mi Home in Figure 6d), during L2 (in A, such as Houseparty in Figure 6a), or during both L1 and L2 (in D, such as Disney+ in Figure 6b).

[2]The two stopping rules measure clustering efficiency in terms of (ideally small) intra-cluster variance and (ideally high) inter-cluster separation. The rules pinpoint the best number of clusters as a high value before a steep decrease that denotes a clear drop in the quality of intra- and inter-clustering.

Cluster	Description	Samples
A	No weekly pattern, increasing in L1, C2 or L2; gaming and messaging apps mainly	WhatsApp, League of Legends
B	Higher usage in weekends, steady over time; video streaming and gaming apps mainly	Netflix, Youtube, Steam, PUGB
C	Higher usage in weekends, increasing over time; gaming apps mainly	MineCraft, Fornite
D	Higher usage in weekends, increasing around L1 and L2; video streaming and gaming apps mainly	Disney+, Apple Video, CounterStrike
E	No weekly pattern, increasing in proximity of L1 and L2	Pinterest
F	No weekly pattern, increasing in L1; gaming and conferencing apps mainly	Zoom, Clash of Kings, Angrywords
G	No weekly pattern, decreasing over time	Battle.net, Shazam
H	Slightly higher usage in working hours, steady over time	N26, Dropbox
I	Higher usage in working hours, decreasing in time; business apps mainly	Evernote, Twitter, Microsoft Office
J	No weekly pattern, increasing in C3	Amazon Prime Video, WeChat
K	No weekly pattern, noisy over time; OS update services mainly	Microsoft Windows Update
L	Slightly higher usage in weekends, noisy in time; gaming apps mainly	World of Warcraft, Playstation
M	No weekly pattern, slightly increasing in C2	Psiphon, Coyote
N	Higher usage in working hours, steady over time; office applications mainly	Gmail, Skype, Google Docs
O	No weekly pattern, increasing in C2 and C3	Telegram, TikTok, Uber
P	No weekly pattern, increasing in C2 and substantially more in C3; location-based services mainly	TripAdvisor, Foursquare, Spotify
Q	Higher usage in working hours, increasing in C2 and C3	Google Maps, Waze, AirFrance
R	No weekly pattern, increasing over time	Twitch, Google Meet

TABLE II: List of 18 micro-clusters issued by the clustering in Figure 5, with a brief description and representative apps.

The reason for such dynamics is simple: videoconferencing, on-demand television, social media, or smart-home managers are all examples of mobile services that are more frequently consumed at home, where people spent a much larger portion of time during lockdowns. Importantly, we recall that we are looking at traffic generated at BTS, hence by mobile devices connected to the cellular network, and not, *e.g.*, to home Wi-Fi hubs. Therefore, the services in the clusters above all demonstrate that a non-negligible portion of the Orange user population is actually employing 2G, 3G and 4G technologies as a way to access the Internet from home.

Another notable behavior induced by COVID-19 is that of micro-clusters that show a dual behavior to the one above. Namely, the apps in O, P, and Q are curbed by lockdowns, and record an increased usage during the relaxed measures in curfews. A closer look reveals that the vast majority of these services are directly or indirectly related to personal mobility: the limitations to movements determined by L1 and L2 clearly reduce their utility. Interestingly, the dynamics are slightly different in apps for general mobility and for a more leisure-oriented mobility. The first case includes services with a marked working-hour pattern (in Q, such as Google Maps or Waze in Figure 6g) and with more regular usage (in O, such as Uber or Apple Maps in Figure 6e), and show moderate increase in usage in both C2 and C3. Instead, apps targeting mobility during free time (in P, such as Foursquare or TripAdvisor in Figure 6f) show a dramatic increase in usage in C3, which we attribute to the combination of more relaxed measures and inviting weather conditions during that period.

To conclude our analysis, we underscore how our approach of considering individual apps is critical to reveal the richness of behaviors above. For instance, Figure 7 shows the traffic time series of a number of popular video streaming services. These mobile services undergo very heterogeneous evolutions in the observed seven months, with volumes that are steady (*e.g.*, YouTube, Netflix) and possibly higher during working hours (*e.g.*, Skype), declining (*e.g.*, Zoom), or heavily dependent on pandemic response measures (*e.g.*, Amazon Prime Video). In fact, these services are classified in *different* micro-clusters by our analysis. Had video streaming been treated as a single category, all this diversity would have been lost.

Insights. *Individual mobile services create an entangled ecosystem of usages, and are affected by pandemic response measures in very diverse ways. Therefore investigations of traffic volumes aggregated over all services, or even over services belonging to a same category, hides the actual complexity of the usage dynamics, and can lead to misleading conclusions. Instead, our first in-depth look at the behavior of hundreds of apps, allows identifying a variety of micro-behaviors. Specifically, we observe how many services are simply not affected by the COVID-19 emergence, whereas others experience substantial and diversified variations in their demands depending on the containment measures enacted. Clear links can be found between the nature of the latter services, and their specific reaction to restrictions.*

B. Demand pattern analysis

An interesting question is whether COVID-19 measures had not only an impact on the volumes of traffic transiting in the mobile network, but also on the temporal distribution of such traffic. For instance, previous works have showed that the typical difference between the hourly traffic pattern in working days and weekends tends to disappear during a lockdown [6]. Here, we look at hourly traffic –*i.e.*, the time index t denotes one specific hour– and examine weekly patterns in both total and per-app traffic. Also, we use equivalent weekly patters computed from a three-month control period in 2019 as a reference, so as to understand if and how the daily activity has changed due to COVID-19 responses. We denote by T_{19} the set of hours t in such a control period.

We first focus on typical weekly dynamics that are known to capture most of the variance in the telecommunication activity of individuals [18], [19], and condense the seven-month traffic dynamics into a *median week signature* [20]. Formally, we compute the median traffic in each hour of the week as $w(t) = \mu_{0.5}\{d(t)|t \in \mathcal{M}(t)\}$, where $\mu_{0.5}$ denotes the median of the argument set, and $\mathcal{M}(t)$ is the set of same hours of the week as t (*e.g.*, Mondays at 8 AM). Then, the median week signature

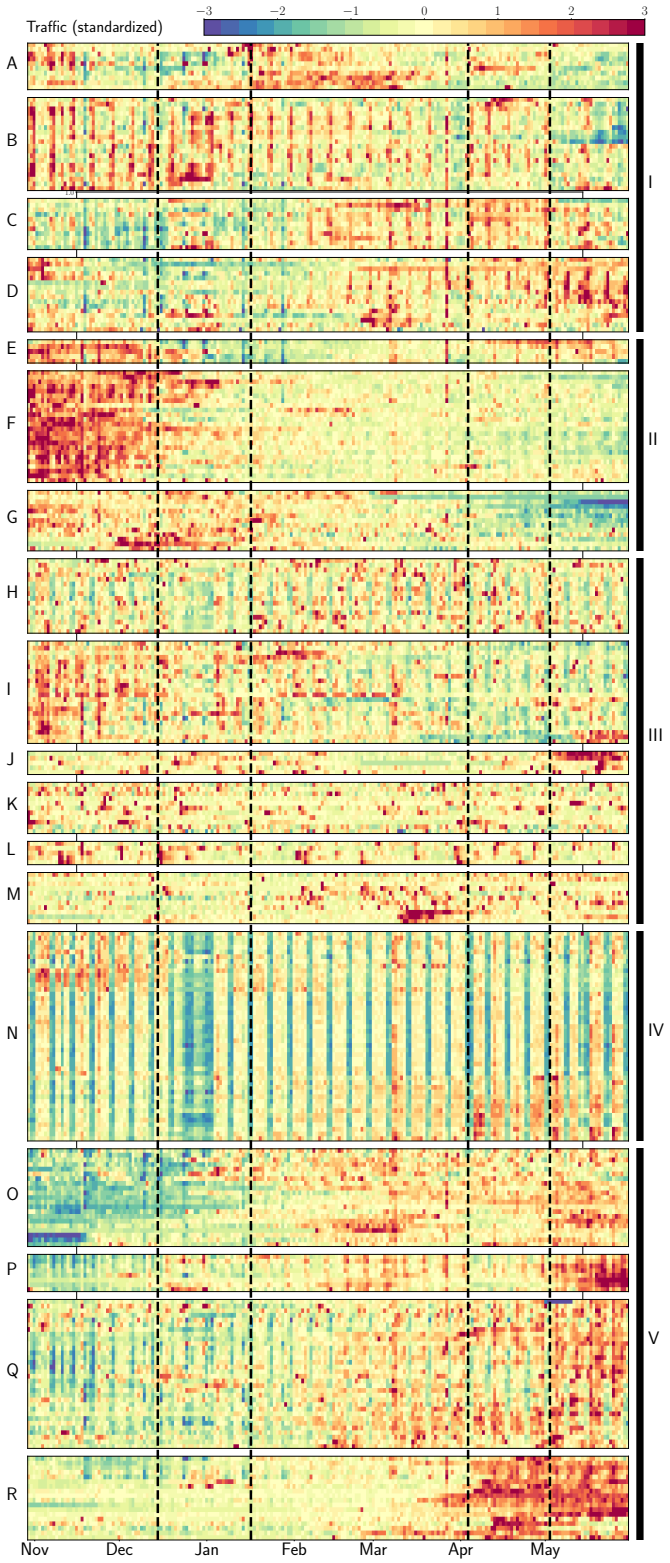


Fig. 5: Normalized time series of the daily traffic volume of individual apps during the seven-month observation period. The left labels highlight 18 micro-clusters A–R, while the right labels mark 5 macro-clusters I–V, obtained with the hierarchical clustering. Dashed vertical lines separate the L1, C1, C2, L2, C3 periods introduced in Figure 1.

is obtained by applying the standard score normalization in (1) to $w(t)$ instead of $d(t)$. Note that disjoint $\mathcal{M}(t)$ are used for T_{21} and T_{19}, so as to obtain independent median weeks during and prior to the COVID-19 pandemic.

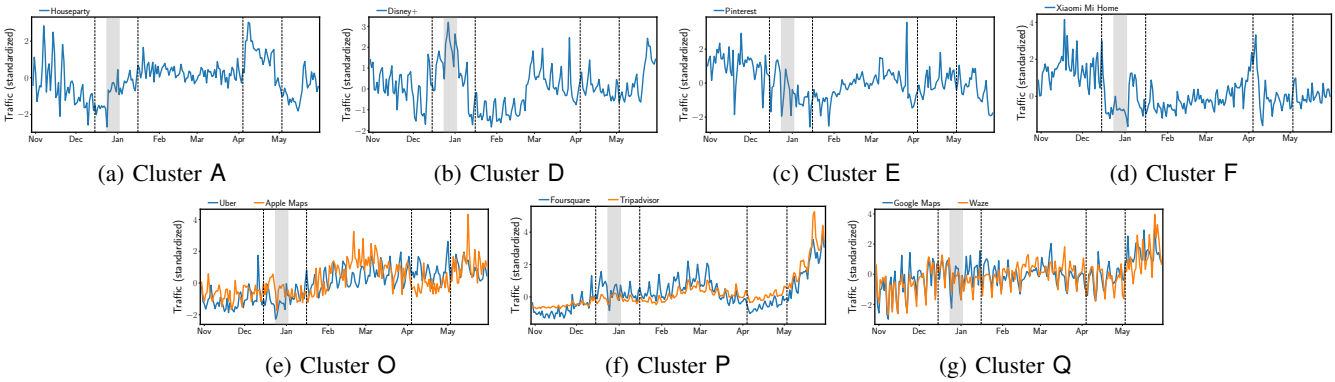
Figure 8 superposes the median week signatures for the considered COVID-19 response period and in the 2019 control period. While peak traffic hours stayed the same, minor changes can be accredited to the enacted restrictions. First, remote working sensibly reduced the need for daily commuting, which explains the disappearance of the early-morning traffic peak in 2020-21. Second, evening peaks during the pandemic are relatively higher than in 2019, which is likely caused by mobility limitations that forced people at home from late afternoon onwards, consistently through the observation period. Third, we confirm the reduced diversity between working and weekend days, which tend to be closer during COVID-19 than they were before.

We then repeat the analysis on a per-app basis. We compute $w^s(t) = \mu_{0.5}\{d^s(t)|t \in \mathcal{M}(t)\}$, for each mobile service s, and derive the app-specific median week signature applying in (2) to $w^s(t)$ instead of $d(t)$. In this case, we also produce separate median weeks for T_{L1}, T_{C1}, T_{C2}, T_{L2}, and T_{C3}, so as to assess the impact of restrictions on the weekly patterns of app usage.

As median week signatures are also standardized, they can be directly compared. We compute, for each service independently, the dynamic time warping between its signatures in different periods, *i.e.*, L1, C1, C2, L2, C3, and in the 2019 control period. Figure 9 shows the result for all mobile services, along columns; the first five rows show the distances of the 2019 median week and those in L1, C1, C2, L2, C3, respectively. The following rows report the distances between different periods in the epidemics, *i.e.*, L1–C1, L1–C2, L1–L2, L1–C3, C1–C2, C1–L2, C1–C3, C2–L2, C2–C3, and L2–C3.

Most apps do not show any significant change in their weekly pattern (*i.e.*, have near-zero or negative distances in all rows), hence the way they are consumed is hardly affected by the pandemic. Among the mobile services that show some diversity (*i.e.*, have positive distances), those on the left (group 1 in the plot) are less popular apps with inherently bursty dynamics that tend to vary all the time, even within weeks of 2019. More interesting is the group of services that show a clear distance between the control and studied period, but no differences in periods within the COVID-19 pandemic (group 2 in the plot). The weekly usage pattern of these apps clearly reacted –in a *uniform* way– to the restrictions.

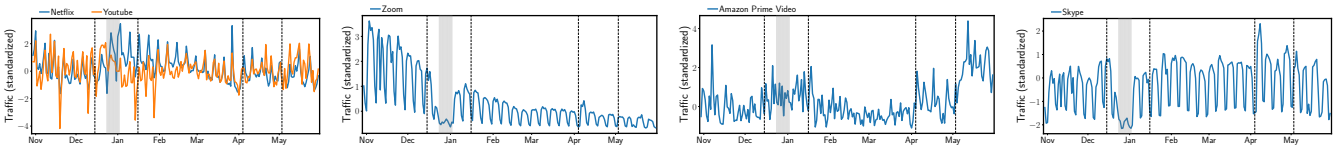
The median weeks of representative mobile services in this group are illustrated in Figure 10. A video calling application such as Skype adjusted to a strongly work-oriented activity pattern, with high peaks in the morning and early afternoon, which also overflew to weekends. Major video streaming services are also in the group of interest. Both YouTube and Netflix saw their early morning peaks disappear, along with home-work commuters who created such demands; given the large volume of traffic of these apps, this also determines the same effect observed in the total traffic in Figure 8. In addition, the incidence of evening traffic grew dramatically during



(a) Cluster A (b) Cluster D (c) Cluster E (d) Cluster F

(e) Cluster O (f) Cluster P (g) Cluster Q

Fig. 6: Time series of traffic volumes for different individual mobile services. The gray shade highlights Christmas vacations, which have been disregarded to avoid biases, as explained in footnote 1. Dashed lines separate the different restriction periods.



(a) Cluster B (b) Cluster F (c) Cluster J (d) Cluster N

Fig. 7: Time series of traffic volumes for different individual mobile services in the video streaming category, by micro-cluster.

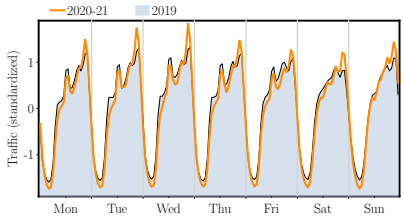


Fig. 8: Total traffic median week in target and control periods.

weekends for Netflix, due to the impossibility for people to enjoy nights out as in pre-pandemic times. Finally, with COVID-19, private cabs were hired with Uber in a new pattern, with reduced hours of operation in working days (owing to no commuting and early curfews), and no evening peaks on Fridays and Saturdays (due to almost absent nightlife).

Insights. *Weekly patterns in the total mobile data traffic show changes during COVID-19, which are however mainly caused by a small subset of popular video streaming apps. In fact, when the vast majority of mobile services are consumed does not change in a significant way during the pandemic.*

V. Spatial Analysis

We now look at whether the temporal changes above occur homogeneously over the French territory, or are the result of geographically diverse effects of the epidemics responses. We leverage data disaggregated at the commune level to this end.

A. Total traffic

We first consider the total mobile data traffic, and compute the average traffic density in each commune during the 2019 control period T_{19} and in the target 2020-21 period. Formally, $\bar{d}_c(T_{19}) = (1/T_{19}) \cdot \sum_{t \in T_{19}} d_c(t)/a_c$, and $\bar{d}_c(T_{21}) = (1/|T_{21}) \cdot \sum_{t \in T_{21}} d_c(t)/a_c$, where a_c is the area of commune c. Since we are interested in understanding if the *relative* geographical distribution of traffic has changed due to COVID-19, we standardize over space the traffic density in each period, as

$$z_c(T_{19}) = \frac{\bar{d}_c(T_{19}) - \frac{1}{|\mathcal{C}|}\sum_{c \in \mathcal{C}} \bar{d}_c(T_{19})}{\frac{1}{|\mathcal{C}|}\sqrt{\sum_{c \in \mathcal{C}} \left(\bar{d}_c(T_{19}) - \frac{1}{|\mathcal{C}|}\sum_{c \in \mathcal{C}} \bar{d}_c(T_{19})\right)^2}}, \quad (3)$$

$$z_c(T_{21}) = \frac{\bar{d}_c(T_{21}) - \frac{1}{|\mathcal{C}|}\sum_{c \in \mathcal{C}} \bar{d}_c(T_{21})}{\frac{1}{|\mathcal{C}|}\sqrt{\sum_{c \in \mathcal{C}} \left(\bar{d}_c(T_{21}) - \frac{1}{|\mathcal{C}|}\sum_{c \in \mathcal{C}} \bar{d}_c(T_{21})\right)^2}}. \quad (4)$$

In practice, z_c is a measure of how the traffic density of commune c compares to that of the average French commune.

The difference in the standardized traffic density between the 2019 control period and the observed pandemic time span in 2020-21 is illustrated in Figure 11. The manifest trend is a significant reduction of the relative importance of cities as the places where the overall mobile traffic demands are generated. Negative differences, in dark blue, pinpoint all large- and medium-sized urban areas in the country. Zoomed views are provided in the bottom part of the figure for the 10 most populated cities: they show even better how the phenomenon is strongly localized in urban centers, whereas the surrounding suburban areas possibly experience a positive difference, *i.e.*, increased contribution to the overall traffic. Indeed, the higher incidence of countryside regions in also visible at a nationwide scale, and is especially strong at locations well known to attract metropolitan inhabitants during vacation periods[3].

The result very neatly demonstrates how COVID-19 measures not only forced inhabitants of major cities at home, so

[3]We recall that we filter out vacations, whose impact would be anyway diluted in seven months: the effect cannot be ascribed to holiday mobility.

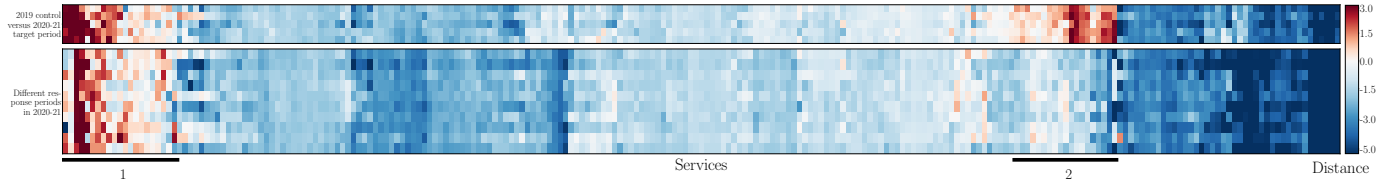
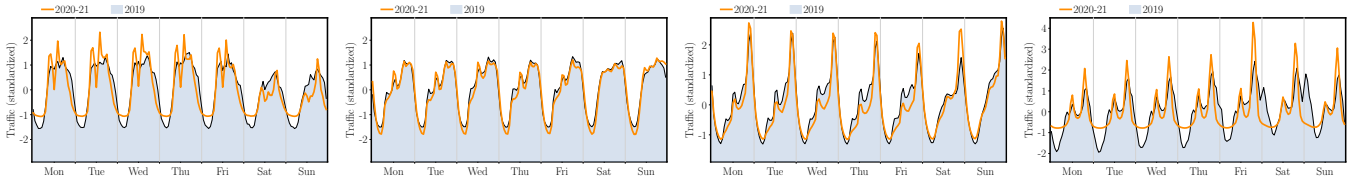


Fig. 9: Distances between the median week signatures of individual apps (columns), comparing pre-pandemic with COVID-19 periods (top rows), and different periods in 2020-21 characterized by varied response measures (bottom rows).



(a) Skype (b) YouTube (c) Netflix (d) Uber

Fig. 10: Median week signatures of representative mobile services in the 2019 control period and in the 2020-21 target period.

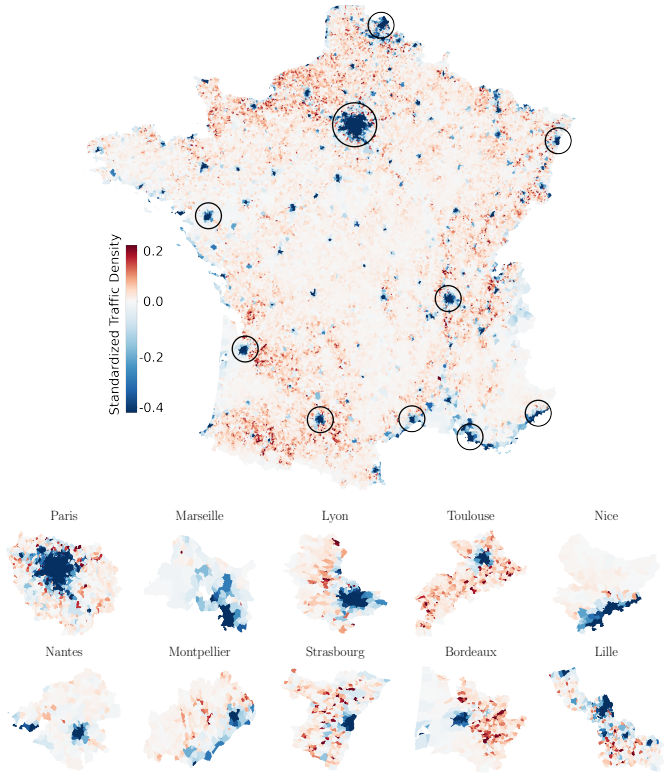
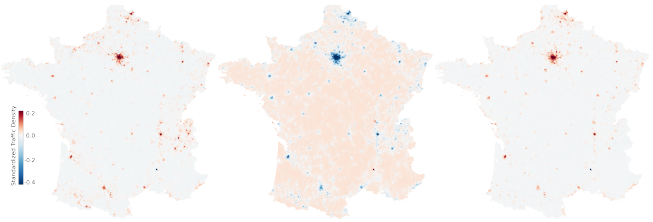


Fig. 11: Difference in the standardized geographical distribution of traffic density between 2019 and 2020-21. Circles highlight the 10 most populated departments in France, for which detailed views are in the bottom part of the figure.

that they increasing relied on Wi-Fi access, and cut down their usual cellular network traffic; rather, it also pushed many people away from city centers, and towards second/vacation homes, or greener places. As most people kept working remotely, and relied on cellular access for Internet connectivity, such a mobility caused the de-urbanization of mobile traffic consumption in Figure 11. In fact, this phenomenon can be even broken down in time, across different periods during the epidemics. Figure 12 shows a similar difference map, but



(a) L1–C2 (b) C2–L2 (c) L2–C3

Fig. 12: Difference in the standardized geographical distribution of traffic density between periods during the pandemic.

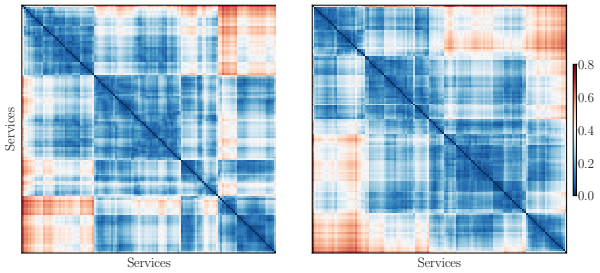


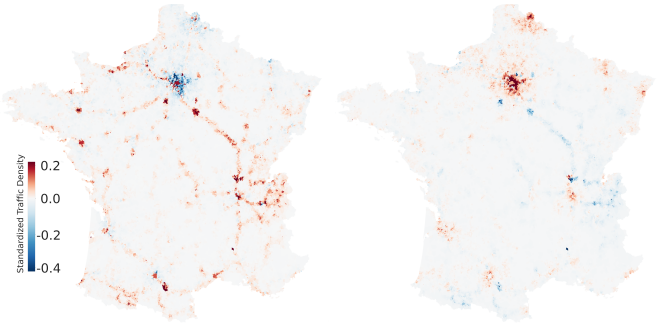
Fig. 13: Sample matrices of pairwise distances between difference maps of each app. Left: L1–C2. Right: C2–L2.

computed for the L1–C2, C2–L2, and L2–C3 period pairs. A striking effect emerge such that entering lockdowns, like C2–L2, determine the effect described above, whereas transitions into more relaxed measures, like L1–C2 and L2–C3 result in a return of traffic towards cities.

Insights. *Restrictive COVID-19 measures exert a very consistent reduction of the contribution of all urban centers to the overall mobile traffic demand. Such de-urbanization of traffic is promptly reverted once restrictions are loosened.*

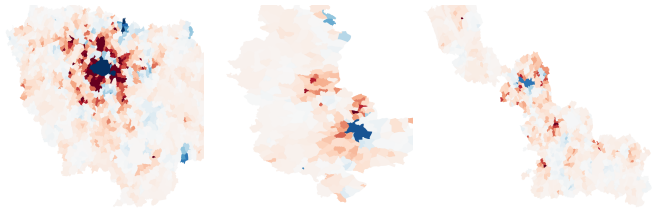
B. Individual mobile services

The dataset we use for our study lets us explore if and how the spatial dynamics above are altered when considering independent mobile services, instead of their aggregated traffic. To this end, we compute standardized traffic densities



(a) L1–C2 (b) C2–L2

Fig. 14: Difference in the standardized geographical distribution of Waze traffic density between example periods.



(a) Paris (b) Lyon (c) Lille

Fig. 15: Difference in the standardized geographical distribution of TripAdvisor traffic density from C2 to L2, for 3 cities.

$z_c^s(T_\star)$ on a per-app basis and for $\star \in \{\text{L1}, \text{C1}, \text{C2}, \text{L2}, \text{C3}\}$, by (i) computing the average traffic density of each service in every commune as $\bar{d}_c^s(T_\star) = (1/T_\star) \cdot \sum_{t \in T_\star} d_c^s(t)/a_c$, and (ii) applying similar equations to (4) where we use $\bar{d}_c^s(T_\star)$ instead of $\bar{d}_c(T_{21})$. This allows producing difference maps like those in Figure 12, for each mobile service separately.

In order to discover macroscopic patterns, we study each combination of periods (*e.g.*, C2–L2) independently. For each case, we summarize the maps above as the probability distributions of the per-commune differences that compose them; then, we compute the pairwise similarity between the distributions of all services, using the Jensen-Shannon distance. This results in distance matrices like those depicted in Figure 13. While those are just two samples, all have a similar structure, highlighting how most services have low distance among them, and spatial dynamics that are aligned to those observed for the total traffic. However, there exist apps for which the alternation of lockdowns and curfews entails fairly unique geographical variations –pointed by high/red values in the matrices.

Due to space limitations, we will only present a couple of interesting examples of such divergent services. The first one is Waze, in Figure 14. Here, changes in the geography of the service traffic are bonded to the road infrastructure: when exiting from the lockdown in L1, an increased incidence of national transportation arteries is observed; this vanishes when entering again in the lockdown during L2, when we can instead note a higher activity around larger cities. This type of behavior is consistent with the fact that long-distance travel is cut out during lockdowns [5]; it also corroborates our previous reasoning on the causes of the de-urbanization of mobile traffic during lockdowns, which push people to move towards the close proximity of their cities of residence –and use apps like Waze to find their way there. The second example is TripAdvisor. Here, *residential* urban areas located around the city centers consistently gain importance as mobile traffic sources when COVID-19 measures are tightened; city centers themselves have the opposite trend. This is shown in Figure 15 for three sample cities entering the L2 lockdown. We argue that these peculiar spatial dynamics can be the results of lockdowns keeping workers far from downtown offices, and forcing them to order delivery food from their homes.

Insights. *While the majority of mobile services follow global spatial trends determined by COVID-19 control measures, some apps are affected in unique ways by the transitions between restrictions. As this is due to the distinctive nature of such apps, an in-depth analysis of their geographies across responses may have interesting applications beyond networking, e.g., to study how the mobility, working, shopping or eating habits evolve in different areas as a result of the pandemic.*

VI. Conclusions

We conducted a first investigation of the impact of COVID-19 response measures on the usage of individual mobile services, targeting a major European country, *i.e.*, France, and a seven-month observation periods that encompasses multiple lockdowns and curfews. Our results build upon substantial measurement data collected in the production network of a main mobile operator, *i.e.*, Orange. They shed new light on the temporal and spatial dynamics of mobile apps consumption at later stages of the pandemic, revealing a number of previously unknown behaviors in both total traffic and specific service demands that result from the alternation of more or less restrictive control strategies enacted by the French government. By doing so, we contribute substantial original knowledge to the (still fairly thin) literature on the effects of COVID-19 on telecommunication systems. We also argue that the demonstrated existence of unique patterns in the reaction to COVID-19 measures for specific apps with a quite narrow scope (*e.g.*, targeting leisure mobility, e-commerce, or work activities) may pave the way to the use of mobile service consumption as a data source for research in other domains. Indeed, such data could offer unique insights to understand the behavior of individuals in presence of pandemic containment measures, across social dimensions like personal movements, shopping habits, or remote work schedules.

The authors have provided public access to their code at https://github.com/afzanella/infocom22_covid.

Acknowledgments

This work of was supported by the ANR-18-CE25-0011 project CANCAN. A.Z. is part of the ITN BANYAN, which received funding from the EU's H2020 research and innovation programme under grant agreement no. 860239. The work of O.M. and M.F. was supported by NetSense, grant no. 2019-T1/TIC-16037 funded by Comunidad de Madrid.

REFERENCES

[1] A. Lutu, D. Perino, M. Bagnulo, E. Frias-Martinez, and J. Khangosstar, "A characterization of the covid-19 pandemic impact on a mobile network operator traffic," in *Proceedings of the ACM Internet Measurement Conference*, IMC '20, (New York, NY, USA), p. 19–33, Association for Computing Machinery, 2020.

[2] A. Feldmann, O. Gasser, F. Lichtblau, E. Pujol, I. Poese, C. Dietzel, D. Wagner, M. Wichtlhuber, J. Tapiador, N. Vallina-Rodriguez, and et al., "A year in lockdown," *Communications of the ACM*, vol. 64, no. 7, p. 101–108, 2021.

[3] W. H. Organization, "Who coronavirus (covid-19) dashboard."

[4] E. Mathieu, H. Ritchie, E. Ortiz-Ospina, M. Roser, J. Hasell, C. Appel, C. Giattino, and L. Rodés-Guirao, "A global database of COVID-19 vaccinations," *Nature Human Behaviour*, vol. 5, pp. 947–953, May 2021.

[5] G. Pullano, E. Valdano, N. Scarpa, S. Rubrichi, and V. Colizza, "Evaluating the impact of demographic, socioeconomic factors, and risk aversion on mobility during COVID-19 epidemic in france under lockdown: a population-based study," June 2020.

[6] A. Feldmann, O. Gasser, F. Lichtblau, E. Pujol, I. Poese, C. Dietzel, D. Wagner, M. Wichtlhuber, J. Tapiador, N. Vallina-Rodriguez, O. Hohlfeld, and G. Smaragdakis, "The lockdown effect: Implications of the covid-19 pandemic on internet traffic," in *Proceedings of the ACM Internet Measurement Conference*, IMC '20, (New York, NY, USA), p. 1–18, Association for Computing Machinery, 2020.

[7] S. Liu, P. Schmitt, F. Bronzino, and N. Feamster, "Characterizing service provider response to the covid-19 pandemic in the united states," 2020.

[8] T. Böttger, G. Ibrahim, and B. Vallis, "How the internet reacted to covid-19: A perspective from facebook's edge network," in *Proceedings of the ACM Internet Measurement Conference*, IMC '20, (New York, NY, USA), p. 34–41, Association for Computing Machinery, 2020.

[9] T. Favale, F. Soro, M. Trevisan, I. Drago, and M. Mellia, "Campus traffic and e-learning during COVID-19 pandemic," *Computer Networks*, vol. 176, p. 107290, July 2020.

[10] M. Candela, V. Luconi, and A. Vecchio, "Impact of the covid-19 pandemic on the internet latency: A large-scale study," *Computer Networks*, vol. 182, p. 107495, 2020.

[11] E. Bembaron, "La 5g représente "moins de 1% du trafic", affirme bouygues telecom," 2021.

[12] U. Paul, A. P. Subramanian, M. M. Buddhikot, and S. R. Das, "Understanding traffic dynamics in cellular data networks," in *2011 Proceedings IEEE INFOCOM*, pp. 882–890, 2011.

[13] EU General Data Protection Regulation (GDPR), "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)," 2016.

[14] Ericsson, "Ericsson Mobility Report," tech. rep., November 2019.

[15] J. H. Ward Jr, "Hierarchical grouping to optimize an objective function," *Journal of the American statistical association*, vol. 58, no. 301, pp. 236–244, 1963.

[16] P. J. Rousseeuw, "Silhouettes: A graphical aid to the interpretation and validation of cluster analysis," *Journal of Computational and Applied Mathematics*, vol. 20, pp. 53–65, 1987.

[17] J. C. Dunn, "A fuzzy relative of the isodata process and its use in detecting compact well-separated clusters," *Journal of Cybernetics*, vol. 3, no. 3, pp. 32–57, 1973.

[18] H. Zang and J. C. Bolot, "Mining call and mobility data to improve paging efficiency in cellular networks," in *Proceedings of the 13th Annual ACM International Conference on Mobile Computing and Networking*, MobiCom '07, (New York, NY, USA), p. 123–134, Association for Computing Machinery, 2007.

[19] F. Calabrese, G. Di Lorenzo, L. Liu, and C. Ratti, "Estimating origin-destination flows using mobile phone location data," *IEEE Pervasive Computing*, vol. 10, no. 4, pp. 36–44, 2011.

[20] A. Furno, M. Fiore, R. Stanica, C. Ziemlicki, and Z. Smoreda, "A tale of ten cities: Characterizing signatures of mobile traffic in urban areas," *IEEE Transactions on Mobile Computing*, vol. 16, no. 10, pp. 2682–2696, 2017.