



HAL
open science

Urban-scale Cellular Offloading through Wi-Fi Access Points: a Measurement-based Case Study

Patrice Raveneau, Razvan Stanica, Marco Fiore, Sandesh Uppoor, Mathieu Cunche, Hervé Rivano, Zbigniew Smoreda

► **To cite this version:**

Patrice Raveneau, Razvan Stanica, Marco Fiore, Sandesh Uppoor, Mathieu Cunche, et al.. Urban-scale Cellular Offloading through Wi-Fi Access Points: a Measurement-based Case Study. RTSI 2015 - 1st International Forum on Research and Technologies for Society and Industry, Sep 2015, Turin, Italy. hal-01201719

HAL Id: hal-01201719

<https://hal.science/hal-01201719v1>

Submitted on 24 Sep 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Urban-scale Cellular Offloading through Wi-Fi Access Points: a Measurement-based Case Study

Patrice Raveneau*, Razvan Stanica*, Marco Fiore*[†], Sandesh Uppoor^{‡§}, Mathieu Cunche*,
Hervé Rivano*, Zbigniew Smoreda[‡]

*Université de Lyon, INRIA, INSA-Lyon, CITI-INRIA, F-69621, Villeurbanne, France

[†]CNR – IEIIT, Corso Duca degli Abruzzi 24, 10129 Torino, Italy

[‡]Sociology and Economics of Networks and Services Department, Orange Labs, F-92794 Issy-les-Moulineaux, France

[§]Sorbonne Universities, UPMC Univ. Paris 06, UMR 7606, LIP6, F-75005, Paris, France

Contact email: razvan.stanica@insa-lyon.fr

Abstract—Wi-Fi offloading is one of the most effective approaches to relieve the cellular radio access from part of the burgeoning mobile demand. To date, Wi-Fi offloading has been mainly leveraged in limited contexts, such as home, office or campus environments. In this paper, we investigate the scaling properties of Wi-Fi offloading, by studying how it would perform on a much larger scope than those considered today. To that end, we consider a real-world citywide scenario, built on data about actual infrastructure deployments and mobile traffic demand, and observe which amount of traffic could be accommodated by the existing pervasive Wi-Fi access infrastructure, were it opened to mobile users. We find that more than 80% of the mobile traffic demand in a large urban area may be easily served by Wi-Fi access points, under a wide range of system settings.

Keywords—Cellular offloading, mobile demand, open Wi-Fi, measurement data.

I. INTRODUCTION

Global mobile traffic is growing at a dramatic pace, with a 81% increase in 2013, and an expected 11-fold further rise by 2018, totalling a compound annual growth rate (CAGR) of 61% in the 2013-2018 period. In the same timespan, mobile network connection speeds promise to double, yet this will not meet the requirements of the thriving mobile demand [1].

In such a context, offloading of cellular traffic to different wireless access technologies is emerging as a very promising solution to relieve the traditional mobile network from its overwhelming load. In fact, offloading is already employed today, and, globally, 45% of total mobile data traffic was offloaded onto the fixed network through Wi-Fi or femtocells in 2013. Projections indicate that, by 2018, more than half of all traffic from mobile-connected devices, i.e., almost 17 exabytes per month, will be offloaded to the fixed network by means of Wi-Fi devices and femtocells [1].

Wi-Fi offloading is typically regarded as a solution for home and work environments, and it is mainly adopted at

individual subscribers' own places, within offices, or in collective spaces such as university campuses. However, the actual potential of Wi-Fi offloading is much larger, and could be unleashed at virtually no cost, by opening the existing Wi-Fi access point deployments to cellular users, through either loose or tight coupling of the two networks [2].

Standardization efforts in this direction are indeed ongoing: 3GPP has been developing Access Network Discovery and Selection Function (ANDSF) mechanisms to trigger handoffs across different radio access technologies [3]. Also, seminal studies have started exploring the performance of large-scale offloading, by considering a user perspective and carrying out smartphone-based measurements [4], [5]. Although such works yielded promising results, they are based on a limited set of measurements, involving 100 users at most, and only rely on received Wi-Fi signal strength information, thus neglecting upper layer operation and the actual traffic demand [6].

In this paper, we take a different approach to investigate how Wi-Fi offloading would scale to citywide scenarios. Specifically, we focus on one case study, represented by the metropolitan area of Lyon, a 2.2-million inhabitants region in Southeastern France. We retrieve information about the actual deployments in the area of (i) cellular base stations of one of the four French mobile operators, and (ii) Wi-Fi access points distributed by a fixed telecommunication operator to its residential customers. These deployments are described and analysed in Sec. II of the paper. In addition, we collect information on the real-world mobile demand generated by 1.3 million users in the Lyon conurbation, for a typical 48-hour period. This dataset is presented in Sec. III of the paper.

The measurement data above allow us to perform a quantitative study of the percentage of mobile traffic in the region that could be accommodated by the existing residential Wi-Fi access infrastructure, were it opened to mobile users. Results are shown in Sec. IV of the paper. Finally, Sec. V summarizes the conclusions of our work.

II. CELLULAR AND WI-FI INFRASTRUCTURES IN LYON

We first present the deployments of cellular base stations (BS) and Wi-Fi access points (AP) in the Lyon metropolitan

This work was partially supported by the French National Research Agency under grant ANR-13-INFR-0005 ABCD.

This work was supported by the LABEX IMU (ANR-10-LABX-0088) of Université de Lyon, within the program "Investissements d'Avenir" (ANR-11-IDEX-0007) operated by the French National Research Agency (ANR).

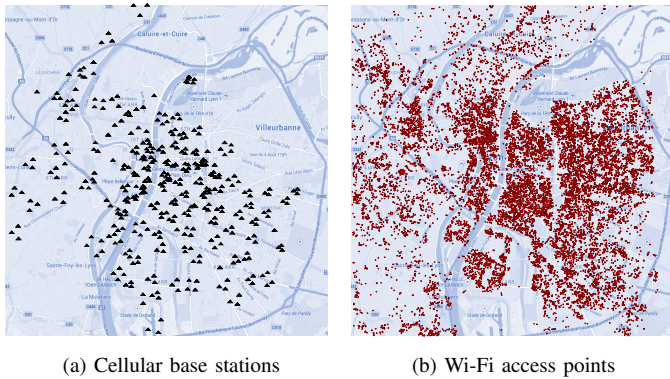


Figure 1: Real-world deployments of wireless access infrastructures in Lyon, France. A random 30% sample of the Wi-Fi access points is shown in (b), for the sake of plot readability.

area. The studied cellular network comprises close to 3,000 BS providing complete coverage of the region through 2G, 3G, and LTE access technologies. The Wi-Fi residential network includes more than 50,000 home AP of a major national network provider.

The geographical distribution of BS and AP are shown in Fig. 1a and Fig. 1b, respectively. The positioning information about the cellular BS was provided by the operator itself, while the equivalent information for the AP was estimated through a geolocation service based on BSSIDs of Wi-Fi access points [7].

The spatial distribution of both BS and AP in Fig. 1 is far from being uniform. Indeed, the BS deployment is expected to be designed to accommodate the mobile traffic demand, and thus is thicker where mobile traffic is denser. On the other hand, the AP deployment is likely to be correlated with the density of population, as one AP maps to the home location of a residential customer.

The factors driving the two access network deployments are thus different, and this is reflected by the diversity among the infrastructure distributions in Fig. 1a and Fig. 1b. An interesting question is then how correlated are the two deployments. In fact, the answer has important implications in terms of offloading capabilities, since it indirectly determines the potential for Wi-Fi offloading across the cellular network. We thus associate each AP to its closest BS, and plot in Fig. 2a the resulting AP-to-BS ratio, for each BS cell (mapped to the equivalent cell in a Voronoi tessellation). We observe that there is a lot of heterogeneity in the AP-to-BS ratio, with different BS covering from a few to several hundreds Wi-Fi AP. Also, there is no clear geographical trend, since AP-dense cells are found both in the city center and in the suburbs of Lyon.

A more rigorous view of the AP-to-BS ratio is provided in Fig. 2b, where the value is aggregated into cumulative Distribution Functions (CDF). Different curves outline the AP-to-BS ratio distribution in each *arrondissement*, i.e., the city districts of Lyon. Overall, the CDF confirm the findings above, since the slopes are fairly similar for each district. However, we note that two divergent behaviors are easily spotted. On the one hand, the curve with the steepest slope maps to a

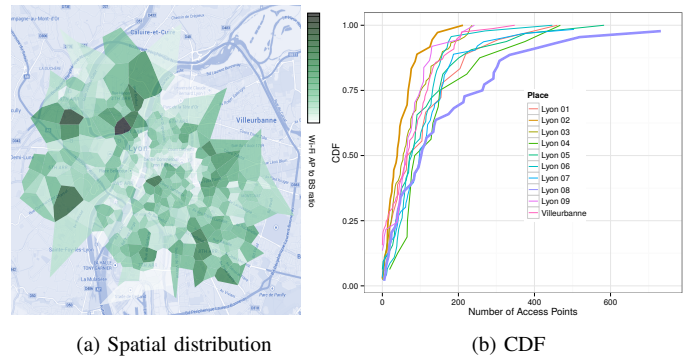


Figure 2: Relationship between BS and AP deployments. (a) Geographical distribution of the AP-to-BS ratio. (b) Per-*arrondissement* CDF of the AP-to-BS ratio.

historic neighborhood of Lyon. The area is very touristic, and the number of locals living there is significantly lower than what recorded in the rest of the metropolitan region: this limits the penetration of home Wi-Fi AP. Conversely, the district is often busy with visitors and inhabitants working in the touristic sector, which leads to a higher volume of mobile traffic and a denser deployment of cellular BS. On the other hand, the curve at the opposite extreme refers to the most residential district of the city. In this case, it is the sheer number of Wi-Fi AP that drives the AP-to-BS ratio to exceptionally high values.

Furthermore, the CDF allow performing some quantitative analysis. Specifically, we remark that there exist BS which do not have any associated AP (5% to 20%, depending on the city district), and that the vast majority of BS (70% to 99%, again depending on the district) have less than 200 associated AP. The distribution of AP-to-BS ratios among such extremes is nearly uniform in most cases.

III. MOBILE TRAFFIC DEMAND IN LYON

In order to have a realistic representation of the typical cellular traffic demand generated by mobile subscribers in the Lyon metropolitan area, we retrieved information on the 2G and 3G data traffic consumption from the operator. At the moment of the study, this mapped to around 90% of the cellular connections and more than 70% of the overall data




Figure 3: Mobile data traffic consumption. Time series.

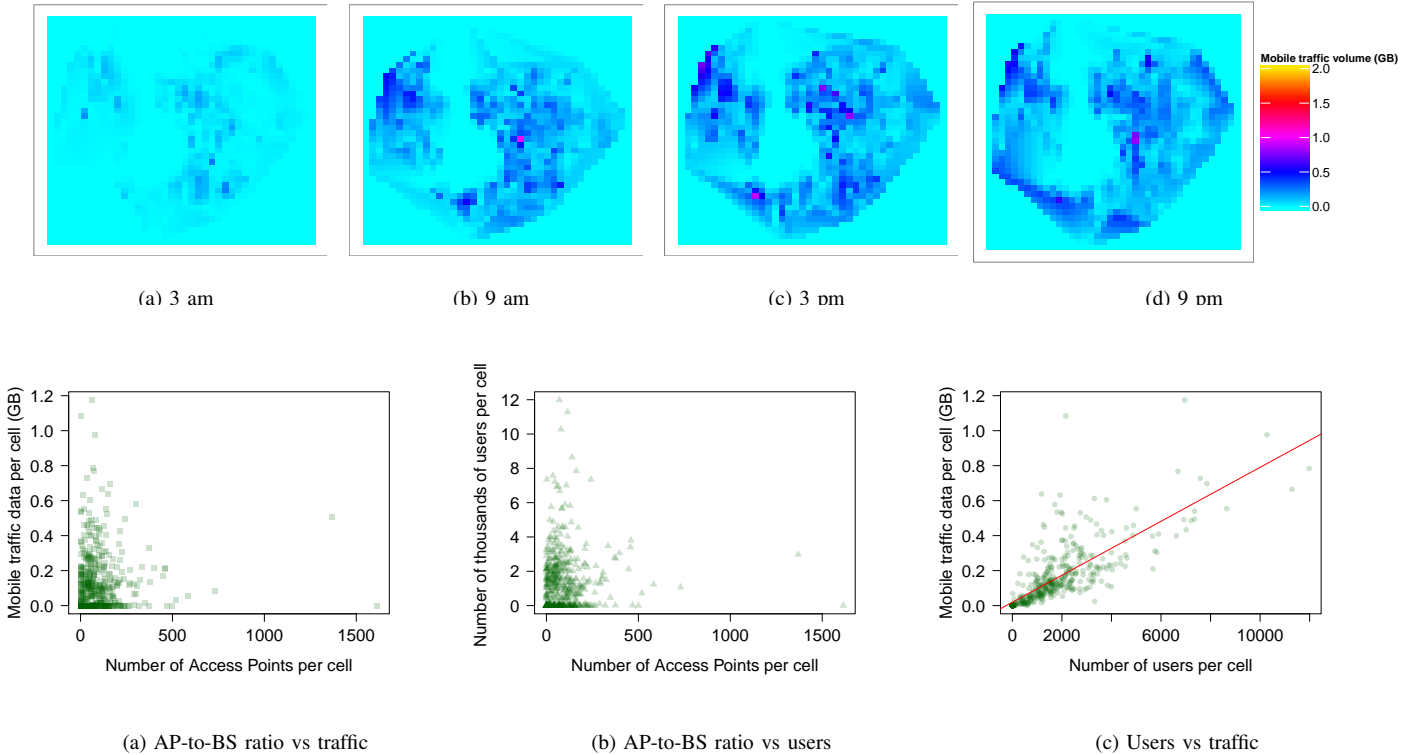


Figure 5: Scatterplots showing the relationship between the AP-to-BS ratio, the mobile traffic volume, and the number of unique users. Each dot refers to one cell, at noon on the first day considered in our study.

traffic on the cellular network, the rest being LTE data, for which no information was available. The collected statistics describe the mobile traffic, expressed in bytes, observed at every hour at each cellular BS, for a two-day period. In addition, we also recorded the number of individual mobile customers generating such traffic volumes. The spatio-temporal granularity of the information suits well our study, and it is sufficiently aggregated not to raise any concern in terms of user privacy breaches.

A. Mobile demand dynamics

We first characterize the evolution of data consumption as a function of time. Fig.3 shows the time series of the mobile demand, aggregated over the whole Lyon conurbation, and expressed in GB/hour. As we could expect, the data consumption is much higher throughout the day than at night. This daily variability, already observed in previous works on mobile traffic analysis, shows that network data consumption follows the local population activity cycle.

Fig.4 portrays instead the data consumption per cell at several hours on day one, and thus provides a representation of the geographical distribution of the mobile demand. The plots refer to significant hours that are representative of very different behaviors, i.e., 3 am, 9 am, 3 pm and 9 pm. We clearly note that the data consumption within the network heavily depends on the precise location within the city, with a spatial dispersal that varies throughout the day: thus, different areas in the city become communication hubs at diverse times of the day. When cross-comparing Fig.4a–4d with Fig.2a, we find that, quite surprisingly, the consumption per cell is not the highest in places with the densest concentration of BS.

B. Mobile traffic and network deployments

Pushed by the last observation above, we analyze if a correlation exists between mobile traffic data consumption, number of active users, and number of Wi-Fi AP. In this sense, we represent in Fig.5 the relationships between the number of Wi-Fi AP associated to a cellular BS and the amount of mobile data and the number of users associated with the same BS. Each point in these figures represents a BS, with the corresponding values for the different parameters. If a correlation exists between any of the two parameters, this would translate into a distribution of the points following the bisector line in Fig.5. From Fig.5a and Fig.5b, we can infer that there is neither a strong correlation between data and number of access points in the cell, nor between number of users and number of access points. Concerning the first association, we could have expected a link with data inversely proportional to the number of access points. Indeed, if there are a lot of access points in an area, people might use available Wi-Fi rather than 3G. The analysis revealed that no such correlation was present in the datasets.

Instead, the correlation between data and number of users per cell is strong, as shown in Fig.5c. In this case We can draw a linear fitting curve that underscores the correlation between these two metrics.

The plots in Fig.5 refer to one specific period in time, i.e., one hour around noon on the first day of our dataset. In order to generalize the result, we also explore the presence of possible correlations emerging when considering different times of the day. Specifically, Fig.6 depicts the Pearson correlation coefficient of the three bindings above, on a hourly basis, over the two days. Although the correlation slightly varies over




Figure 6: Two-day hourly Pearson correlation coefficients.

time, the results confirm that the correlation is weak for the two first associations (i.e., those involving the AP-to-BS ratio), but it is consistently strong between the mobile demand and the number of users per cell.

IV. OFFLOADING PERFORMANCE

Building on these two real-world datasets, we assess how much data could be offloaded in these settings from the cellular network to the Wi-Fi AP.

A. Assumptions

In order to evaluate the performance of offloading-based mechanisms, we need to make a number of assumptions. First of all, we consider each Wi-Fi AP has an individual offloading capacity c_{io} , which we vary between 100kb/s and 1Mb/s. These figures might seem conservative, considering current Wi-Fi technologies announce data throughputs of several hundred Mb/s. Nevertheless, we recall that we propose to use for offloading purposes residential AP, which highly reduces the available throughput. Indeed, such a residential Wi-Fi network might offer a significant capacity on the access side, but it is usually connected to the core network through limiting technologies, such as cable or xDSL. Moreover, a part of this capacity may already be used by the Wi-Fi AP owners.

We also point out that not all the residential AP might be available for offloading purposes. French Wi-Fi providers currently propose important incentives to users in order to share their AP, trying to increase as much as possible their network coverage and propose a public Wi-Fi service using these residential AP. However, some of the users might not be attracted by this, and keep their AP closed. For this reason, we only use a percentage of p_{AP} AP in the offloading process, where p_{AP} varies between 0% and 100%, and the participating AP are randomly picked from the AP dataset.

To assess the offloading capacity of such a residential Wi-Fi network, we need to map a part of the cellular data traffic to the Wi-Fi AP. However, the information we have regarding




Figure 7: Percentage of offloaded data as a function of the presence and shared capacity of open AP. Full value ranges.

data consumption on the cellular network is aggregated at a cell level, for user privacy protection reasons. Lacking any information on the user distribution within a cell, we propose a simplified offloading model. For each BS i , we consider the number of associated AP N_{AP}^i , as described in Sec. II. In this case, if we denote as D_i the data consumption in BS i , we compute the offloaded data as: $D_o^i = \min(D_i, c_o \cdot N_{AP}^i)$.

Practically, we assume that an AP will be able to fully exploit its offloading capacity whenever needed. We acknowledge that our assumption ignores the impact of the geographical distribution of the users and of the AP within the cell. However, this simplified model allows us to detect some important properties and performance bounds of the cellular offloading mechanisms, as discussed below.

Finally, we consider that cellular data can only be offloaded if a Wi-Fi network is present in the area at the time the data request is transmitted. Unlike previous studies on offloading, we do not assume any delay tolerance by offloaded services.

B. Offloading results

We use the two datasets described in Sec. II to evaluate the limits of a citywide residential Wi-Fi network in terms of offloading data from the cellular network. Therefore, our metric of interest will be the percentage of offloaded data.

In this sense, Fig. 7 shows the impact of two parameters: the Wi-Fi AP offloading capacity, and the percentage of AP used for offloading. The results are aggregated both in space (over all the cellular BS) and in time (over the two days we studied). Overall, we observe that a residential Wi-Fi network could relieve the cellular network from more than 90% of its load, if a fair percentage of AP, i.e., around 20%, were open to cellular users and thus participated in the offloading process. Moreover, in such a scenario, the shared capacity of the open Wi-Fi AP does not have a significant impact, as the difference between 100 kb/s and 1 Mb/s Wi-Fi AP capacity is negligible. Overall, we conclude that opening one every five AP is enough to offload the entire data from the corresponding BS, even when AP share a low fraction of their actual capacity.




Figure 8: Percentage of offloaded data as a function of the presence and shared capacity of open AP. Zoom on the case where the percentage of open AP is low, between 0 and 2%.

The impact of the AP offloading capacity becomes noticeable only when the participation of AP is low, i.e., below 2%, as shown in Fig. 8. Still, the performance remains very good even in this case, with 2% of the AP being able to offload between 25% (at 100kb/s) and 55% (at 1Mb/s) of the entire cellular traffic.

We further delve into the analysis, by disaggregating the results over time and space. First, we focus on the temporal evolution of the offloading capacity. Fig. 9 shows the percentage of offloaded data for each hour in the two days we studied, with a varying number of usable Wi-Fi AP. The results confirm that, when 100% of AP can be used in the offloading process, the data from most of the BS can be entirely offloaded, even at peak hours. However, the offloading performance does not reach 100% even in this case, as some cells have only a few, or even no associated AP. When the percentage of usable AP is significantly reduced, i.e., falls below 10%, a significant difference can be noted between night and day hours. In the former case, the amount of data transmitted on the cellular network is reduced, and even a reduced number of AP can offload an important percentage of data. On the other hand, during the day more data is transferred, saturating the AP.

Finally, Fig. 10 provides an insight on the spatial performance of the offloading process, showing the offloading distribution per BS at noon, on the first day of the study. Once again, for a high percentage of usable AP, all the BSs can offload their entire traffic on the Wi-Fi network. For a smaller ratio of usable AP, the offloading capacity of the residential Wi-Fi network is attained, and some differences can be noticed between the different cellular BS.

To summarize, the results show that a network formed by residential Wi-Fi AP could offload up to 80% of the cellular data in the city of Lyon. Even when a reduced number of AP is used in the offloading process, the gains are considerable: 50% of the cellular data at peak hours can be offloaded by using only 5% of the Wi-Fi AP in the city.




Figure 9: Time series of the percentage of offloaded data, for a varying presence of AP, each sharing 100 kb/s.




Figure 10: CDF of the offloaded data per cellular BS at noon, for a varying presence of AP, each sharing 100 kb/s.

V. CONCLUSION

In this paper, we studied the performance bounds of Wi-Fi offloading in real-world settings, considering the deployments of a cellular network and the one of a city-scale Wi-Fi residential network in the Lyon metropolitan area. We show that, if collaboration between the two networks is possible, around 90% of the cellular data can be offloaded through the Wi-Fi network, without any delay tolerance requirements. Moreover, even when only a small number of residential AP is used (less than 2%), more than 50% of the cellular data can still be offloaded.

REFERENCES

- [1] Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2013–2018.
- [2] F. Rebecchi, M. Dias de Amorim, V. Conan, A. Passarella, R. Bruno, M. Conti, “Data Offloading Techniques in Cellular Networks: A Survey”, to appear.
- [3] 3GPP, “3GPP TS 24.312: Access Network Discovery and Selection Function (ANDSF) Management Object (MO) (Rel. 10)”, 2011.
- [4] P. Fuxjager, I. Gojmerac, H. R. Fischer, P. Reichl, “Measurement-based small-cell coverage analysis for urban macro-offload scenarios”, IEEE VTC-Spring, Yokohama, Japan, 2011.
- [5] K. Lee, J. Lee, Y. Yi, I. Rhee, S. Chong, “Mobile data offloading: How much can WiFi deliver?”, IEEE/ACM Transactions on Networking, 21(2):536–550, 2013.
- [6] S. Liu, A. Striegel, “Casting Doubts on the Viability of WiFi Offloading”, ACM SIGCOMM CellNet Workshop, Helsinki, Finland, 2012.
- [7] P. Rouveyrol, P. Raveneau, M. Cunche, “Large Scale Wi-Fi tracking using a Botnet of Wireless Routers”, SAT Symposium, Philadelphia, PA, USA, 2015.