POLITECNICO DI TORINO
Repository ISTITUZIONALE

Understanding the Present and Future of Cellular Networks through Crowdsourced Traces

Original
Understanding the Present and Future of Cellular Networks through Crowdsourced Traces / Malandrino, Francesco; Chiasserini, Carla Fabiana; Kirkpatrick, Scott. - STAMPA. - (2017), pp. 1-9. (Intervento presentato al convegno IEEE WoWMoM 2017. IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks tenutosi a Macao (China) nel June 2017) [10.1109/WoWMoM.2017.7974294].

Availability:
This version is available at: 11583/2667103 since: 2017-11-02T10:06:21Z

Publisher:
IEEE

Published
DOI:10.1109/WoWMoM.2017.7974294

Terms of use:

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

Publisher copyright

(Article begins on next page)

12 October 2024

Understanding the Present and Future of Cellular Networks through Crowdsourced Traces

Francesco Malandrino
Politecnico di Torino
Torino, Italy

Carla-Fabiana Chiasserini
Politecnico di Torino
Torino, Italy

Scott Kirkpatrick
The Hebrew University of Jerusalem
Jerusalem, Israel

***Abstract*—We focus on today's LTE systems and use real-world, crowdsourced traces to understand (i) how present-day LTE networks are deployed and to which extent they are suited to the current traffic load; (ii) how well they will withstand the traffic demand forecasted within 2020; (iii) which techniques to improve them should be pursued and how aggressively. To this end, we use two datasets, coming from WeFi and OpenSignal and available under commercial terms. We find that today's networks are composed of tangled, medium to large-sized cells, characterized by fairly high interference. Also, current networks are typically overprovisioned, but the future traffic load will pose a significant strain on them. To accommodate the forecasted mobile traffic, our study highlights the efficacy of: (i) traffic offloading for pedestrian and stationary users, (ii) increasing the available bandwidth through, e.g., spectrum refarming, (iii) mitigating interference and improving link quality for edge users through coordinated downlink transmissions. By putting in place these actions, only a negligible amount of additional cellular infrastructure will be required. Our results come from the combination of real-world traces, experimental measurements, and ITU-recommended propagation models. Each step we take is backed by real-world facts and data.**

I. Introduction

In the coming years, there will an impressive growth in mobile data due to the ever increasing usage of mobile devices such as smartphones and tablets, and to the popularity of mobile cloud services. This is a serious challenge for mobile network operators (MNO), who risk to see their networks choke with data traffic. Fortunately, several strategies for technology improvement are foreseen in the LTE-Advanced (LTE-A) specifications [1], as well as in the on-going design of next-generation mobile networks [2]. Among these are network densification through multiple cell tiers (macro, micro, femto), the use of different access technologies (e.g., cellular, WiFi), coordinated downlink transmissions for boosting link quality, and the exploitation of additional spectrum portions. The result is that MNOs are faced with a plethora of choices to revamp their current network systems.

In this context, one may wonder *where* mobile networks currently stand with respect to the requirements posed by today's mobile traffic demand, and *what* will be actually needed in the near future to cope with the expected growth in mobile data. Understanding these issues would also help to design future cellular networks and guide MNOs towards the most effective and profitable solutions.

In this paper, we address these issues by analyzing real-world, crowdsourced traces of mobile traffic, coming from WeFi and OpenSignal [3] and collected in urban areas of the United States. Although we have considered other cities such as Los Angeles, Boston and Atlanta, for the sake of brevity, here we present our analysis in the case of San Francisco, California. We focus on the *LTE* technology and, exploiting the information provided by the above traces, we develop a methodology to characterize present LTE networks and understand their future.

Our study is challenging for three reasons. First, we need to clean and process large cellular traffic datasets and carefully combine the two traces we work with. In this way, data can be cross-checked where the two traces overlap and complemented wherever possible. Second, the information we derive from the above traces has to be integrated with suitable signal propagation models, experimental data, FCC license records and known facts about real-world cellular networks, in order to obtain a full-fledged, reliable representation of the system. To make things worse, cells that appear in the traces across the geographical area have different characteristics (e.g., macro/micro-cells) thus requiring the model to be tailored accordingly. Third, when studying the future development of cellular networks, the diversity of factors that affect the system performance and of strategies foreseen for throughput increase, ask for as many different solutions to be considered and assessed.

The main contribution of this paper is indeed the methodology that we adopt in order to cope with all the above issues. Importantly, our methodology is general and each of its steps is backed by real-world facts and data. The main steps of the procedure are illustrated in Fig. 1, along with the input each step requires and the answers it yields. In particular, steps 1–3 reveal how *present-day* LTE networks are deployed and utilized. Steps 4–6, instead, explore to which extent these networks will be able to withstand their load in the *near future*, i.e., in the 2016–2020 time span, and which enhancements will be necessary given the forecasted growth of mobile cellular traffic. The motivation behind the choice of such time span is that reliable forecasts on traffic trends are typically limited to the next four years [4]. In addition, beyond 2020, 5G is expected to take over, bringing about major (but not yet well defined) changes to the mobile network technology. It follows that in this work we will consider those enhancements that are or will be available by 2020. However, the insights we provide into cellular networks and their needs, as well as into the potentiality of different technological strategies,

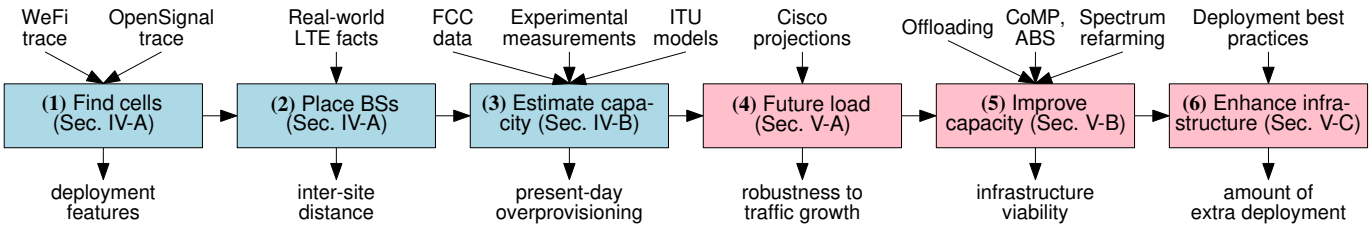

Fig. 1. The processing steps we perform (represented by boxes), the input data they use (arrows entering the boxes), and the information we obtain (arrows exiting from the boxes). Blue boxes deal with today's networks, while pink boxes refer to their future evolution.

represent useful guidelines for the development of 5G systems too [5]. At a more abstract level, our study improves the ability of *all* researchers, including those lacking strong ties with mobile operators, to understand how mobile networks are built, operated, and evolved.

Our study is unique in three main ways:

- its reliance on real-world, large-scale, crowdsourced traces that are available (albeit for a fee) to the general public, as opposed to information provided by mobile operators under non-disclosure agreements;
- its methodology, which easily generalizes to other information sources and network types;
- its taking into account both the present and the forecasted traffic load, thus helping to understand the strategies currently adopted by MNOs and how they could be enhanced in the future.

The rest of the paper is organized as follows. Sec. II discusses previous work, while Sec. III presents the traces that we use in our study. Sec. IV characterizes current LTE networks and their load. The future development of such networks is analyzed in Sec. V, which provides useful insights and guidelines for the design and enhancement of LTE networks. Finally, Sec. VI concludes the paper highlighting directions for future research.

II. Related work

Our work is mainly related to studies on mobile network planning and enhancement, and to the body of work analyising real-world measurements of cellular network traffic.

Out of the vast literature existing on network planning, the works [6], [7] aim at optimizing the deployment of LTE base stations considering both coverage and capacity. A similar problem is addressed in, e.g., [8] in the case where multiple operators share the network infrastructure. The goal of our study is fundamentally different from all these works: we aim not to optimize infrastructure deployment or sharing, but to develop a methodology to characterise real-world LTE networks and study their performance and potential compared to the present and future traffic demand.

As far as capacity increase of mobile networks is concerned, several studies have focused on physical-layer techniques that can enable cellular networks to meet the growing traffic load. These include Coordinated Multipoint (CoMP), mmWave communications and MIMO [9]–[11]. Relevant to our work are also the studies on traffic offloading, such as [12], [13]. We refer to the above works in order to get input values on the performance gains that these techniques are expected to yield.

Finally, several papers have appeared on the analysis of cellular traffic data traces, tackling different aspects. As an example, mobile traffic patterns of cellular towers are modeled through an empirical study in [14], while the geospatial and temporal dynamics of mobile traffic are studied in [15], [16]. User mobility and temporal activity patterns, as well as the usage of radio resources by different applications, are studied in [17], [18] for 3G networks. In [19], the aggregate temporal behavior of calling activity in a mobile phone network is used to infer daily mobility patterns in an urban area. Within this body of work, the spirit of [20] is the closest to that of our study. Indeed, [20] characterizes the operational performance of a 1-tier cellular network during high-profile crowded events, the experienced performance degradation in user service, and possible remedies.

TABLE I
The WeFi and OpenSignal San Francisco datasets

Metric	WeFi	OpenSignal
Time of collection	March 2016	
Number of records	9 millions	2 millions
Unique users	7,182	n/a
Unique Wi-Fi BSSIDs	21,196	5,890
Mobile operators (number of cells)	MNO 1 (7,998) MNO 2 (6,154) MNO 3 (2,338)	MNO 1 (2,123) MNO 2 (2,104) MNO 3 (1,294)

III. Input data

We begin our analysis from two *crowdsourced* mobile network traces, coming from the WeFi and OpenSignal apps [3], respectively. In particular, we consider the traces related to the city and county of San Francisco, corresponding to a geographical area of about 11×11 km^2. This is a dense urban environment, challenging for any MNO to adequately serve. We focus on three major, US-wide MNOs, hereinafter randomly labeled from 1 to 3. Tab. I summarizes the main features of both datasets.

a) WeFi: WeFi collects information about the user's position, connectivity and activity. Each record within the dataset contains the following information:

- day, hour (a coarse-grained timestamp);

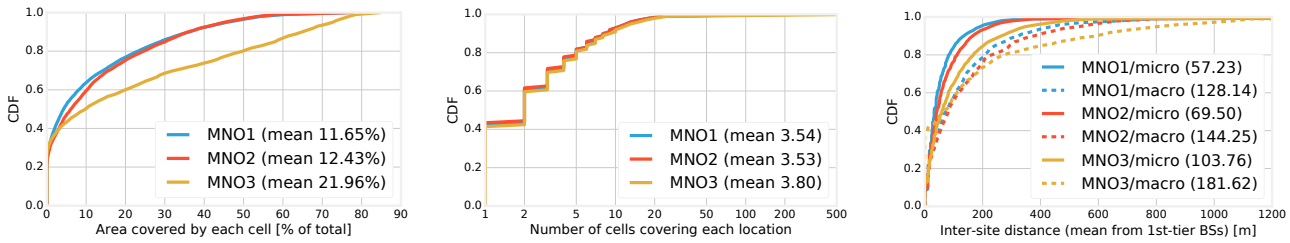

(a) (b) (c)

Fig. 2. Distribution of the area covered by each cell (a); number of cells covering each location (b); inter-site distance for macro and micro-cells (c). In (a), covering 10% of the total area means a cell radius of roughly 2 km, 20% 2.8 km and 50% 4.5 km. Legends also report mean values within parentheses.

- anonymized user identifier and GPS position;
- MNO, cell ID[1], cell technology (e.g., 3G/4G) and local area (LAC) the user is connected to (if any);
- Wi-Fi network (SSID) and access point (BSSID) the user is connected to (if any);
- active app and amount of downloaded/uploaded data.

If the position of the user or the networks she is connected to change within a one-hour period, multiple records are generated. Similarly, one record is generated for each app that is active during the same period.

The fact that records in the WeFi trace include mobility information allows us to track how much each user moves over time, and therefore to categorize users as static, pedestrian, or vehicular.

b) OpenSignal: The objective of OpenSignal is to construct a publicly available, operator-independent map of worldwide connectivity. To this end, users of the OpenSignal app volunteer to share their position and connectivity information, both cellular and Wi-Fi. Furthermore, users can decide to run speed tests, whose outcome – upload/download speed and latency – further enriches the map.

As highlighted in Tab. I, OpenSignal data includes neither user identifiers nor application information. Furthermore, due to its smaller user base, it does not report some cells and Wi-Fi access points that are reported by WeFi. Notice that instead WeFi reports over 97% of the cells reported by OpenSignal, coming very close to being a superset thereto. Thus, we mostly base our study on the WeFi trace, owing to the larger amount of information it contains. We will use the OpenSignal trace whenever appropriate to complement and cross-check our results and observations.

c) Availability and reproducibility: Compared to traditional traces collected by MNOs, the datasets we use enable a more comprehensive vision of mobile networks, spanning different technologies (Wi-Fi and cellular) and multiple mobile networks. Another, non-technical, advantage is that our datasets are collected by commercial companies and are available under commercial terms. This makes our work easier to reproduce, and our findings easier to generalize. All the code needed to generate the results presented in this paper is available online at [21].

IV. A DATA-DRIVEN LOOK AT LTE NETWORKS

Our purpose here is to use the information at our disposal to study (i) the deployment of present-day LTE networks, (ii) the load they serve, and (iii) to which extent the first suits the latter. To this end, we perform the steps 1–3 in Fig. 1.

A. Network deployment

Let us first consider the number of cells that appear in the trace for each MNO, as per Tab. I. We note that such a number is fairly high considering the geographical extension covered by the trace. We then look at the size of the cells, expressed as the fraction of the total area they cover. We (conservatively) assume the coverage area of a cell to be the *convex hull* of all locations from which users report being covered by the cell itself (i.e., they report the corresponding Cell ID). The results are presented in Fig. 2(a), which shows a quite high number of large cells. More than 50% of all cells cover over 10% of the whole area under study, and the coverage of the 10% biggest cells reaches (for MNO 1 and MNO 2) or exceeds (for MNO 3) half of the whole area. Recall that we are looking at 11×11 km^2, so a cell covering half of this surface has a radius of 6.5 km – fairly commonplace for LTE macro-cells, even in urban scenarios. Also, a 10% coverage translates into roughly 2 km cell radius, thus MNOs have between 50% and 60% small/medium sized cells in their networks.

Since there are so many cells (see the last row of Tab. I) and they are fairly large, the resulting coverage is very dense. As it can be seen from Fig. 2(b), 20% of all locations[2] are covered by more than 5 cells, and it is not uncommon to find areas covered by as many as 10 cells. Importantly, similar observations hold for both the WeFi and the OpenSignal trace. Fig. 3 shows which zones exhibit a denser deployment: as expected, they turn out to be downtown areas (e.g., the financial district in the north-east) and the main thoroughfares (e.g., Market street immediately south of the financial district).

In the following, we take the widely accepted [22], [23] value of 2 km as a watershed[3] between macro and micro-

[1]Cell IDs uniquely identify each cell within the MNO's network. They are not to be confused with physical cell IDs, i.e., integer numbers in $[0, 503]$ used for data scrambling on control channels.

[2]In order to show location-based results, the geographical area has been discretized by superimposing a 10-m grid thereto.

[3]We remark that 2 km is the value widely considered (see, e.g., [22], [23]) as *maximum* radius of micro-cells.

Fig. 3. Number of cells covering each location for MNO 2; lighter areas correspond to denser coverage. Maps for other operators (omitted for brevity) show a similar behavior.

cells: cells whose range exceeds 2 km are classified as macro-cells, while the others (about 50-60%) are micro-cells. As we will see in Sec. IV-B, this choice is also backed by the traffic service data available in the WeFi dataset. However, it is important to stress that the macro and micro-cells layout that emerges from our traces does not resemble a typical 2-tier deployment, rather smaller cells often span across the coverage of several macro-cells leading to a quite tangled structure.

Next, we need to determine the location of the base stations (BSs) serving each cell (step 2 in Fig. 1), a piece of information that is not present in our datasets. Real-world macro-cells typically employ tri-sectoral sites – a fact that is captured in the system models recommended by ITU [24] and confirmed by the shape of most of our macro-cells. Therefore, in the case of macro-cells, we place groups of three BSs at one corner of the coverage area of as many cells, specifically, the one that minimizes the average distance from the covered users.

Microcells, on the other hand, are known to employ both directional and omnidirectional antennas. For each cell, we compute a *roundness* metric, defined as [25] $4\pi\frac{A}{P^2}$, where A is the size of the cell coverage area, and P is the perimeter thereof. The metric takes value 1 for circles, and 0 for segments. We then assume an omnidirectional antenna at the center of the coverage area for the cells with a roundness exceeding 0.5, and a sectoral antenna for the others.

We then compute the inter-site distance for macro and micro-BSs, expressed as the average distance of a macro (resp. micro) BS from its first-tier neighboring macro (resp. micro) BSs. The results are depicted in Fig. 2(c). Consistently with the high number of cells, the inter-site distance is quite short for both macro and micro-BSs, and the mean values are in agreement with those characterizing *dense* deployments in 5G systems (namely, 200 m for macro-cells in urban scenarios and 50 m for micro-cells [26]). This is an interesting fact: densification is commonly thought of as a future trend, that will come to maturity as small cells (including femto-cells) become commonplace. Also, it is usually foreseen in two-tier scenarios, such as those recommended by 3GPP [27]

TABLE II
FREQUENCIES ASSIGNED TO MNOS. SOURCE: FCC [28]

Operator	700 MHz	1700 MHz	1900 MHz
MNO 1	12 MHz, band 17	—	15 MHz, band 2
MNO 2	5 MHz, band 12	15 MHz, band 4	—
MNO 3	10 MHz, band 13	15 MHz, band 4	—

for performance evaluation of cellular networks. Our data instead suggests that densification is already happening, at least in urban areas, and is carried out with tangled, medium to large-sized cells. This implies that not all results obtained in simplified reference scenarios may still hold in real-world networks: verifying this is one of the goals of our study.

B. Network capacity vs. traffic load

One may rightly wonder what the capacity of such a dense network could be and how it stands with respect to today's traffic load. In order to answer these questions, we focus on downlink data transfers, which represent the greatest fraction of the traffic reported in the WeFi trace and are deemed to be the most critical component also in the future [4]. We approach the above nontrivial task as follows:

1) computing the attenuation between geographical locations in the topology and BSs;
2) computing the signal-to-interference-plus-noise ratio (SINR) at each location in the WeFi trace, from every BS covering the location;
3) mapping the SINR onto the throughput associated with the LTE radio resource unit, i.e., a resource block (RB);
4) validating the results at the locations reported in the WeFi trace, against the traffic volume received by the users from their serving BS, and evaluating network overprovisioning.

a) Signal propagation: The first step is accomplished by exploiting the ITU models recommended for LTE networks serving urban areas [24]. We remark that the models and the parameters we set are also in line with those foreseen for 5G urban environments [26].

- Micro-BSs, line-of-sight (LOS): $PL_{dB}{=}40\log d{+}7.8{-}18\log h_{BS}{-}18\log h_{UE}{+}2\log f$;
- Micro-BSs, non line-of-sight (NLOS): $PL_{dB} = 36.7\log d + 22.7 + 26\log f$;
- Macro-BSs NLOS: $PL_{dB} = 22\log d + 28 + 20\log f$.

In the equations above, f is the frequency, d is the distance between BS and user, h_{BS} and h_{UE} are the antenna heights of, respectively, BSs and users. We set $h_{BS} = 25$ m for macro-BSs, $h_{BS} = 10$ m for micro-BSs, and $h_{UE} = 1.5$ m [24], [27]. Following [24], we consider only the NLOS model for macro-BS and we use the LOS expression for micro-BS with probability $P_{LOS} = \min\left\{1, 18d^{-1}\right\}\left(1 - e^{-\frac{d}{36}}\right) + e^{-\frac{d}{36}}$.

Our datasets do not include the frequencies used by each BS (parameter f in the equations above). We look for this information in the FCC license database [28]: as summarized

in Tab. II, we find that all MNOs use frequencies at 700 MHz, and then some more at either 1700 or 1900 MHz. These values naturally map into large and medium-sized cells, respectively: in the following, we assume that macro-cells use 700-MHz frequencies, while micro-cells use higher frequencies (whichever are available to their owner). It is worth stressing that FCC licenses can have a limited geographical scope, e.g., a single state or county. Tab. II only includes those licenses whose scope includes San Francisco; licenses valid for other areas are not included therein.

As aimed at by LTE MNOs, we initially assume frequency reuse factor of 1, i.e., all macro (micro, resp.) BSs use all the frequencies available to an MNO for macro (micro, resp.) cells. Also, in line with [24], [27], we assume a transmission power of 43 dBm for macro-cells, and 30 dBm for micro-cells. Using such a model, we can then compute the SINR that is experienced at each geographical location.

b) Matching SINR with service data: We now need to validate our signal propagation model by using the information included in the WeFi trace on the traffic volume served to the users. To this end, we first map the SINR experienced by a user at a given location onto the amount of data that can be carried by one RB. We use experimental measurements [29] collected in the case of 2×2 MIMO – a fairly common setting in LTE networks –, and obtain the per-RB throughput. The number of available RBs is computed using the frequency allocation in Tab. II. Then, in line with real-world LTE systems, we consider that proportional-fair scheduling is in place and obtain the throughput that can be offered at each location. Importantly, the experimental data in [29] shows that, in order to have a BS-user data communication, the SINR should be above -10 dB (which is also in accordance with the figures reported in [30]).

Fig. 4(a) depicts the distribution of the SINR for user-BS pairs that, in the WeFi trace, exchange data. The dashed lines therein refer to the case where we apply the path-loss equations to our data and set the frequency reuse factor to 1. We see that over 50% of communications that we observe in the WeFi trace are deemed impossible by our model, having SINR lower than -10 dB. This is a consequence of the dense deployment, which, under frequency reuse factor equal to 1, yields a very high interference. Note that decreasing the radius value taken as watershed between macro and micro-cells only worsen the situation (these results have been omitted for brevity). We therefore need to refine our model, in order for it to match the service coverage that emerges from the WeFi trace. Specifically, we need to account for interference mitigation techniques on the data plane, which in today's systems[4] mainly consist of flexible frequency reuse.

To this end, we relax the assumption on frequency reuse factor being equal to 1 and make *local*, per-BS decisions on which frequency bands to use. We adopt a hill-climbing approach, starting from those areas where users experience the lowest SINR. Then, given an area and initially setting the reuse factor to 1, we consider the individual BSs therein, starting with the ones with larger coverage. If we find it beneficial, we increase the reuse factor K so that the BSs will use only a fraction $1/K$ of the available frequencies thus reducing the interference towards neighboring BSs and, at the same time, their own capacity. We found that, in order to match the downlink service data reported in the WeFi trace, K should be typically increased to 3 for 8–18% of micro-cells and 40–48% of macro-cells, depending on the MNO. The high number of macro-cells and the fact that also micro-cells were involved, reflect the dense and tangled deployment we observe.

The final result is shown by the solid lines in Fig. 4(a), where the SINR at virtually every served location is above -10 dB. This means that the SINR that our model yields is in substantial agreement with the data transfers we observe from the trace. The fact that a few user-BS pairs still have a low SINR tells us that our model is slightly more conservative, a desirable property since, as detailed next, we are looking into worst-case, peak-time performance.

We now proceed and assess where the capacity of our networks stands with respect to their current traffic load. To do so, we need to find out the system *peak-time* load. A straightforward solution would be to consider the date and time with the highest total load, and use that snapshot as a reference. However, this would make us neglect that traffic load varies over both time and space. We thus construct a *combined* peak-load snapshot, where we consider the maximum load of each cell, and then combine together these local peak loads.

Fig. 4(b) depicts the distribution of the *pressure*, i.e., the ratio of the traffic demand to the throughput available at different locations. Consistently with the well-known fact that LTE networks are overprovisioned, pressure values average below 1%, and only exceptionally exceed 10%. Fig. 4(c) shows the moderate- and high-pressure areas for MNO 2 (maps for the other MNOs show similar results; they can be found in [21]). In accordance with common sense, we can clearly observe that downtown areas and main thoroughfares have higher pressure, and are thus more likely to become problematic in the future.

C. Summary

We performed the processing steps 1–3 in Fig. 1. Specifically, we characterized our LTE networks using WeFi and OpenSignal traces and complementing them with real-world LTE facts, experimental measurements, ITU propagation models, and FCC license records. So doing, we observed a much denser deployment than we expected (Fig. 2(b), Fig. 3), made of medium to large-sized cells (Fig. 2(a)), deployed fairly close to each other (Fig. 2(c)).

While simpler models would predict a very poor performance for such a dense deployment, we properly accounted for present-day interference mitigation techniques, obtaining SINR values (Fig. 4(a)) that are consistent with the data traffic reported in the WeFi trace. We also found that the network capacity far exceeds today's peak demand (Fig. 4(b)). As further confirmation of the correctness of our methodology,

[4]Note that the Almost-blank subframe (ABS) technique (i.e., one of the eICIC strategies) is not implemented yet in the networks of the considered MNOs. It will be considered as a way to enhance LTE networks in Sec. V.

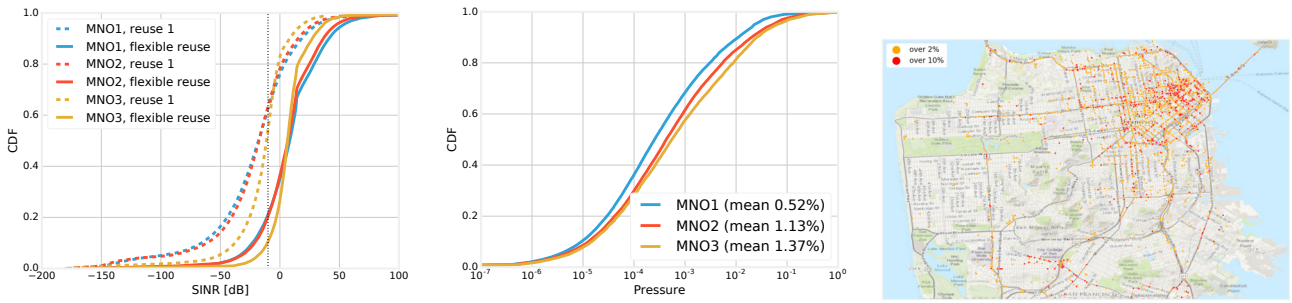

(a) (b) (c)

Fig. 4. Distribution of the SINR with (solid lines) and without (dashed lines) flexible frequency reuse (a); distribution of the pressure for different MNOs (b); locations where MNO 2 has a pressure exceeding 2% (yellow) or 10% (red) (c).

downtown areas and main thoroughfares are the locations where demand and capacity are the closest (Fig. 4(c)).

V. Enhancing and extending LTE networks

We now describe the processing steps corresponding to blocks 4–6 in Fig. 1, focused on the *future* demand and the ability of LTE networks to deal with it.

A. Future demand and pressure

Cisco [4] is a prime source of information on future network demand. The figures below[5] are especially relevant to us:

- cellular traffic from non-mobile users will grow with a 57% compound annual rate (CAGR) [4, Fig. 2];
- cellular traffic from mobile users will grow with a 61% CAGR [4, Tab. 5].

As discussed in Sec. III, the WeFi trace contains reasonably accurate location information on individual users. This allows us to mark as mobile any user that moves by more than 1 km in each one-hour period. (Needless to say, the same user can be marked as mobile in a time period and as non-mobile in others.) We multiply today's combined peak traffic, obtained in Sec. IV-B, by the CAGR figures provided by Cisco: either 1.57^5 (non-mobile users) or 1.61^5 (mobile users), thus obtaining the projected future demand. Similarly, we multiply today's combined peak 3G demand by the same factors, and add that to the future LTE load. This way we account for the current trend of user traffic migrating from 3G to LTE.

We compare the demand values we obtain to the throughput that our networks can provide, and identify the *struggling* locations, i.e., the locations wherein the former exceeds the latter. The majority of locations will be able to deal with the future traffic – as one might expect from Fig. 4(b). However, as Fig. 5 exemplifies, each MNO will have several hundreds of struggling locations, where the network capacity cannot meet the demand, and action will be needed. We remark that the reason why certain locations are struggling is essentially twofold:

[5]The data load that could come from M2M is not considered as we would need information (currently unavailable) on the fraction of this data with stringent latency constraints, which necessarily should be supported by cellular. However, our study can be easily extended to account for such traffic component.

Fig. 5. Struggling locations for MNO 2 and the reason why they struggle.

(1) the experienced SINR is low, hence each RB can carry a small number of bits and the provided data rate is not enough to support the future traffic demand. In particular, two factors contribute to a low SINR: (1.a) the location is highly interfered by neighboring BSs; (1.b) the signal received from the serving BS is weak, as in the case of cell-edge locations;

(2) the experienced SINR is satisfactory but the traffic load is exceedingly high, compared to the amount of available radio resources.

Our model, combined with the WeFi trace, reveals that, quite consistently across the different MNOs, interference (case (1.a) above) is the main cause of struggle for more than 60% of locations, along with about 37% of locations being at the cell edge (case (1.b)). Struggling locations with a good SINR – higher than 5 dB – (case (2)) amount to only few percentage points. The different reasons why locations struggle in the case of MNO 2 are represented with different colors in Fig. 5, which also highlights that struggling locations include mainly downtown areas and thoroughfares. This is in agreement with the above percentages, as these areas exhibit a particularly dense network deployment (see Fig. 3) – hence many locations therein suffer high interference –, and they are burdened with high traffic demand. Thus, their SINR is insufficient to carry the required data load, as we can see from Fig. 4(c).

Below, we first aim to "heal" struggling locations without extending the present-day network deployment (Sec. V-B).

Then, in Sec. V-C, we study where and how such a deployment needs to be complemented with new cells.

B. Enhancing the network

In order to understand how the *existing* network can be improved to cope with the future traffic load, we investigate the following three strategies:

1) traffic offloading;
2) spectrum extension;
3) SINR increase.

We cascade the above strategies starting from those that aim to accommodate the additional traffic load without acting on the SINR (i.e., traffic offloading and spectrum extension). Then we target SINR increase and consider coordinated multipoint (CoMP) to mainly heal cell-edge locations, and almost-blank subframes (ABS) to mitigate interference. The reason for such an order is twofold. First, both traffic offloading and spectrum extension are, at least partially, already in place, as demonstrated by offloading toward Wi-Fi and spectrum refarming. Thus, it is worth investigating to which extent they should be further pursued and enhanced. Second, this order turned out to give the best overall results. Indeed, CoMP and ABS increase the SINR at the expense of BS capacity; thus, they can be extensively applied when the network performance is not limited by the number of available RBs.

a) Traffic offloading: One of the simplest and most straightforward ways to deal with overburdened LTE networks is offloading traffic to other networks, with Wi-Fi being an obvious destination. To this end, we need to know:

- the existing Wi-Fi networks, and their coverage area;
- their capacity;
- how much of such capacity will be available for offloading, also considering the forecasted growth of Wi-Fi traffic.

The first piece of information is readily available from the WeFi and OpenSignal traces. About the capacity of Wi-Fi networks, we assume they all use the 802.11n technology – a fair assumption, given how fast this technology is being adopted –, with a per-access point aggregate throughput of 300 Mbit/s[6] [31], [32].

As far as the spare capacity available for offloading is concerned, we proceed as follows: (i) we increase today's Wi-Fi traffic according to the Cisco projections [4], and then (ii) we subtract the traffic generated by Wi-Fi static users, which will be served by technologies to come such as mmWave. We further assume that only static and pedestrian LTE traffic can be offloaded to Wi-Fi.

The results are summarized in the third column of Tab. III. They are very encouraging: about two thirds of struggling locations (for any reason) stop struggling as Wi-Fi offloading is enabled. This confirms the widespread belief that offloading is a remarkably effective way of easing the load of cellular networks. Many of the healed locations lie in the city center, and many of the still-struggling ones along the thoroughfares. This is consistent with the abundance of Wi-Fi networks in the first area, and their relative scarcity, as well as the presence of higher-mobility users, in the latter.

b) Spectrum extension: MNOs are already extending the bandwidth available to LTE by *refarming* some of their spectrum: they are changing the destination of some frequency bands from GSM to LTE, and the same can be foreseen for 3G. We therefore focus on refarming as our spectrum extension strategy, and assess its efficacy.

Tab. III (fourth column) reports how struggling locations fare after 5 MHz (e.g., of GSM spectrum) are refarmed to LTE for each MNO, in addition to traffic offloading. Refarming 5 MHz yields a fairly high gain, especially for MNO 1. We then try to add an extra 5 MHz (e.g., from 3G spectrum) to LTE: doubling the new spectrum available to LTE yields substantially more healed locations. This suggests that spectrum refarming is a strategy worth pursuing aggressively, however - good news - 5–10 MHz are already enough to significantly improve the network performance.

We remark that similar qualitative conclusions can be drawn when refarming is applied to the whole set of struggling locations, i.e., in absence of traffic offloading. However, cascading these strategies as done in Tab. III allows us to alleviate the traffic pressure at struggling locations through traffic offloading in the case of static/pedestrian users, and through spectrum refarming in the case of higher-mobility users, who could not be served efficiently by other networks.

At last we stress that, in spite of the above efforts, Tab. III reports a significant number of locations that are still struggling. These are the locations affected by very low SINR, compared to the forecasted traffic requirements. Refarming the spectrum means adding more RBs, but it does nothing to increase the amount of data each RB can carry – hence it may be not enough to heal locations with remarkably poor SINR. We also underline that such locations exhibit quite a high pressure already in the present (as per Fig. 4(c)), but the future increase in demand will exacerbate their situation. Consequently, below we proceed with two strategies that aim to increase the experienced SINR.

c) Coordinated downlink transmissions: Here we first[7] focus on CoMP, which, using multiple BSs to serve a single location, helps to boost the power level of the useful signal and reduce interference at the same time. CoMP is therefore particularly beneficial to cell-edge locations, many of which appear to struggle. However, other techniques such as coordinated beamforming or MIMO could be considered as well.

We assign to each struggling location one additional BS: the one from which the location receives the strongest signal, among those that (i) cover the location and (ii) have sufficient spare capacity. The results are reported in the fifth column of

[6]This is a rough, yet reliable, average value, as shown by several technical studies on 802.11n performance. Also, our study could be easily extended to the case where LTE traffic generated by static users is offloaded toward other network technologies than Wi-Fi.

[7]The order in which the techniques presented in this section are applied is the one yielding the best performance, although swapping them makes very little difference.

TABLE III

NUMBER OF STRUGGLING LOCATIONS HEALED BY THE STRATEGIES DISCUSSED IN SEC. V-B, WHEN APPLIED IN THE ORDER REPORTED BELOW (PERCENTAGES ARE GIVEN W.R.T. THE NUMBER OF LOCATIONS STRUGGLING AFTER THE PREVIOUS STEP). GREEN BACKGROUND HIGHLIGHTS THE STRATEGIES THAT HEAL OVER 40% OF STRUGGLING LOCATIONS, RED BACKGROUND THOSE THAT HEAL LESS THAN 20%. IN THE THREE RIGHTMOST COLUMNS, WE CONSIDER THAT 5-MHZ REFARMING IS ENABLED

Operator	Struggling	Healed by				Residual
		Offloading	Refarming	CoMP	ABS	
MNO 1	1083	796 (75%)	120 (42%) [with 10 MHz: 149 (52%)]	67 (40%)	35 (35%) [with K=1: 35%]	65
MNO 2	2131	1413 (66%)	190 (26%) [with 10 MHz: 226 (31%)]	209 (39%)	159 (50%) [with K=1: 49%]	160
MNO 3	1390	968 (69%)	130 (31%) [with 10 MHz: 154 (36%)]	13 (10%)	63 (25%) [with K=1: 25%]	193

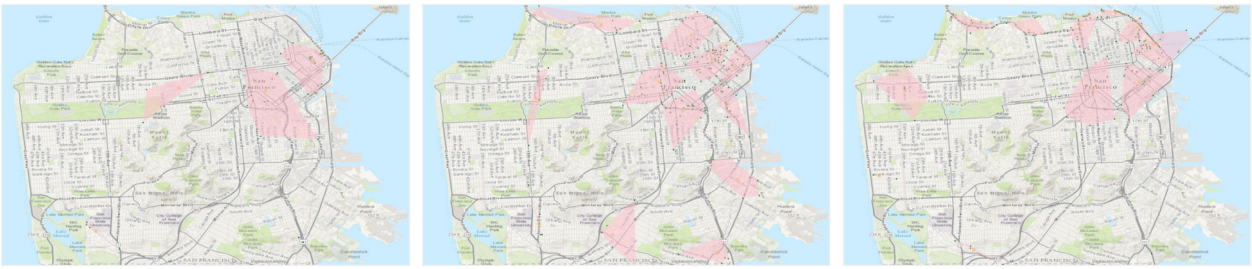

Fig. 6. New cells to deploy: 5 for MNO 1, 11 for MNO 2, and 10 for MNO 3.

Tab. III. For MNO 1 and MNO 2, CoMP heals about 40% of struggling locations. For MNO 3, instead, CoMP avails little, essentially because CoMP requires multiple BSs covering the same location, and this is less likely to happen for this MNO, as we can discern from Tab. I.

Next, we consider ABS, a technique standardized by 3GPP but not currently implemented by the MNOs. According to ABS, BSs can refrain from transmitting in some subframes[8]. In our scenario, we make per-BS decisions on whether to implement ABS or not. If to be applied, in accordance with the surveyed literature [33], downlink transmissions are muted in 25% of subframes. We proceed in a simple hill-climbing fashion, starting from the BSs causing the most interference, skipping those lacking enough spare capacity, and stopping when implementing ABS stops being beneficial.

It is important to mention that, owing to the tangled deployment of our networks with no clear distinction of roles between macro and micro-cells, we considered that any BS can perform ABS if beneficial. However, we found that less than 5% of micro-cells need to perform ABS, versus 60–70% of macro-cells. This is in accordance with the fact that this technique is foreseen mainly for macro-cells, and it further validates the distinction we operate between micro and macro-cells.

The potential of ABS to improve performance is shown in the sixth column of Tab. III, when applied on top of CoMP. ABS heals roughly 30% of the struggling locations for MNO 1 and MNO 3, and as many as 50% for MNO 2. Interestingly, although ABS was developed with classic two-tier deployments in mind, it works well also in the more tangled deployment we are observing.

Finally, we check what happens if, while enabling ABS, we disable the flexible spectrum reuse we introduced in Sec. IV-B, i.e., we set $K = 1$ for all BSs. In this case, ABS proves to be very effective: not only it makes up for the lack of flexible frequency reuse, but it also heals virtually the same number of struggling locations as before. This conforms with the notion that ABS and spectrum reuse serve mostly the same purpose in two different domains (time and frequency, respectively), and they are seldom both needed.

[8]An LTE subframe is defined as a 1-ms time period.

C. Deploying new cells

As we can see from the last column in Tab. III, even when all previous techniques are in place, a small number of locations will still struggle. Such locations are typically those that are currently served by only one BS and they are at the cell edge. We remark that, even considering a further spectrum extension, namely, a total of 20 MHz per MNO, our results (omitted for brevity) showed that several struggling locations still remain. We therefore have to take the plunge, and deploy some extra BSs to serve these unfortunate locations.

Making decisions on where to deploy additional BSs and of what kind, is a difficult task. Here, we are merely interested in getting an idea of *how many* extra cells operators would need to add, in order to heal the remaining struggling locations. To this end, we adopt a simple clustering-like approach, and group struggling locations in sets that could be served by a single micro-cell.

The resulting extra deployment (considering a spectrum extension of 5 MHz per MNO) is presented in Fig. 6, which shows that a limited number of new cells could heal all locations. While there may be different, and potentially better, ways of placing these cells – for instance, using larger cells –, our results strongly suggest that present-day LTE networks, with appropriate management and only a small addition to

their deployment, can face the challenges coming from the forecasted load increase.

D. Summary

In this section, we proposed a methodology to evaluate how LTE networks can withstand their *future* load. Our first step was to construct a conservative, worst-case snapshot of such future load, using the WeFi trace and the Cisco projections [4]. As exemplified in Fig. 5, MNOs will be unable to provide the required capacity in more than one thousand locations each.

We studied to which extent this situation can be eased by cascading traffic offloading, spectrum refarming, CoMP and ABS (Tab. III), and we found that different strategies have different impact, also depending on the reason why locations struggle. Offloading and refarming (especially, when an extra bandwidth of 10 MHz can be added) are both very effective on all locations. As we might expect, CoMP and ABS are mostly, although not exclusively, successful with cell-edge and highly-interfered locations, respectively. Finally, we have seen in Sec. V-C that MNOs will need to deploy only a small number of cells to solve the residual capacity problems.

VI. Conclusions

By leveraging two large datasets (both available under commercial terms), along with ITU propagation models, FCC license records and experimental data, we have developed a methodology to investigate current LTE networks and their ability to support today's traffic load. We then exploited projections on the growth of mobile data traffic, and evaluated how LTE networks can cope with that. We considered several strategies for performance improvement, which are or will be available within 2020, and assessed their efficacy.

Our study indicates that today's LTE networks are already quite dense, with a tangled deployment of macro and micro-cells. They will be able to cope with the forecasted traffic growth once they are enhanced with physical-layer and traffic management techniques, without requiring significant additions to the cellular infrastructure. In particular, traffic offloading is the most effective strategy, followed by spectrum increase through, e.g., spectrum refarming. CoMP and ABS significantly benefit, respectively, cell-edge and highly-interfered locations, and work best when applied after additional room has been made for the forecasted load increase.

A prominent way to extend our work is to consider different environments, including rural ones, and to investigate the practical implementation and efficacy of additional technologies, such as those that will be specified for 5G systems.

References

[1] "LTE-Advanced Pro: Part 3," *IEEE Comm. Mag.*, 2016.

[2] NGMN Alliance, "NGMN 5G white paper," https://www.ngmn.org/uploads/media/NGMN_5G_White_Paper_V1_0.pdf.

[3] http://www.wefi.com, http://www.opensignal.com,.

[4] Cisco. Visual Networking Index 2016. http://www.cisco.com/c/en/us/solutions/service-provider/visual-networking-index-vni/index.html.

[5] Qualcomm, "The 4G LTE network we have will get us the 5G network we need," 2016, https://www.qualcomm.com/news/spark/2016/02/24/4g-lte-network-we-have-will-get-us-5g-network-we-need.

[6] A. Guo and M. Haenggi, "Spatial stochastic models and metrics for the structure of base stations in cellular networks," *IEEE Trans. on Wireless Comm.*, 2013.

[7] H. Ghazzai, E. Yaacoub, M. S. Alouini, Z. Dawy, and A. Abu-Dayya, "Optimized lte cell planning with varying spatial and temporal user densities," *IEEE Trans. on Veh. Tech.*, 2016.

[8] P. D. Francesco, F. Malandrino, T. K. Forde, and L. A. DaSilva, "A sharing- and competition-aware framework for cellular network evolution planning," *IEEE Trans. on Cognitive Comm. and Netw.*, 2015.

[9] B. Mondal, E. Visotsky, T. A. Thomas, X. Wang, and A. Ghosh, "Performance of downlink CoMP in LTE under practical constraints," in *IEEE PIMRC*, 2012.

[10] G. C. Alexandropoulos, P. Ferrand, J. m. Gorce, and C. B. Papadias, "Advanced coordinated beamforming for the downlink of future LTE cellular networks," *IEEE Comm. Mag.*, 2016.

[11] C. Dehos, J. L. Gonzlez, A. D. Domenico, D. Ktnas, and L. Dussopt, "Millimeter-wave access and backhauling: the solution to the exponential data traffic increase in 5G mobile communications systems?" *IEEE Comm. Mag.*, 2014.

[12] K. Lee, J. Lee, Y. Yi, I. Rhee, and S. Chong, "Mobile data offloading: How much can wifi deliver?" *IEEE/ACM Trans. on Netw.*, 2013.

[13] S. Hoteit, S. Secci, G. Pujolle, S. Wieltholter, A. Wolisz, C. Ziemlicki, and Z. Smoreda, "Quantifying the achievable cellular traffic offloading gain with passpoint hotspots," in *ACM WiMobCity*, 2014.

[14] H. Wang, F. Xu, Y. Li, P. Zhang, and D. Jin, "Understanding mobile traffic patterns of large scale cellular towers in urban environment," in *ACM IMC*, 2015.

[15] D. Willkomm, S. Machiraju, J. Bolot, and A. Wolisz, "Primary users in cellular networks: A large-scale measurement study," in *IEEE DySPAN*, 2008.

[16] M. Z. Shafiq, L. Ji, A. X. Liu, J. Pang, and J. Wang, "Geospatial and temporal dynamics of application usage in cellular data networks," *IEEE Trans. on Mob. Comp.*, 2015.

[17] U. Paul, A. P. Subramanian, M. M. Buddhikot, and S. R. Das, "Understanding traffic dynamics in cellular data networks," in *IEEE INFOCOM*, 2011.

[18] J. Yang, Y. Qiao, X. Zhang, H. He, F. Liu, and G. Cheng, "Characterizing user behavior in mobile internet," *IEEE Trans. on Emerging Topics in Comp.*, 2015.

[19] P. Dazzi, M. Dell'Amico, L. Gabrielli, A. Lulli, P. Michiardi, and M. Nanni, "Improving population estimation from mobile calls a clustering approach," in *IEEE ISCC*, 2016.

[20] M. Z. Shafiq, L. Ji, A. X. Liu, J. Pang, S. Venkataraman, and J. Wang, "Characterizing and optimizing cellular network performance during crowded events," *IEEE/ACM Trans. on Netw.*, 2016.

[21] Code and maps. https://dl.dropboxusercontent.com/s/ssqr1u6xxm4rpnq/index.html.

[22] *Small Cells - Technologies and Deployment, Second and Expanded Edition*. John Wiley & Sons, 2014.

[23] S. Peric and T. Callahan, "Challenges with Microcell Deployment & Configuration," 2013, https://www.wirelessdesignmag.com/article/2013/04/challenges-microcell-deployment-configuration.

[24] ITU. Recommendation ITU-R P.1411-1. https://www.itu.int/dms_pubrec/itu-r/rec/p/R-REC-P.1411-1-200102-S!!PDF-E.pdf.

[25] NISR. Roundness Measurements. http://www.itl.nist.gov/div898/handbook/mpc/section3/mpc344.htm.

[26] EU Project METIS-II:. Living document on 5G-PPP use cases and performance evaluation models. https://5g-ppp.eu/.

[27] 3GPP Technical Report 36.814, "Further advancements for E-UTRA physical layer aspects," 2010.

[28] FCC. License View API. https://www.fcc.gov/general/license-view-api.

[29] FP7 IP EARTH, "Energy efficiency analysis of the reference systems, areas of improvements and target breakdown," 2010.

[30] H. Holma and A. Toskala, *WCDMA for UMTS: HSPA Evolution and LTE*. John Wiley & Sons, 2010.

[31] L. Bedogni, M. Di Felice, F. Malabocchia, and L. Bononi, "Cognitive modulation and coding scheme adaptation for 802.11 n and 802.11 af networks," in *IEEE Globecom Workshops*, 2014.

[32] E. Tanghe, "Capacity simulation and analysis of an IEEE 802.11 n system in a residential house," in *IEEE EuCAP*, 2016.

[33] S. Deb, P. Monogioudis, J. Miernik, and P. Seymour, "Algorithms for enhanced inter cell interference coordination (eICIC) in LTE HetNets," *IEEE/ACM Trans. on Netw.*, 2014.