Mobile Network Sharing Between Operators: A Demand Trace-Driven Study

Paolo Di Francesco
CTVR
Trinity College, Dublin, Ireland
pdifranc@tcd.ie

Francesco Malandrino
CTVR
Trinity College, Dublin, Ireland
malandrf@tcd.ie

Luiz A. DaSilva
CTVR
Trinity College, Dublin, Ireland
dasilval@tcd.ie

ABSTRACT

Network sharing is often hailed as a promising and cost-effective way to tackle the ever-increasing load of cellular networks. However, its actual effectiveness strongly depends on the correlation between the networks being joined – intuitively, there is no benefit in joining two networks with exactly the same load and exactly the same deployment. In this paper, we analyse the deployment and traffic traces of two Irish operators to (i) study their correlation in space and time, and (ii) assess the potential benefit brought by network sharing. Through our analysis, we are able to show that network sharing is remarkably effective in making the load more regular over space, improving the operations and performance of cellular networks.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless communication*

Keywords

Cellular traffic data; Resource sharing; Hotspot analysis

1. INTRODUCTION

These are happy days for owners and users of such mobile devices as smartphones and tablets. Interactive services, high-quality multimedia contents, involving games of all sorts are readily available at their fingertips. Said users consistently prove willing and ready to pay for higher resolution screens and better cameras, more mind-blowing apps and more interesting multimedia contents. These are happy days for device manufacturers and operators as well.

Amidst such happiness, mobile network operators have a stern challenge to face. On the one hand, the growing demand for bandwidth and capacity prompts costly infrastructure enhancements. On the other hand, having grown accustomed to services like mobile streaming and mobile video uploading did not make subscribers any more willing to tolerate an increase in their fees. While the network disruptions of 2010, caused by the first iPhones [1], are unlikely to repeat, this situation is endangering the very profitability of running a cellular network [2].

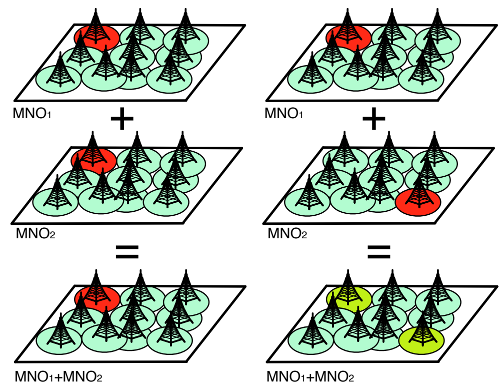

Figure 1: Network sharing: combining two networks with very similar load patterns (left case) yields little or no benefit. Instead, combining two networks with different load patterns (right case) results in a more evenly distributed load for both networks.

Operators and researchers are exploring different ways to reduce the OPEX and CAPEX associated to cellular networks, including deploying micro- and femto-cells and supporting device-to-device communications. A parallel, accelerating trend is represented by *network sharing*. It can come in the guise of joint ventures aimed at developing new network, as in the Polish and Danish cases [3], or bilateral agreements to jointly manage existing ones. The latter works in a similar way to roaming: two network operators agree to serve each other's users indifferently.

Sensible as it sounds, there are several issues that could undermine the practicality and effectiveness of network sharing. Some are related to commercial agreements or competition issues, and fall beyond the scope of our paper. Some, instead, are technical: intuitively, network sharing makes sense if the networks being joined and their demand are *different* enough. Joining two networks with very similar deployments and very similar loads has no effect on their ability to accommodate the peak load (see Fig. 1).

As mentioned, load and deployment are the foremost aspects to account for in studying the potential effectiveness of network sharing. In this paper, we leverage the data from two real-world traces, provided by two Irish network operators. Using such information allows us to assess the

CSWS'14, August 18, 2014, Chicago, IL, USA.

ACM 978-1-4503-2991-0/14/08 ...$15.00.
http://dx.doi.org/10.1145/2630088.2630090 .

practicality and the potential performance benefit of sharing capacity through real data, without the need to rely on (potentially, oversimplified) models and (potentially, unrealistic) assumptions.

The remainder of the paper is organized as follows. In Sec. 2, we review the related work. In Sec. 3, we describe our traces and their relevance to our problem. Sec. 4 contains a discussion of the temporal and spatial correlation of the network loads, and what it portends concerning the effectiveness of network sharing. In Sec. 5 we give a quantitative estimation of such effectiveness. Finally, we conclude the paper in Sec. 6.

2. RELATED WORK

There are a number of studies in the literature that aim at analysing traffic dynamics in cellular networks; they can be grouped into two categories: field measurements-based and large-scale dataset-based.

Field measurement studies have the advantage to capture the actual channel occupancy [4, 5, 6]. All these works offer interesting insights about the demand and its fluctuations. However, their foremost limitation is the difficulty to provide concurrent measurements at more than a few locations. On the other hand, studies based on large-scale datasets offer a broader view of the characteristics of a network. Such a global viewpoint has been adopted in but a limited number of papers, mainly because of the difficulty to obtain such data from network operators. One of the first attempts was made by Willkomm et al. [7], who characterised the primary usage in cellular voice networks using information from a US CDMA-based cellular operator. Such data was used to study the call arrival process, and to propose a random walk model capturing the aggregate load dynamics. In [8], Keralapura et al. analyzed the browsing behavior of mobile users in an American 3G data network, by monitoring 24 hours of IP traffic. Paul et al. in [9] looked at individual subscriber behavior and traffic patterns, studying a nation-wide 3G network at the base station level.

Our work differs from the aforementioned ones in two main ways. First, our goal is more specific: instead of characterizing the general behavior of cellular networks and their users, we aim at assessing the effectiveness of network sharing techniques . Furthermore, our work is unique in that we have access to *multiple* traces, coming from different operators: we are therefore able to check whether the deployment and load patterns are different enough to make network sharing an attractive option.

In this work, we study network sharing in a setting that is very close to present-day networks. Our results serve as a motivation and enabling factor for more advanced network sharing schemes. In [10] the authors provide a comprehensive survey of the radio access network (RAN) sharing functionality currently standardized and discussed in 3GPP, while [11] analyses feasible sharing options in the near-term in LTE. In [12] the authors introduced Network without Borders (NwoB), a new concept of wireless networks, characterised by an extreme sharing regime. Operators construct their networks in a service-oriented fashion, exchanging resources – base stations, spectrum blocks, hot spots... – from a shared pool, through a virtual marketplace. This new vision also entails a business paradigm shift [13], with operators having their role completely re-defined.

Technology	MNO_1	MNO_2	Total
3G (W-CDMA/HSPA)	5656	6679	12335
2G (GSM/GPRS)	5423	4040	9463

Table 1: Number of transmitters included in our traces, for each operator and technology.

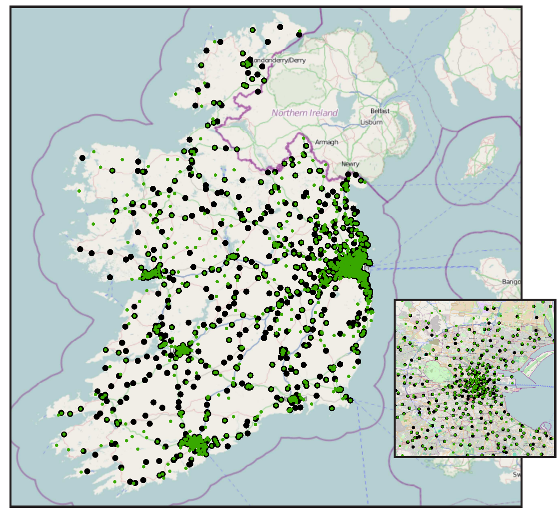

Figure 2: 3G Deployment. Dark points represent MNO_1 transmitters; light green points MNO_2 transmitters. The densely covered area in the East corresponds to Dublin (zoomed in in the box).

3. OUR TRACES

Our traces come from two Irish operators. They include a one week long call-detail record (CDR) information for both data and voice, concerning over 10,000 2G (i.e., GSM/GPRS) and 12,000 3G (i.e., W-CDMA/HSPA) transmitters distributed over the entire Republic of Ireland, as shown in Table 1.

For each transmitter, we know its position, azimuth and sectorization information, as well as its (approximate) coverage area. For each voice call and data session, we know the transmitter it is initiated and terminated at, its duration, and amount of transferred data.

Fig. 2 summarizes the nation-wide 3G deployment. We can already observe that, outside of urban centres, different operators tend to cover different areas.

3.1 Shortcomings and workarounds

Precious as they are, our traces have two main shortcomings. The first one is that they lack information on the user position and mobility, i.e., we do not know whether users move during their call or data session. We circumvent this limitation by associating each call and data session to the transmitter it is initiated at [7]. Owing to the short average duration of both calls and data sessions, this is not a significant limitation.

Furthermore, the traces come from different time periods (respectively 2011 and 2013), and therefore the magnitude of the traffic they represent changes substantially. To deal with this issue, we normalize both traces, so as to study the *fluctuations* of the demand and not its absolute value. Notice that this also bars us from reconstructing the global

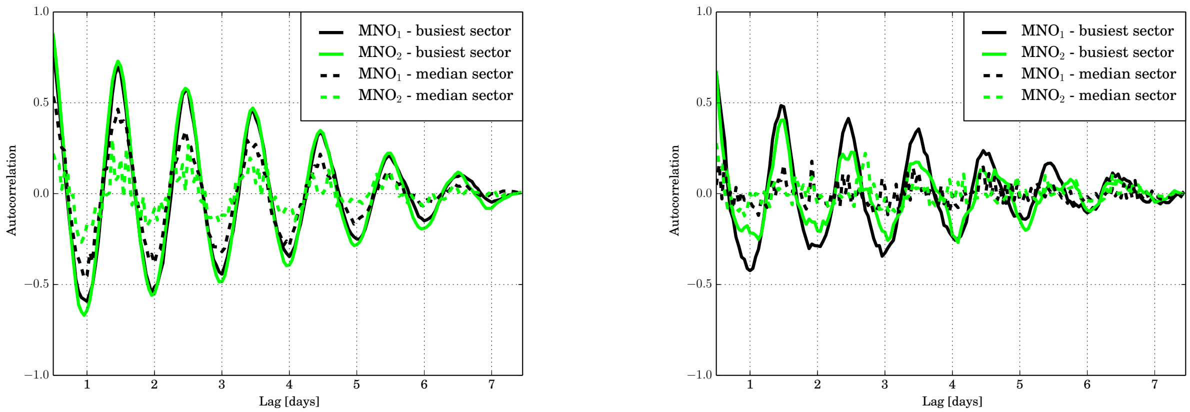

(a) (b)

Figure 3: Autocorrelation for 3G voice (a) and data (b).

load by *summing* the demand of each operator, as it would be natural to do. As we will see in Sec. 5, we resort to spatial locality analysis to work around this (potentially critical) shortcoming.

4. GLOBAL CORRELATION

As discussed in Sec. 1 and summarized in Fig. 1, network sharing is ineffective when the networks being shared are too similar to each other. In this section, we analyse the correlation between the load of MNO_1 and MNO_2, in both space and time. A high degree of correlation would mean that the potential benefit of network sharing is limited; on the other hand, a lower degree of correlation would bode well.

4.1 Time correlation

We represent the load of each sector (i.e., the area covered by each transmitter) of each operator through a time series. Their time resolution is one hour. We consider as *load* the duration of voice calls and the amount of data exchanged in data session. The traces do include the duration of data sessions, but such information is often unreliable, e.g., there are many hour-long sessions with no data exchanged. As far as normalization is concerned, we do not need any: all the metrics we will compute work with raw, unnormalized data; furthermore, as discussed in Sec. 3.1, we are not going to directly compare the two traces.

The first aspect we study is the autocorrelation of the time series, shown in Fig. 3. The shape of all curves reflects well known daily patterns: there is high positive correlation at 24-hour intervals (and, to a decreasing degree, 48-, 72-, etc.), highly negative correlation at 12-hour intervals (and, decreasing in magnitude, at 36-, 60-, etc.). Similar effects were observed in [9]. Also notice how the two operators exhibit virtually the same behavior.

What is less expected and more interesting is the sharp difference between voice (Fig. 3(a)) and data (Fig. 3(b)), with the latter having a much lower correlation. Intuitively, data traffic tends to have a more irregular time evolution; this translates into a higher probability that different operators experience different load levels at a given time. This bodes well for the effectiveness of network sharing in current networks, where most of the load is due to data rather than voice, and even more so in future ones, with additional services such as gaming and tele-presence coming into play.

Still focusing on Fig. 3(b), let us look at the difference between the busiest and median sectors: the correlation for the busiest sector is much higher. Intuitively, this suggests that the load of busy sectors follow very regular patterns, while less-used sectors have more changing loads. This is a potential issue, as busy sectors are exactly the ones that should benefit more from network sharing. We need a more clear view of how busy sectors are distributed in space, as described next.

4.2 Space correlation

Our purpose now is to understand how strong is the space correlation of the demand. In other words, if a sector is highly loaded, how likely is it that its neighboring sectors will also be highly loaded? Similarly to time correlation, space correlation is relevant to understand the effectiveness of network sharing: if busy (i.e., potentially overloaded) sectors come in large, compact clusters, then it is less likely that combining networks from different operators can do much about them.

Moran's index

Contrary to time correlation, there is no unique definition of space correlation. We employ Moran's index [14], also used in [9, 15] to study spatial aspects of network phenomena. In our context, we can define it as:

$$I_G = \frac{n}{S_0} \frac{\sum_{i=1}^{n} \sum_{j=1, j \neq i}^{n} w_{i,j} (x_i - \bar{X})(x_j - \bar{X})}{\sum_{i=1}^{n} (x_i - \bar{X})^2},$$

where n is the number of sectors, x_i represents the load of sector i and $\bar{X}$ is the average load, and

$$S_0 = \sum_{i=i}^{n} \sum_{j=i, j \neq i}^{n} w_{i,j}.$$

The weights $w_{i,j}$ represent in general the *distance weight* between two elements; in many cases, the Euclidean distance

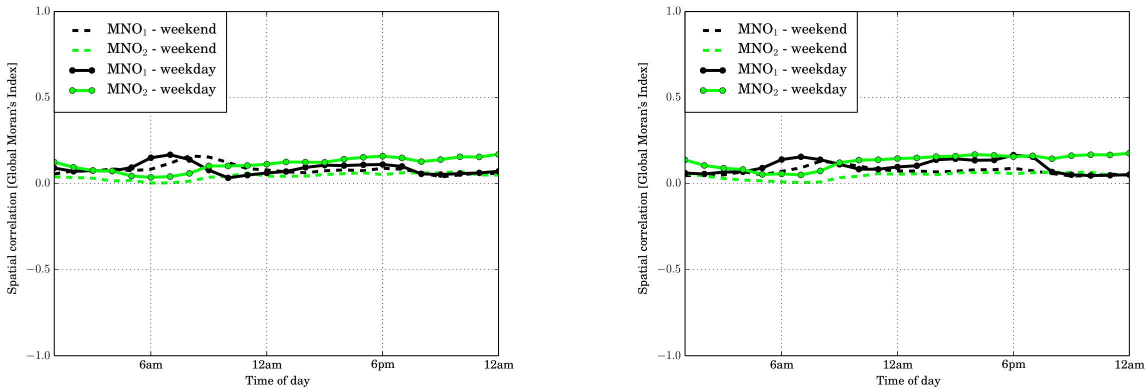

(a) (b)

Figure 4: 3G data: space correlation (Moran's index) at different times of the day and for Dublin (a) and for all of Ireland (b).

is used.

In network sharing scenarios, load can only be shared between *overlapping* sectors. Therefore, we adopt the following alternative definition of distance weight:

$$w_{i,j} = \frac{|A_i \cap A_j|}{|A_i \cup A_j|},$$

where A_i is the area covered by sector i. From our viewpoint, two sectors that do not overlap are infinitely distant from each other, as there is nothing network sharing can do about their load.

The resulting correlation is plotted in Fig. 4. We can see that it is slightly higher during weekdays and during peak hours (around 8am and 6pm). However, the most important aspect to observe is that correlation levels are always very low.

Recall [14] that Moran's index is 0 for complete spatial randomness, 1 perfect correlation, and −1 for perfect negative correlation. Our values seldom exceed 0.15, corresponding to positive but very weak correlation. We can expect that highly loaded sectors from different MNOs are not likely to overlap, thus, their load can be successfully relieved through network sharing.

5. THE EFFECTIVENESS OF NETWORK SHARING

So far, we have found several hints that network sharing is a promising way of tackling the load in cellular networks. Now, we want to go one step further, and assess *how much* we can actually gain from it.

The most straightforward way of doing this would be considering the aggregated load of the two operators, and see how well their joint networks would fare against it. The shortcomings of our traces, described in Sec. 3.1, rule this option out.

We therefore take a longer route, and start by computing the local version of the Moran's index [16]. The index for sector i is defined as:

$$\frac{x_i - \bar{X}}{S_i^2} \sum_{j=1, j\neq i}^{n} w_{i,j}(x_j - \bar{X}),$$

where

$$S_i^2 = \frac{\sum_{j=1, j\neq i}^{n} (x_j - \bar{X})^2}{n-1} - \bar{X}^2.$$

Combining the index values for neighboring sectors, we can divide them into four classes, namely:

HH high-load sectors surrounded by other high-load ones;

HL high-load sectors surrounded by low-load ones (*hot spots*);

LH low-load sectors surrounded by high-load ones (*cold spots*);

LL low-load sectors surrounded by other low-load ones.

Notice that the classification is made on a per-operator basis, i.e., as discussed in Sec. 3.1, we do not mingle together the traces of the two operators.

We are especially concerned with hot spots, i.e., sectors in class HL. These sectors are linked to the so-called *flash crowds*, i.e., groups of people sharing the same location that suddenly become interested in downloading some data. Such events are often impossible to foresee, and represent the most significant threat for the operations of cellular networks [17].

Therefore, we look for HL sectors (hot spots), that overlap with sectors of the other operator that have low load, i.e., that are in LL or LH class – just like in the right case described in Fig. 1. For these sectors, network sharing can effectively reduce the load, and thus improve the network performance.

Fig. 5 shows the number of hot spots when the MNO_1 and MNO_2 networks are operated separately or jointly, for different times of the day, during weekends and weekdays. The most important aspect to observe is the sharp decrease in the number of hot spots brought by network sharing. This holds for all times of the day, for both networks, for both weekdays and weekends: enabling network sharing invariably translates into fewer hot spots.

This is clearly very good news: as we mentioned, hot spots represent one of the most significant challenge that cellular networks have to face, and a technique as simple and cost-effective as network sharing has proven very effective in curbing it.

5.1 Broadening the focus

So far, our results have focused on the Dublin area alone. This is sensible, as Dublin is the biggest and most densely populated urban area of Ireland, and that is where overloading issues are most likely to happen. However, for the sake of completeness, we also present the number of hot spots nation-wide, and how it is changed by network sharing.

Fig. 6 shows two interesting facts. First, Dublin does not host the majority of the Irish hot spots. This is a bit counterintuitive, as Dublin does account for most of the traffic. Recall, however, that the metric defined earlier in the section is *local*; it follows that hot spots in rural areas of Ireland can be, so to speak, *colder* than ordinary sectors in Dublin. Notice that, whatever their *temperature*, hot spots always represent a problem for the network.

The second aspect that we can observe is that the effectiveness of network sharing in reducing the number of hot spots in rural areas is remarkable. Comparing the solid lines in Fig. 6 and Fig. 5, we can conclude most of the rural hot spots disappear when network sharing is enabled. This is consistent with what we would expect: hot spots are fairly uncommon in rural areas, and overlapping hot spots even more so.

Tab. 2 confirms these data. Enabling network sharing removes virtually all the hotspots in rural areas, and many of the ones in Dublin. Even in the most challenging setting, i.e., weekdays in Dublin, at least one third of the hot spots are removed.

6. CONCLUSION AND FUTURE WORK

This paper addresses the problem of the ever-increasing load on cellular networks, and the viability of network sharing as a cost-effective technique to tackle it.

After observing that sharing is only effective for networks whose deployment and load patterns are different enough, we studied such an issue with the help of two real-world deployment and traffic traces, provided by two Irish operators. We started by looking at the correlation of the demand in both time and space, finding that, especially for data, it is low enough to warrant the effectiveness of network sharing.

We moved one step further, trying to assess how much exactly we can gain through network sharing. Specifically, we

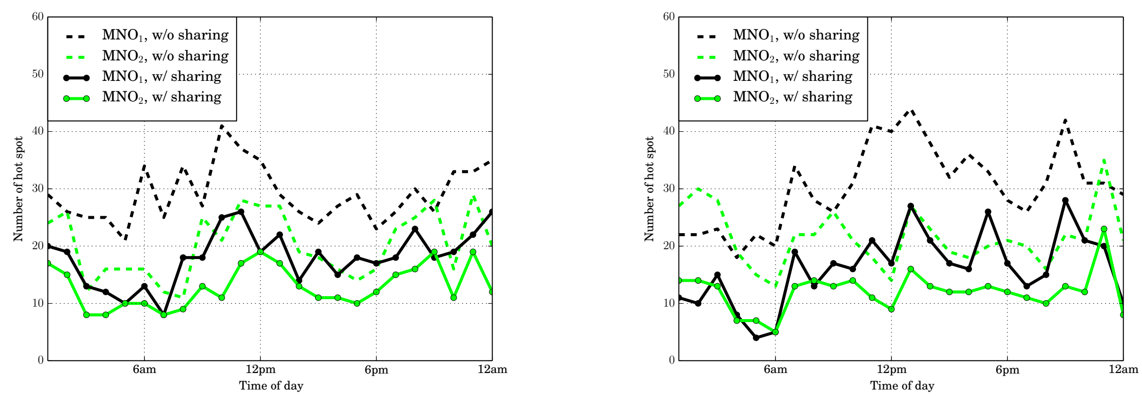

(a) (b)

Figure 5: 3G data, Dublin area: number of hot spots with and without network sharing, during weekdays (a) and weekends (b).

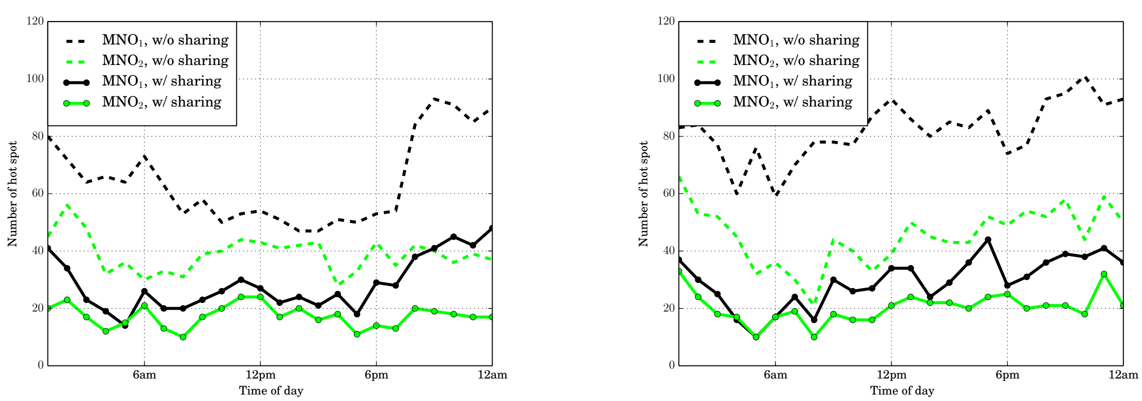

(a) (b)

Figure 6: 3G data, all of Ireland: number of hot spots with and without network sharing, during weekdays (a) and weekends (b).

		Operator	Ireland	Urban	Rural
Deployment density [sectors/km^2]		MNO_1	0.080	4.488	0.040
		MNO_2	0.095	5.615	0.042
Space correlation [Moran's Index]	we	MNO_1	0.10	0.08	0.11
		MNO_2	0.13	0.11	0.25
	wd	MNO_1	0.07	0.08	0.10
		MNO_2	0.04	0.04	0.11
hot spot reduction	we	MNO_1	-55%	-38%	-93%
		MNO_2	-55%	-35%	-50%
	wd	MNO_1	-64%	-46%	-96%
		MNO_2	-54%	-44%	-93%

Table 2: Deployment density, spatial correlation and reduction in the number of hot spots, for the whole of Ireland, urban areas (Dublin) and rural areas. Figures are differentiated for weekdays (wd) and weekends (we).

addressed the problem of hot spots, i.e., sudden spikes in the cellular load. We found that combining the networks of the two operators can greatly reduce their number. Such a reduction is massive in rural areas, and altogether remarkable in urban areas as well.

A first, natural prosecution of this work deals with *how* to actually implement and manage a shared network. There are many possible sharing regimes, and they call for appropriate managing algorithms. However, in order to do that we will consider other aspects of network sharing other than capacity, such as coverage (spatial or service specific) and cost savings (i.e. energy saving through infrastructure sharing).

Our findings can also motivate and drive the development of network virtualization schemes, where virtual wireless access networks are dynamically built to meet time- or location-specific content demands, such as the ones we detect through hot spots.

7. REFERENCES

[1] B. S. Arnaud, "iPhone slowing down the Internet – Desperate need for 5G R&E networks," 2012. [Online]. Available: http://billstarnaud.blogspot.com/2010/04/iphone-slowing-down-internet-desperate.html

[2] Credit Suisse, "U.S. wireless networks running at 80% of capacity," 2011. [Online]. Available: http://benton.org/node/81874

[3] L. Tung, "It's 4G, but they don't like to talk about it: How Denmark's LTE rose from the ashes of a price war," 2013. [Online]. Available: http://www.zdnet.com/ its-4g-but-they-dont-like-to-talk-about-it-how-denmarks-lte-rose-from-the-ashes-of-a-price-war-7000014585

[4] S. Yin, D. Chen, Q. Zhang, and M. Liu, "Mining Spectrum Usage Data: a Large-scale Spectrum Measurement Study," *IEEE Transactions on Mobile Computing*, vol. 11, no. 6, pp. 1033–1046, 2012.

[5] O. Holland, P. Cordier, M. Muck, L. Mazet, C. Klock, and T. Renk, "Spectrum Power Measurements in 2G and 3G Cellular Phone Bands during the 2006 Football World Cup in Germany," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, 2007.

[6] T. Kamakaris, M. Buddhikot, and R. Iyer, "A Case for Coordinated Dynamic Spectrum Access in Cellular Networks," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, 2005.

[7] D. Willkomm, S. Machiraju, J. Bolot, and A. Wolisz, "Primary Users in Cellular Networks: A Large-Scale Measurement Study," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, 2008.

[8] R. Keralapura, A. Nucci, Z. L. Zhang, and L. Gao, "Profiling users in a 3G network using hourglass co-clustering," in *ACM International Conference on Mobile Computing and Networking (MobiCom)*, 2011.

[9] U. Paul, A. P. Subramanian, M. M. Buddhikot, and S. R. Das, "Understanding Traffic Dynamics in Cellular Data Networks," in *IEEE International Conference on Computer Communications (INFOCOM)*, 2011.

[10] X. Costa-Perez, J. Swetina, T. Guo, R. Mahindra, and S. Rangarajan, "Radio Access Network Virtualization for Future Mobile Carrier Networks," *IEEE Communications Magazine*, vol. 51, no. 7, 2013.

[11] J. S. Panchal, R. D. Yates, and M. M. Buddhikot, "Mobile Network Resource Sharing Options: Performance Comparisons," *IEEE Transactions on Wireless Communications*, vol. 12, no. 9, pp. 4470–4482, 2013.

[12] L. A. DaSilva, J. Kibiłda, P. Di Francesco, T. Forde, and L. Doyle, "Customized Services over Virtual Wireless Networks: The Path towards Networks without Borders," in *Future Network and Mobile Summit*, 2013.

[13] L. Doyle, J. Kibiłda, T. Forde, and L. A. DaSilva, "Spectrum without Bounds, Networks without Borders," *Proceedings of the IEEE*, vol. 102, no. 3, pp. 351–365, 2014.

[14] P. A. Moran, "Notes on continuous stochastic phenomena," *Biometrika*, vol. 37, no. 1-2, 1950.

[15] A. Palaios, J. Riihijärvi, O. Holland, A. Achtzehn, and P. Mähönen, "Measurements of Spectrum Use in London: Exploratory Data Analysis and Study of Temporal, Spatial and Frequency-Domain Dynamics," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, 2012.

[16] L. Anselin, "Local indicators of spatial association – lisa," *Geographical analysis*, vol. 27, no. 2, 1995.

[17] F. Malandrino, M. Kurant, A. Markopoulou, C. Westphal, and U. C. Kozat, "Proactive seeding for information cascades in cellular networks," in *IEEE International Conference on Computer Communications (INFOCOM)*, 2012.

Acknowledgements

This work is supported by the Science Foundation Ireland under Grant No. 10/IN.1/I3007.

The authors would like to thank Telefónica and Eircom for their cooperation. They are also grateful to Mr. Robert Mourik and Mr. Eddie Gleeson for their insightful discussions.