An evaluation of IEEE 802.11 community networks deployments

German Castignani, Lucien Loiseau, Nicolas Montavont

An Evaluation of IEEE 802.11 Community Networks Deployments

German Castignani
IT/TELECOM Bretagne
Univ. Europ. de Bretagne
Cesson-Sévigné, France
german@castignani.com.ar

Lucien Loiseau
IT/TELECOM Bretagne
Univ. Europ. de Bretagne
Cesson-Sévigné, France
lucien.loiseau@telecom-bretagne.eu

Nicolas Montavont
IT/TELECOM Bretagne
Univ. Europ. de Bretagne
Cesson-Sévigné, France
nicolas@montavont.net

Abstract—**IEEE 802.11 appears as the most popular technology for wireless access because of its low cost and high data rates capacity. In the recent years a new communication paradigm emerged, called the Community Networking, which takes advantage of existing residential 802.11 access points and is based on sharing some part of the bandwidth to members of the community. We evaluate the 802.11 environment, the quality of link, and the potential of these community networks in the city of Rennes, France. We aim at determining if these networks are offering enough coverage area to be considered as an alternative to cellular networks.**

I. INTRODUCTION

Within the fourth generation (4G) wireless mobile environment, 802.11-based technologies are playing a major role because of their low-cost and high data rate capacity, allowing users to deploy their own networks and directly connect them to the Internet. This is confirmed by the extraordinary growth of WiFi products demand, estimated by market forecasts [1] to be augmented by 33% in 2010, with more than 60% of this demand being only caused by 802.11n-compliant products, due to its recent standardisation.

This remarkable growth is not only related to new user devices but to APs as well, producing an increasing number of deployments. We can roughly categorize these deployments into two models: hot spot and community networks (CN). In the first case, APs are deployed in specific areas of cities, offering a geographically limited wireless access to users close to the APs. These APs are either plugged to a broadband Internet connection, or inter-connected via a wireless connection independent from the 802.11 wireless link dedicated to the users (which form a mesh architecture). The first city-wide WiFi deployments were proposed by municipalities (or a single operator) by installing APs in so-called hot spots across a city. In the second case, the CN paradigm is defined as the sharing of residential APs to a given community of people for accessing the Internet. A community is defined as clients from a single operator, or users that are collaborating together. Different residential Internet Service Providers (ISPs) provide users with 802.11 APs which have the ability to broadcast different network identifiers (ESSID). One of these ESSIDs is common to all clients and is broadcasted by all APs from the same ISP that have activated the CN service. Thus a user walking in a city can benefit from Internet access by using the CN. Even if current CN deployments are managed by ISPs, the origin of this kind of networks comes from small user communities having specific knowledge in wireless LANs, neighborhood deployments and associations. Additional information could be found in the Wireless CN forum (http://wcn.cnt.org/).

CNs thus appear as an interesting alternative to provide Internet access everywhere and for low cost. However their performance are still unknown and unpredictable; the deployment is random, as it depends on the number of clients, their willing to share, and their locations. Moreover, APs are deployed indoor, which may reduce the APs' coverage area and result in bad connection from a user on the street. In this paper, we propose an evaluation of these CNs and analyse whether they might be considered as a reliable and usable solution for an ubiquitous Internet access.

A. The 802.11 deployment problematic

802.11 operates in the free ISM band, so deployments are never totally controlled since anyone could deploy any device on that frequency band. This means that 802.11 devices can suffer from intra-system interferences (with other 802.11 devices) and inter-system interferences (with other technologies using the same frequency band). In the 2.4GHz band, 802.11 defines 13 channels (from 2.412GHz to 2.472GHz) in Europe. Each channel is 22MHz large and they are spaced 5MHz apart. This implies that adjacent channels are strongly overlapping, and that two APs operating in the same area should at least be separated by 22 MHz, which corresponds to 5 channels. For example, if an AP is using channel 1, then another AP in the same area should avoid channels 2 through 5. Due

to this large overlapping, it is often recommended to only use channels 1, 6 and 11 [2]. Though not controlled by a central entity, channel assignment may be optimized by good practices, shared knowledge, or default equipment configurations. Nowadays, APs often come with automatic channel selection, which consists in determining the less congested channel to use. In this paper, one of our goal is to identify which channels are the most used and to determine whether the 1-6-11 deployments are a myth or match the reality. The used channels have an important impact on user mobility, since a mobile user should probe channels one after the other to find APs. By identifying which channels are the most used, we could prioritize those channels whenever a user is searching for an AP and thus minimizing the interruption time due to a handover from one AP to another.

B. The experiment

Within this work we aim to study the current 802.11 deployment in the city of Rennes, France. Rennes has a population of almost 208.000 inhabitants with an average density of $4.126 hab/km^2$. It is also a well-known interest area in computer sciences and telecommunications, with a vast number of research centers, universities and technological companies. The overcrowding of 802.11 deployments provides a high level of ubiquity to users, but a number of limitations may arise due to the unmanaged deployment of such networks. Some of them are related to interference, security issues and multipath loss, among others.

The experiment results we are presenting in this paper show the current and real wireless deployment in different parts of the city. We focused on the frequency usage, the level of security and the performance in terms of link quality. We also analyze the coverage of CNs to check whether a continuous connectivity could be used by a user moving in the city.

The outline of this paper is as follows. In section II we present the related work on hot-spot and community wireless networks, considering deployment issues and measurement studies. Then, in section III we present our experimentation platform, metrics and results. We also compare our results with a state of the art measurement study. We finally conclude the paper and delineate future work in section IV.

II. RELATED WORK

A. Hot spot Deployments

The first city-wide WiFi deployments were proposed by municipalities (or a unique operator) by installing APs in so-called hot spots across a city. These APs were either plugged to a broadband Internet connection, or inter-connected via a wireless connection independent from the 802.11 wireless link dedicated to the

users (which form a mesh architecture). Some experimental and operational architectures can be cited: MIT Roofnet [1], Champaign - Urbana Network [2], Oklahoma City Network [3], Madison MadMesh [3] and the Google WiFi Network [4], among others.

Aguayo et al. [4] measure the performance of wireless links in a 38-node Roofnet deployment in Massachusetts. While presenting several metrics and analysis, they focus on the packet loss problem, concluding that intermediate loss rates is probably due to multi-path fading rather than attenuation or interference. Another measurement study performed by Brik et al. [5] during eights months focus on the Madison MadMesh Network. By using both passive and active measurements, the authors analyse the deployment's robustness in terms of link quality and mesh routing, client connectivity and network usage. They found that the bottleneck of the network is related to the client access side rather than the backbone side.

More centered on users activity, Afanasyev et al. [6] and Arjona et al. [7] studied the Google WiFi Network, a 500 APs mesh deployment in Mountain View (California). Afanasyev's study was more oriented on user activity and usage of the network (type of application and device, time of access, etc). Arjona et al. evaluated Voice over IP (VoIP) applications in mobile outdoor conditions to measure the potentials of Voice over (wireless) IP compared to cellular voice services. Testing VoIP did not give good results in term of mobility, because the network had several gaps of coverage. The authors also made a comparison with cellular networks using simple propagation models, concluding that AP density of the Google WiFi Network should increase from $12AP/km^2$ to $81AP/km^2$ to perform as well as cellular networks. Hot spots require a high level of investment and a dedicated infrastructure for network maintenance and operation. Because of this barrier, some Municipal WiFi (mesh) in the US failed [8], as is the case of Philadelphia city, where the city-wide deployment mutated to isolated hot spots around the city. With a totally different approach, the notion of CNs as presented in the next subsection does not suffer from this cost of deployment and maintenance.

B. Community Network Deployments

Recently, CNs emerged as a possible solution to provide Internet connectivity for a vast number of users in urban areas (see section I). CNs stand as self-organized and fully decentralized deployments, since APs are mainly connected to a residential broadband line. Due to this characteristic, APs are mainly deployed indoor without any deployment strategy, while

[1]MIT Roofnet: http://pdos.csail.mit.edu/roofnet/doku.php

[2]Urbana network: http://cuwireless.net/

[3](Madison MadMesh: http://madcitybroadband.com/

[4](Google WiFi network: http://WiFi.google.com/

hot spot APs are deployed outdoor. This is one of the aspect that motivated our experiment - to study if a user from a street can connect to these APs, and how long he/she can stay connected. Since these networks are based on opening the residential broadband access to the community, security and management issues arise as a challenge. These critical issues are studied by Vidales et al. [9], where an enhanced architecture of CN deployments is proposed.

From a commercial point of view, the FON Wireless company (www.fon.com) appears as the leader of the community networks business, since they deployed one of the first shared-wireless infrastructure. They also set up an incentive model to motivate people to open their connections to the general public in a secured manner. In France, numerous community-aware ISP can be found on the market like Free (http://www.free.fr) and Neuf-SFR (http://www.sfr.fr). Moreover, users may find some APs deployed by these ISP as hot-spots in public spaces (not residential APs), that are accessible by using the same community identifiers.

C. Evaluation

One of the most interesting and complete measurement studies of metropolitan WiFi deployments (not limited only to mesh deployments) is proposed by Bychkovsky et al. [10]. In this work, the authors focused on estimating the expected performance of 802.11b community deployments in the context of vehicular networking. This measurement study took 290 driving hours (within 232 days) by using up to nine cars implementing an embedded computer (CarTel) with a 5.5dBi omni-directional antenna. These cars performed active scanning to discover 802.11 APs. Then, when it was possible to associate, they configured an IP address and performed some data communication. Within this experience, they discovered over 32,000 APs, but they could only associate and ping in the 6.6% of the cases. Some interesting metrics have been presented in this work. The median amount of TCP data successfully uploaded by a moving car was $216Kb$, with a median throughput of $30Kbytes/s$. Regarding mobility, the average time of link-layer connectivity of a moving vehicle is $13s$, independently from the speed of the vehicle up to $60km/h$. The measured average time between two connections was quite high, around $75s$. This study also identifies the main concerns the research community should focus on before these deployments become a reality. Commercial and legal issues may be studied. For instance, one hot topic in CN is related to incentives models, in order to motivate the introduction of the community paradigm. A credit based approach is proposed in [11]. Morover, Manshaie et al. [12] proposed a game-theoretic model to study the evolution and potential market share of CNs, considering the competition with traditional licensed-band operators (3G, LTE).

Within this paper we aim to present and analyze a set of metrics defining 802.11 deployments in a medium-size European city. We focus on CNs since we observed an important relevance of these deployments on the whole 802.11 environment.

III. EXPERIENCE

A. Description

During the spring 2010, we performed an 802.11 APs deployment analysis along an $8km$ path in the city centre and some neighborhoods of Rennes, in France. We aim at identifying how many APs are deployed and how they are deployed, what are the coverage areas of the public hot spots and the CNs and what are the potential usages of this deployment for a mobile user. We are trying to answer the question: "Is it possible to access the Internet from everywhere in the city using an 802.11 access?". We also expect to use these results to document and further develop algorithms needed when a mobile user is roaming from one AP to another. As we have shown in [13], these algorithms highly depend on the channels used by the APs. In order to answer those questions, we used a laptop as a sniffer to capture the 802.11 frames along the defined itinerary. The materials and methodology of the experiment are presented in this section.

1) Materials: We captured the 802.11 frames with an Asus N10JB running Ubuntu 9. We extended the internal antenna by using an external DLink $7dBi$ omnidirectional antenna. The internal wireless card runs a modified version of the ath9k open source driver. In order to log the location of the captures, we used an external USB GPS receiver and osm-gps-map (http://nzjrs.github.com/osm-gps-map/) for placing the captures on a map. With regards to the results extraction, we implemented a set of scripts in Perl and C++.

2) Methodology: In order to maximise the success of AP discovery, we implemented some modifications in the MAC layer of the ath9k (mac80211 kernel module). We forced the Asus computer to perform Active Scanning every two seconds while moving around the city. This scanning procedure was performed by listening each channel for a longer duration than the one defined in the default ath9k implementation. First, the station tunes its radio on the first channel ($2412Mhz$) and sends three Probe Requests on that channel. Then the station waits for $50ms$ for the reception of Probe Responses and Beacons frames, containing the information about different APs in the channel. This procedure is performed for each channel up to channel 13 ($2472Mhz$). Therefore the time required to scan all the 13 channels is around 700ms (including the time to switch channels).

(a) APs Deployment along the path

(b) Community Network ESSID Distribution

(c) Security at MAC Level

(d) Channel Distribution

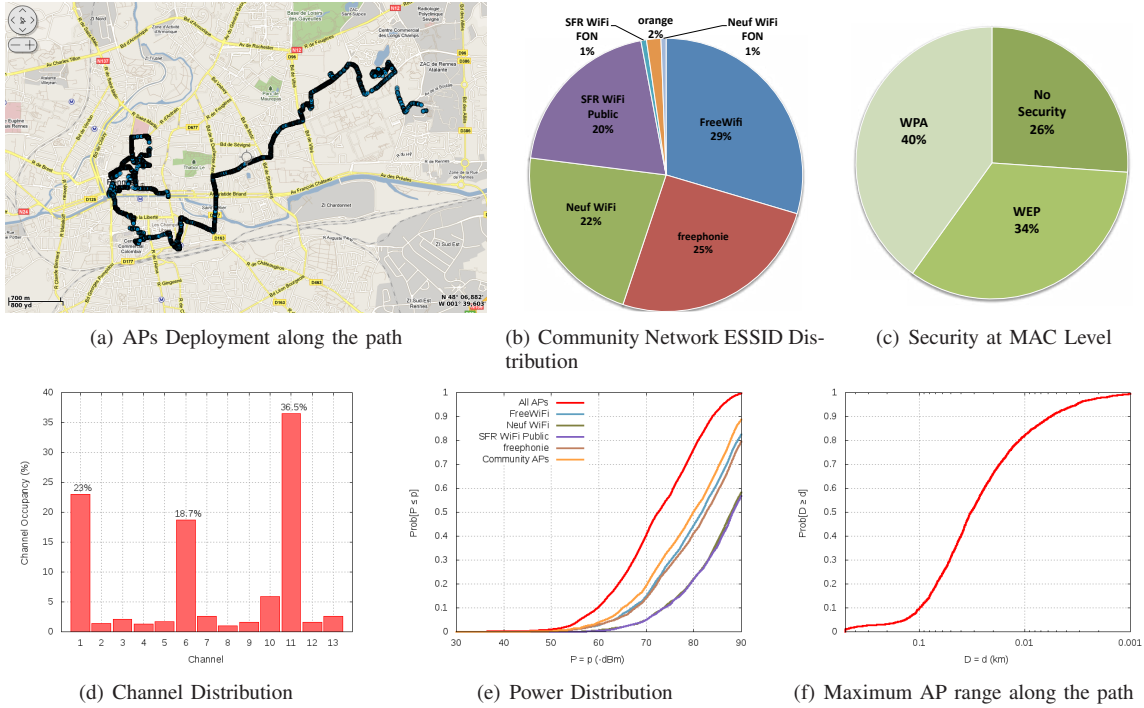(e) Power Distribution

(f) Maximum AP range along the path

Fig. 1. General Results

At the same time, the station captures every second the geolocation information in a separate file. This file and the one with the responses from the APs are merged offline in order to obtain a complete view of the information of each AP, including the geographical data of the measure point. Regarding the mobility pattern, we mainly perform the experimentation by walking. However some captures outside the city center have been done by car, driving at a low speed, always below $45km/h$.

B. Results

1) General Results: Fig. 1 summarizes the general results of the experience. Along the $8km$ path (Fig. 1(a)) we discovered 6220 different APs and 3305 different network identifiers (ESSID). 29% of these belong to CNs, FreeWiFi being the most frequently found ESSID (Fig.1(b)). Considering security issues at MAC level, in Fig. 1(c) we see that WPA-based authentication is the most frequently implemented with 40%, followed by WEP (34%) and finally a significant 26% of APs without authentication. This is mainly due to the presence of CNs, since its authentication is based on a captive portal. However, data encryption is not provided at the wireless link.

Fig. 1(d) shows the channel allocation and more than 78.21% of the APs are deployed in the (recommended) non-overlapping channels 1, 6 and 11. The percentage of usage of the rest of the channels is always between 1% and 5.9%. This result perfectly matches the non-overlapping channel occupancy

calculated by Gerla et al. [14] for urban wireless deployments in California (USA), in which they found 77.98%, even though additional combinations of non-overlapping channels could be used in France since two additional channels are available (see section I-A).

With regard to the link quality, Fig. 1(e) shows the distribution of the received power (in dBm). The value of the power was obtained by extracting the RSSI field from the Beacons and Probe Responses from APs. The all-APs curve represents the distribution for all discovered APs, and we also present the power distribution for the most relevant CNs. Within the limits of the operative power range of 802.11 (-80dBm) we can find more than 75% of the APs.

Another useful metric is related to the maximum distance a user can use an AP. Depending on the speed of the user, this can approximately determine how long an AP is available. This metric is computed by finding the longest geographical distance between two Probe Responses received from the same AP. Fig. 1(f) illustrates the distribution of this metric, showing an average distance of $32.3m$. Regarding this last metric, results presented by Bychkovsky et al. [10] suggest an AP coverage of $96m$ in an 802.11b environment in Boston (USA). The difference between this value and our estimation ($32.3m$) could be related to the transmission power of the devices, which was $200mW$ in [10] against $10mW$ (without the gain from the external antenna) in our experiment.

(a) Number of APs per scan   (b) Inter Channel Interference   (c) Intra Channel Interference
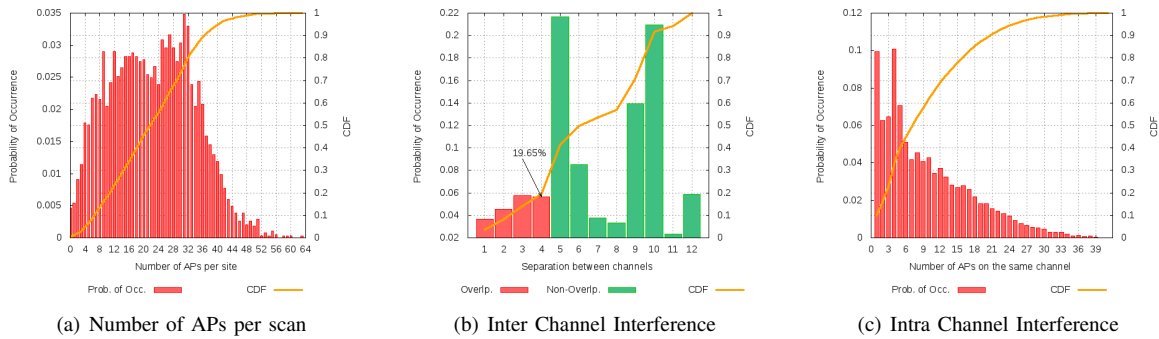
Fig. 2. Channels Occupancy

2) Interferences: In an European 802.11 deployment, up to 13 channels are available in the 2.4GHz band. Following the discussion given in I-A, only three channels can be simultaneously used without interferences. The most recommended channels are 1, 6 and 11 [2], but other combinations could be used such as 2, 7 and 12. Sources of interference can be generated by several APs sharing the same physical channel or being separated by less than 22 MHz. In order to study this issue, we propose to evaluate the distribution of the number of APs per scan (Fig. 2(a)), the distribution of the number of APs on the same channel per scan (Fig. 2(c)) and the inter-channel separation between APs per scan (Fig. 2(b)).

Fig. 2(a) shows that 22 APs are discovered in average during a scan. Moreover, the most frequent value is 32 APs per scan, showing an overcrowded wireless environment. Complementing Fig. 2(a), we propose two interference metrics. A first metric named intra-channel interference represents the distribution of the number of APs that are being deployed on the same channel per scan. Fig.2(c) shows that 7 APs are deployed in average on the same channel, and we observed up to 41 APs operating on the same channel. Secondly, we measure the minimal distance in number of channels from an AP to another operating in a different channel to observe how the APs are spread in the channels sequence. We name inter-channel interference (see Fig. 2(b)) the separation between active neighboring channels (those having at least one deployed AP). For almost 20% of the cases, APs are deployed with an inter-channel distance lower or equal to 4 (which causes interferences). For the rest of the cases, APs' deployment follows a non-overlapping pattern, since a distance greater or equal to 5 channels (i.e. 25 MHz) is observed. The peak observed for a separation of 10 channels could be generated by deployments in which channel 6 is not used, and only channels 1 and 11 are occupied by at least one AP. Similarly, the peak observed at a separation of 5 channels stands for the cases where channels 1 and 6 are used, or 6 and 11. These results

are coherent with those in Fig. 1(d) where we observed that channels 1, 6 and 11 were the most used ones. The analysis of this set of results clearly shows that the 802.11 environment is overcrowded, and while APs are generally deployed on the non-overlapping channels 1,6 and 11, there is a high number of APs deployed on the same channel.

3) Network Potential: Large 802.11 deployments come up as a key point in 4G heterogeneous environments, because 802.11 may serve as an alternative technology to 3G in some areas. A CN can easily provide Internet access to a wide number of users in urban areas at a low price. In our study, we identified that FreeWiFi, FreePhonie, Neuf WiFi and SFR WiFi Public are the most relevant CNs around the city. The similarity of the results from FreeWiFi and FreePhonie (Free) and Neuf WiFi and SFR WiFi Public (Neuf/SFR) respectively may be explained by the fact that both pairs of ESSIDs are advertised by the same physical AP. They only represent two different services, namely Internet access and voice calls. For all these identifiers, we analyzed two metrics, as shown in Table I. The Global Coverage is the percentage of locations where the network was found to be active. The Longest Path metric aims to estimate the distance a mobile user remains connected to at least one AP from the given CN while moving. Regarding the global coverage, we can appreciate that Free leads the community deployments with 67% of the path covered by at least one AP. A user having Free identifiers may find an available AP with higher probability than the second operator, Neuf/SFR. An important consideration is that the growth of such CNs is directly related to the growth of the broadband access market. This means that the efforts to deploy this kind of networks are not related to a single entity (government or operator) but to the whole population, making deployments more feasible.

Concerning user mobility, we found that the FreeWiFi network has the greatest potential. A mobile user could have uninterrupted access to at least one FreeWiFi AP along $292.12m$. If we consider a user

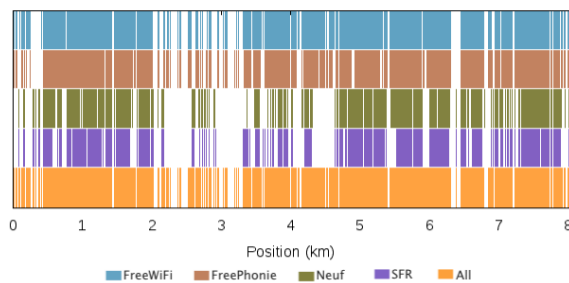| Network	Global Coverage (%)	Longest Path (m)
FreeWiFi	67.43	292.12
FreePhonie	63.52	283.18
Neuf WiFi	41.06	125.89
SFR WiFi Public	38.59	125.89
All	74.77	451.82

Fig. 3.   Presence of APs along the path

having access to both Free and Neuf/SFR platform, he could move along $451.82m$ (as a maximum) without interrupting the connection. However, for now mobility is not supported between APs, which means that a mobile user still has to re-connect to each AP and thus interrupting his/her session. We computed the presence of CNs along the path in Fig. 3. We can observe than Free and Neuf/SFR deployments are quite complementary. The lowest stripe from Fig. 3 illustrates the case of all CNs being considered together, showing a lower number of uncovered zones.

IV. Conclusions and Perspectives

We presented a first set of results to evaluate the performance and the usability of 802.11 networks deployed in Rennes, France. We proposed a number of metrics to measure how many APs are reachable, what is the level of interference, how the channels are used, and how long a user remains connected to an AP. These metrics describe what are the conditions that a user faces while attempting to connect to an AP.

We found that the 802.11 environment is over-crowded (6220 APs over an 8km path), where a large number of APs (78%) are operating on the well-known non-overlapping channels (1, 6 and 11). However, the number of APs per scan is so significant (32 APs in average) that many APs are sharing the same channel leading to poor quality. These findings are important for the development of optimized channels scanning algorithms (for mobile users), by giving priority to those channels. When concerning the link quality, not all the discovered APs provide a power level that guarantees a stable communication. This could be due to regulations on transmission power limits and the fact that APs are mainly deployed indoor.

Focusing on CNs, we can state that this new paradigm is greatly present around the city, giving users the possibility to rapidly access the Internet at a low cost and to obtain reasonable performance. We observed that the two most important operators are quite complementary, giving the possibility to a user to remain connected up to $450m$ (assuming the user roams between APs). In terms of potential of these networks, we can consider that its growth may follow the residential broadband access market progress, and so the coverage of these networks should increase as well. As future work, we aim to develop a framework not only to perform discovery of the wireless environment, but also to evaluate the performance at the application level and study mobility using different CNs.

References

[1] L. Snol, "Wi-fi forecast to fly high this year." PCWorld, 2010.
[2] P. Fuxjager, D. Valerio, and F. Ricciato, "The myth of non-overlapping channels: interference measurements in ieee 802.11," in *Wireless on Demand Network Systems and Services, 2007. WONS '07. Fourth Annual Conference on*, pp. 1 –8, 24-26 2007.
[3] "Ock: It's all in the air." A Tropos network Case Study, 2008.
[4] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, (New York, NY, USA), pp. 121– 132, ACM, 2004.
[5] V. Brik, S. Rayanchu, S. Saha, S. Sen, V. Shrivastava, and S. Banerjee, "A measurement study of a commercial-grade urban wifi mesh," in *IMC '08: Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*, (New York, NY, USA), pp. 111–124, ACM, 2008.
[6] M. Afanasyev, T. Chen, G. M. Voelker, and A. C. Snoeren, "Usage patterns in an urban wifi network," *Networking, IEEE/ACM Transactions on*, vol. PP, no. 99, pp. 1 –1, 2010.
[7] A. Arjona and S. Takala, "The google muni wifi network–can it compete with cellular voice?," in *Telecommunications, 2007. AICT 2007. The Third Advanced International Conference on*, pp. 11 –11, 13-19 2007.
[8] N. Patel, "Municipal wifi." http://www.engadget.com/, 2008.
[9] P. Vidales, A. Manecke, and M. Solarski, "Metropolitan public wifi access based on broadband sharing," in *Computer Science (ENC), 2009 Mexican International Conference on*, pp. 146 – 151, 21-25 2009.
[10] V. Bychkovsky, B. Hull, A. Miu, H. Balakrishnan, and S. Madden, "A measurement study of vehicular internet access using in situ wi-fi networks," in *MobiCom '06*, (New York, NY, USA), ACM, 2006.
[11] X. Ai, V. Srinivasan, and C.-K. Tham, "Wi-sh: A simple, robust credit based wi-fi community network," in *INFOCOM 2009, IEEE*, pp. 1638 –1646, 19-25 2009.
[12] M. Manshaei, P. Marbach, and J.-P. Hubaux, "Evolution and market share of wireless community networks," in *Game Theory for Networks, 2009. GameNets '09. International Conference on*, pp. 508 –514, 13-15 2009.
[13] G. Castignani, A. Arcia-Moret, and N. Montavont, "An evaluation of the resource discovery process in ieee 802.11 networks," in *MobiOpp '10: Proceedings of the Second International Workshop on Mobile Opportunistic Networking*, (New York, NY, USA), pp. 147–150, ACM, 2010.
[14] E. Giordano, R. Frank, G. Pau, and M. Gerla, "Corner: a realistic urban propagation model for vanet," in *Wireless On-demand Network Systems and Services (WONS), 2010 Seventh International Conference on*, pp. 57 –60, feb. 2010.