# An Experimental Evaluation of Packet-Level Measurements of Hidden Traffic Load

José Núñez-Martínez†, Marc Portoles-Comeras†, Albert Cabellos-Aparicio§, Daniel López-Rovira§, Josep Mangues-Bafalluy†, Jordi Domingo-Pascual§

†Centre Tecnològic de Telecomunicacions de Catalunya
Castelldefels (Barcelona), Spain
§Universitat Politècnica de Catalunya
Departament d'Arquitectura de Computadors
Barcelona, Spain

**Abstract.** Hidden traffic interference has been identified as an important source of network instabilities in dense wireless network deployments. This paper follows recent research results that relate packet loss and hidden traffic interference and explores, through extensive experimentation, the possibility to infer hidden traffic load using loss measurements in practical WLAN deployments. Furthermore, we present a preliminary tool able to provide online estimates of the hidden traffic load combining active and passive measurements.

**Key words:** Wireless, Measurements, Hidden traffic, Experimental

## 1 Introduction

WLAN deployments have spread at an unprecedented pace during the last years. The high availability of off-the-shelf hardware and software solutions and the configuration flexibility that they offer have driven the use of IEEE 802.11 devices in multiple and varied scenarios.

However, as a side consequence of this success, it is common to find multiple wireless devices or WLAN deployments coexisting in shared spaces. As a result, the interference between WLAN transmissions has been identified as an important source of unexpected problems in wireless networks. This has been specially remarked and studied in the context of wireless mesh networks (e.g. [1],[2]) where it highly increases the complexity of tasks such as routing (e.g. [4], [5]). However, the undesirable effects of interference also extend to traditional WLAN access networks, specially in dense urban areas or offices, where it has fostered the development of strategies for channel and power allocations (e.g. [3]).

This paper focuses on the measurement of hidden traffic interference from a packet level perspective. We consider as hidden interference those transmissions that are not considered (i.e. sensed) harmful from a sender-side perspective when transmitting a packet, but that lead to losing this packet at the receiving side of the communication. Note here that the study does not restrict to a particular interference model [8], but considers the measurement of hidden interfering transmissions from a generic perspective.

More specifically, the paper studies the inter-relation that exists between packet losses and the presence of hidden transmission and takes base in our previous work [9]. In [9] we used renewal theory arguments to show how the losses of a probing packet sequence can be used to measure (infer) the actual hidden traffic load affecting the communication between a pair of nodes. We provided expressions of the bias that such a measure presents and suggested methods to avoid it.

In this paper we extend our previous work and we present extensive experiments validating the assumptions taken in a wide range of scenarios. Our experiments confirm that loss samples are biased measurements of hidden traffic, and that this bias can be quite large, eventually. The experiments serve also to analyze the use of simple estimators to measure hidden traffic in practical scenarios. Finally the paper presents a preliminary tool to gather online measurements of hidden traffic load.

The results of this study provide fundamental insights to understand some of the results obtained in related literature (e.g. [1], [6]). From a use case perspective, the approach presented here can be used to model the effects of hidden traffic interference on those routing metrics that rely on active probing of wireless channels [2] (e.g. ETX and ETT), and to tune WLAN optimization strategies that rely on loss characterization [7].

## 2 The interrelation between probing losses and hidden traffic load

The study presented in [9] reveals how packet losses constitute, in fact, biased samples of hidden traffic load. Let us denote as $U(t)$ the process describing the *busy/idle* state of a wireless channel. If this channel utilization is hidden to a probing station, that sends probing packets at instants $\{T_n\}$, the *successful/lost* transmissions of these probing packets constitutes a sample of the hidden traffic load such that,

$$U_n = U(T_n) = \begin{cases} 1 & \text{probe packet n is lost} \\ 0 & \text{probe packet n is received} \end{cases} \tag{1}$$

Taking the long term hidden traffic utilization as $u_r = \lim_{t\to\infty} P[U(t) = 1]$, and assuming that the sequence of probing arrivals $\{T_n\}$ follows a Poisson distribution, the study in [9] reveals how the limiting average of the probability of losing a packet is biased with respect to the channel utilization. This can be expressed as

$$\lim_{n\to\infty} P[U_n = 1] = u_r + \varepsilon(T_p), \tag{2}$$

where $\varepsilon(T_p)$ is the bias term and depends on $T_p$, the time that it takes to transmit the probing packets over the air.

### 2.1 Building loss based estimators of hidden traffic load

The previous results provide fundamental insights that can be used to develop packet probing based estimators of the hidden channel utilization. Here we review two basic approaches that help illustrating the impact of the bias term.

Following the definitions in the paper, a sample-mean estimator of the hidden traffic utilization based on the losses of a Poisson sequence, can be defined as

$$\widehat{u}_s(T_p) = \frac{1}{N}\sum_{i=0}^{N-1} U_i = u_r + \varepsilon(T_p) + w, \tag{3}$$

where $N$ samples are used for the estimation and where $w$ denotes the noise associated to the measure (assumed with a zero mean and independent of the hidden traffic). As it can be seen the sample mean constitutes, in fact, a biased estimator of the hidden traffic utilization.

Another important result of [9] is that it shows that when $T_p$ is sufficiently small, $\varepsilon(T_p)$ can be modeled as being linear with respect to $T_p$. As a consequence, the following proposed estimator uses two different probing sizes $T_p^1 > T_p^2$ to obtain a (linear approximation) based unbiased estimation of the utilization,

$$\widehat{u_r}(T_p^1, T_p^2) = \frac{\widehat{u}_s(T_p^2)T_p^1 - \widehat{u}_s(T_p^1)T_p^2}{T_p^1 - T_p^2}. \tag{4}$$

Next sections provide extensive experimental results validating, in practical cases, the statements of this section.

## 3 Experimentation setup

This section describes the methodology employed to carry out experiments. Precisely, we have used a controlled experimentation platform called EXTREME Testbed® [13] and a sniffer attached to a real operational WLAN network.

### 3.1 Experimentation platform

All the experiments were carried out using the EXTREME Testbed ® [13] deployed at CTTC. This is a multi-purpose networking experimental platform featuring high automation capabilities that support automatic execution of the experiments, data collection and data processing.

The EXTREME testbed is composed of a cluster of computer nodes. All these nodes are Pentium 4 PCs with a 3GHz processor, 512MB of RAM memory, and running Linux with kernel 2.6.27. Every node can be equipped with up to two wireless Network Interface Cards (NICs). The type of wireless NIC employed for the experiments is the LevelOne WNC-0300, which is based on the Atheros 11b/g

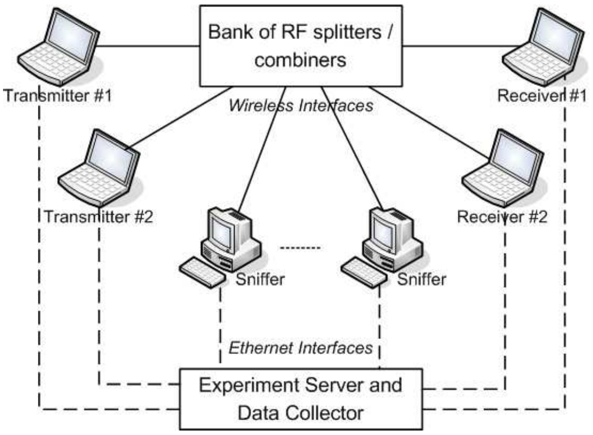

**Fig. 1.** Setup of EXTREME Testbed ®.

chipset. The automated experiment setup in EXTREME makes an extensive use of the wireless extensions API [14] to configure and control wireless devices. Specifically, the Madwifi driver [10] supports this API and controls LevelOne cards.

Since the main objective of this study is devoted to study interference due to hidden nodes, the testbed setup was designed as to minimize the effects of channel propagation on the measurements. To this aim, all the communications between wireless devices are done through coaxial wires, and all wireless devices are connected to a central bank of splitters and combiners (see Figure 1). This bank of splitters replicates with very low attenuation (in comparison to open-air propagation) all signal inputs in each of its ports to the rest of ports. The bank of splitters and combiners is composed of minicircuit ZX10-4-27 splitters (with 4 ports) and minicircuit ZFSC-2-10G splitters (with 2 ports).

### 3.2 Traces from a campus network

We have also gathered data from a production WLAN AP belonging to a real university campus network. The traffic has been collected using a Linux box (PIV-2GB RAM) equipped with an Atheros card running MadWifi drivers [10] in monitor mode. In particular, we have collected traffic in the following scenarios[1]:

- *Filetransfer_far:* In this scenario, a contending machine was setup to download a large file from a remote server located at many hops from the campus network.
- *Filetransfer_close:* This scenario mimics the previous one, but the remote server was located in the same autonomous system than the campus network.

[1] The traces used can be downloaded from [11].

- *P2P:* In this scenario, the contending node was running a popular P2P application (BitTorrent).
- *Contending_n:* In this final scenario, we setup $n$ contending nodes (ranging from 1 to 4) running a wide set of applications (filetransfer, e-mail clients, P2P...). This setup reproduces a WLAN link under heavy load.

## 4 Experiment observations to validate the model

This section validates some of the statements presented in the previous section and in [9]. First we validate, in a controlled environment, the inter-relation between loss and hidden traffic. Second, we use measurements gathered in an operational campus network to show that the renewal theory assumptions taken in [9] correctly model practical wireless channel utilization. Third, we provide some results that relax the requirement to use Poisson probing in order to gather accurate measurements.

### 4.1 Validation in a controlled environment

We have used the EXTREME platform (see subsection 3.1) to study the inter-relation between packet loss and hidden traffic load in the absence of propagation losses.

Figure 2 reports the measurements of hidden traffic load, for different rates of hidden traffic, when using the sample mean and the estimation based on linear interpolation. The figure has been generated using more than 15k loss samples for each measurement point. To gather these measurements we have used a probing sequence with packets of 128 bytes in the sample mean case, and an additional sequence with packets of 64bytes in the linear approximation case. Hidden traffic is sent at a phy rate of 12Mbps, with Poisson distribution. Probing packets are sent at 6Mbps with Poisson distribution and broadcast.

As it can be seen, first, the sample mean provides biased estimates of the actual hidden traffic load. Second, when using the packet sizes we have chosen, the linear approximation of the bias is consistent and provides accurate estimates of the utilization using the linear approximation.

### 4.2 Validation of the model assumptions on WLAN traffic

The purpose of this subsection is to validate some of the assumptions taken in section 2 and to analyze the accuracy of the proposed estimators using the datasets described in subsection 3.2, that have been collected in an operational campus network. Specifically, we are interested in obtaining realistic figures of the distributions of *busy/idle* states of channel occupancy in operational networks. We take the downlink traffic transmitted through the monitored AP as hidden to an eventual monitoring node and study whether the assumptions hold in such a case. Using the datasets we measure the actual traffic load and the measure

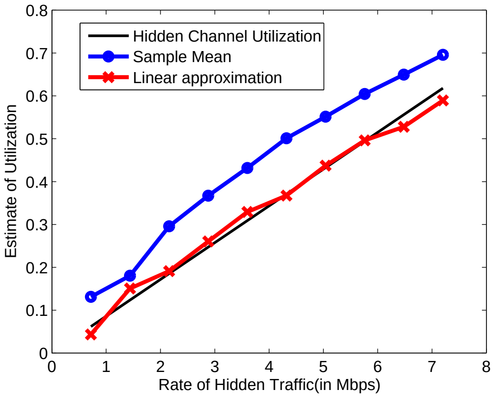

**Fig. 2.** Measurements of utilization using probing losses in a controlled environment

that would be obtained using loss samples from an eventual probing sequence with mean inter-arrival time of 100ms.

Figures 3 and 4 show the results of these experiments. Please note that for readability we do not show all the results, the interested reader can find the remaining ones in [11]. Left plots show the accuracy of the sample-mean estimators (eq. 3) while right plots show the accuracy of the linear approximation proposed to obtain unbiased measurements (eq. 4).

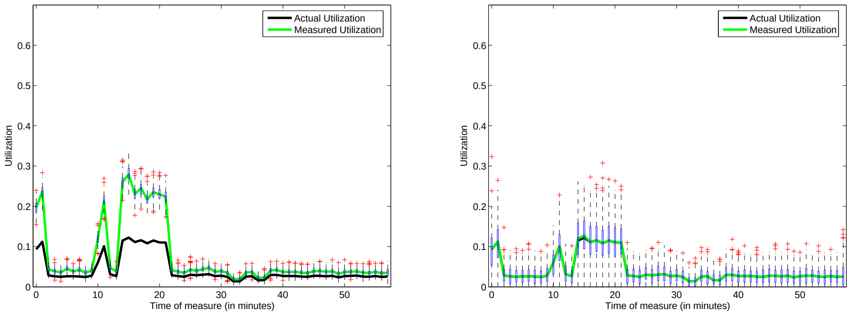

(a) Measurement of utilization using sample-mean.

(b) Measurement of utilization using unbiased estimator.

**Fig. 3.** Comparison of the accuracy of sample-mean vs unbiased estimator (filetransfer_far scenario).

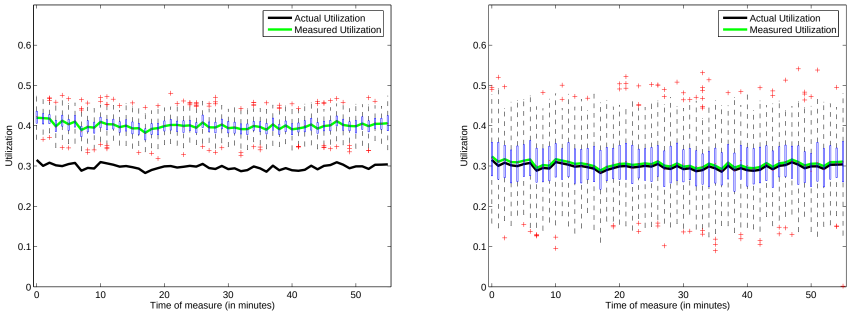

(a) Measurement of utilization using sample-mean.

(b) Measurement of utilization using unbiased estimator.

**Fig. 4.** Comparison of the accuracy of sample-mean vs unbiased estimator (contending_3 scenario).

From the figures we can see that all the experiments validate our analytical model, and show the large bias incurred by sample-mean based estimators. This bias is present in all the cases and may grow up to 50% in some scenarios. Concerning the linear approximation (right side of the plots) the experiments shows that it performs remarkably well. This also confirms that the linear assumption taken holds with a wide range of traffic profiles and with different amounts of contending (hidden) stations. It is worth noting here that the plots also reveal the higher variance of the (approximately) unbiased estimator in comparison to the pure sample-mean measure.

### 4.3 Do we really need to use Poisson sequences?

One of the main assumptions taken in section 2 and used throughout the validation study is the use of Poisson probing sequence to sample channel use. The PASTA property associated to Poisson sequences guarantees the convergence of some of the findings and the accuracy of the estimators proposed. The objective of this section is to relax this assumption and show that channel utilization presents mixing properties [15] that allow using alternative probing patterns. Relaxing the Poisson assumption is useful in practice, as we will see in the next section, as it allows using, for example, *periodic* beacon broadcasts to get samples of hidden traffic utilization.

The hypothesis here is that although the sampling process is not poisson, the underlying traffic is mixing enough, and will produce similar results. For this we have repeated the experiments of the previous section but instead of using poisson probing, we use periodic probing. In particular each plots shows, in each case the accuracy of the proposed estimator using poisson probing (left) and periodic probing (right).

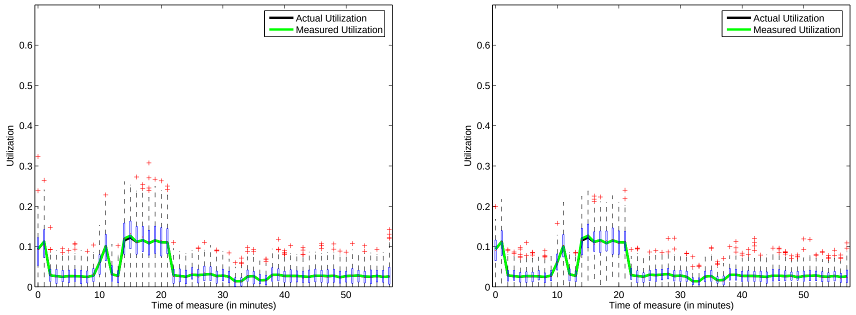

(a) Measurement of utilization using poisson probing (unbiased estimator).

(b) Measurement of utilization using periodic probing (unbiased estimator).

**Fig. 5.** Comparison of the accuracy of periodic vs. poisson probing (filetransfer_far scenario).

Figures 5 and 6 show the results for this experiment. Again, we only present the plots of the same scenarios than in the previous subsection, the interested reader can find the remaining ones in [11]. The results show that, for all the scenarios, the accuracy of the unbiased estimator is not affected by periodic probing and in fact, shows similar accuracy than when using poisson probing.

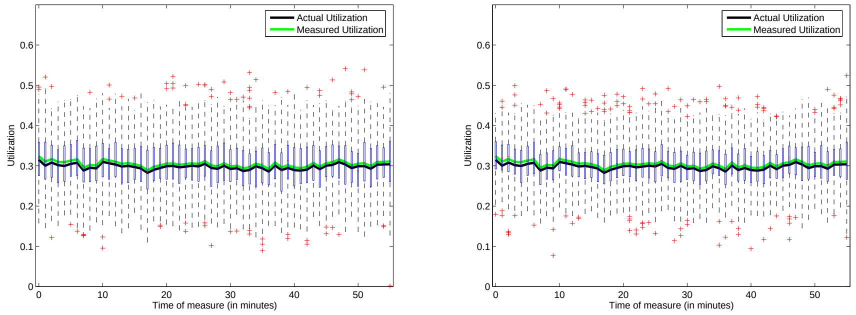

(a) Measurement of utilization using poisson probing (unbiased estimator).

(b) Measurement of utilization using periodic probing (unbiased estimator).

**Fig. 6.** Comparison of the accuracy of periodic vs. poisson probing (contending_4 scenario).

# 5 Design options to build a tool to measure hidden traffic load

This section presents some preliminary observations on the design options to the design of a tool to obtain online measurements of hidden traffic load in practical scenarios.

## 5.1 Periodic samples using beacon frames

Beacon frames can be considered as periodic samples of the wireless channel. Indeed, in a vast majority of cases, Access Points send beacons with a periodic interval that comes in the TBTT field. As argued above in section the periodic nature of beacon broadcasts does not prevent their use, in a practical scenario, to gather samples of hidden traffic load. Note that the authors of [12] suggested previously the use of the TBTT field to gather information from beacon measurements.

We have implemented a tool that keeps track of the TBTT field in order to determine how many beacons might have been lost between two successive reception of beacon frames. However, as figure 7 reveals, beacon losses have to be handled with care, as APs are prone to reconfiguration procedures that may be easily mistook with periods of high traffic occupancy.

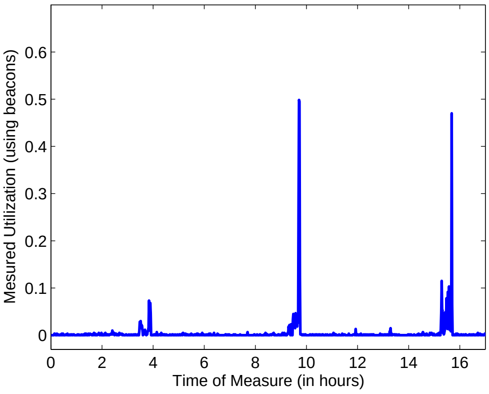

**Fig. 7.** Measurements of the beacon frame losses during night hours

## 5.2 Getting samples using retry counts

In order to obtain unbiased measurements of hidden traffic utilization we need packets of different sizes (and sufficiently small) that sample the occupancy of the wireless medium. However, generating broadcast frames without proper control of the WLAN network is generally a difficult task.

The solution that we investigate here takes advantage of the retry field present in the IEEE 802.11 packet header. In this case we propose the use of unicast packets as indicators of the presence of hidden traffic. The basic idea consists in that (1) a remote station sends packets of two different sizes to the monitoring stations. (2) the monitoring station determines how many packets are not successfully transmitted at their first attempt (i.e., they have their retry field to 1). Then, the packets with the retry field equal to 1 are used as indicators of frame losses, and thus, of hidden channel use.

Figure 8 shows the efficacy of using this method in a controlled environment. The figure plots the number of 128-byte packets that have not been successfully transmitted at their first attempt (sample mean) and a linear approximation of the hidden utilization using a second sequence of 64-byte long packets. As it can be seen the approach is accurate.

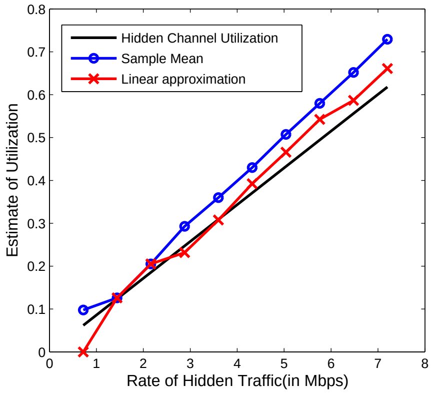

**Fig. 8.** Retry based measurements of hidden traffic in a controlled environment

Figure 9 shows the results of applying retry based measurements to gather unbiased measurements of hidden traffic load in an operational environment.

### 5.3 Issues to implement the tool in practice

There are a number of issues that difficult the implementation of a practical tool to measure the present of hidden traffic transmissions.

First, the solution to use unicast frames and retry counts suffers largely from adaptive mechanisms implemented in commercial WLANs. As an example, rate adaptation algorithms and variable retransmission strategies prevent the possibility to use a single value for $T_p$ to be used in the estimation.

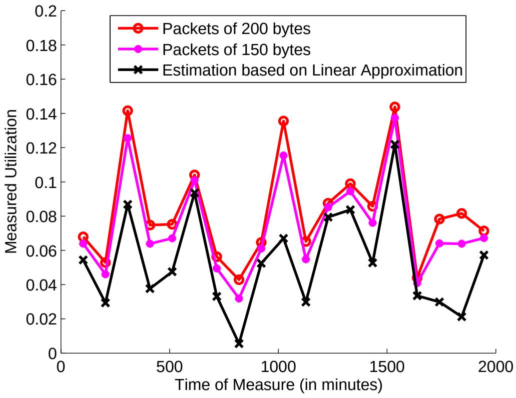

**Fig. 9.** Retry based measurements of hidden traffic in an operational network

Second, channel propagation errors constitute an important source of distortion to the mechanisms introduced here. In many cases, channel propagation errors are more important than hidden transmissions which results in the impossibility to gather reliable measurements of hidden traffic load. A typical solution is to use low rate modulations to transmit probing packets in order to increase their reliability.

Finally, beacons present the desirable properties of being short and transmitted at a sufficiently low rate so that the impact of channel propagation is low. However, without having a certain level of control of a WLAN network, it is difficult to introduce additional broadcast transmissions to build unbiased estimators.

## 6 Conclusions

Hidden traffic interference has been identified as an important source of network instabilities in dense wireless network deployments. This paper explores the possibility to infer hidden traffic load using loss measurements in practical WLAN deployments. It takes base on [9], where we proposed an unbiased estimator for the hidden load interference. At present, existing estimators account for the hidden load computing the ratio of lost packets from a sequence of periodic probing and using a pure sample-mean estimator. In [9], we proposed an approximately unbiased estimator based on poisson sampling.

In this paper we have validated through extensive experimentation the validity of the assumptions taken in [9] along with the accuracy. The validation has been carried out both in a controled testbed (EXTREME® [13]) and using a dataset collected in an operational campus network. The results confirm the assumptions and show that the proposed estimator performs remarkably well.

Further, the results also show that the proposed estimator can also operate when we sample the path using a periodic probing process, instead of a poisson one. This hypothesis has helped us to design a tool able to provide online estimates of the hidden traffic load combining active and passive measurements.

**Acknowledgments.** This work has been partially funded by the Spanish Ministry of Science and Innovation under grant number TEC2008-06826/TEC (project ARTICO), and by the Catalan Regional Government under grants 2009SGR-1140 and 2009SGR-940.

## References

1. J. Padhye, S. Agarwal, V. Padmanabhan, L. Qiu, A. Rao, and B. Zill. *Estimation of Link Interference in Static Multi-hop Wireless Networks.* In ACM/USENIX Internet Measurement Conference (IMC) 2005.
2. Y. Li, L. Qiu, Y. Zhang, R. Mahajan, Z. Zhong, G. Deshpande, and E. Rozner, *Effects of Interference on Wireless Mesh Networks: Pathologies and a Preliminary Solution.* In HotNets, 2007.
3. B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, C. Diot, *Measurement-Based Self Organization of Interfering 802.11 Wireless Access Networks.* in IEEE INFOCOM, 2007.
4. A. P. Subramanian, M. M. Buddkihot, and S. Miller, *Interference Aware Routing in Multi-Radio Wireless Mesh Networks.* in Proc. of IEEE WiMesh, 2006.
5. J. Nuñez-Martinez, J. Mangues-Bafalluy *A Survey on Routing Protocols that really Exploit Wireless Mesh Networks* in Proc of Journal of Communications, 2010
6. N. Ahmed, U. Ismail, S. Keshav, and K. Papagiannaki, *Online estimation of RF interference.* In CoNEXT, 2008.
7. P. Acharya, A. Sharma, E. Belding, K. Almeroth, and K. Papagiannaki *Rate Adaptation in Congested Wireless Networks through Real-Time Measurements* To appear in IEEE Transactions on Mobile Computing
8. M. Takai, J. Martin, and R. Bagrodia. *Effects of wireless physical layer modeling in mobile ad hoc networks.* In MobiHoc 2001.
9. Marc Portoles-Comeras, Marc Sole, J. Nuez-Martinez, Albert Cabellos-Aparicio, Josep Mangues-Bafalluy, Jordi Domingo-Pascual "Packet level measurement of hidden traffic interference in WLAN networks" *submitted for publication*, 2010
10. MadWifi drivers `http://madwifi.org`
11. WLAN Measurements, dataset and Technial report `http://personals.ac.upc.edu/acabello/WLAN\_measurements/`.
12. S. Vasudevan, K. Papagiannaki, C. Diot, J. Kurose, and D. Towsley, Facilitating access point selection in IEEE 802.11 wireless networks, in Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC '05), p. 26, 2005.
13. EXTREME Testbed ® - System level testbed featuring IP mobility, CTTC `http://www.cttc.es/en/projects/testbeds/project/EXTREME.jsp`
14. Wireless Extensions for linux. `http://www.hpl.hp.com/personal/JeanTourrilhes/Linux/`
15. Baccelli, F.; Machiraju, S.; Veitch, D.; Bolot, J.; , "The Role of PASTA in Network Measurement," Networking, IEEE/ACM Transactions on , vol.17, no.4, pp.1340-1353, Aug. 2009