Coupling QoS provision with interference reporting in WLAN sharing communities

Pantelis A. Frangoudis and George C. Polyzos

Mobile Multimedia Laboratory
Department of Computer Science
Athens University of Economics and Business
{pfrag,polyzos}@aueb.gr

***Abstract*—Driven by their low cost and ease of deployment, as well as their operation in unlicensed spectrum bands, IEEE 802.11-based Wireless Local Area Networks (WLANs), also termed Wi-Fi, have been established as the de facto access technology for local area wireless connectivity. Especially in densely-populated urban areas, WLAN presence is ubiquitous. Residential WLAN owners, municipalities and venue owners, among others, set up wireless hotspots for private or public use. However, one can identify two important problems that have to be efficiently tackled. On the one hand, while Wi-Fi is present practically everywhere in modern metropolitan areas, access to roaming users is typically restricted. At the same time, the broadband Internet connections WLAN Access Points (APs) are attached to may have excess capacity, thus leaving resources underutilized. On the other hand, unplanned Wi-Fi deployment leads to significant interference problems among neighbor WLANs. In this work, we exploit our prior work on WLAN sharing communities to jointly tackle the above problems. The solution we propose is based on offering users QoS benefits as an incentive to perform spectrum sensing and supply interference reports to the WLAN APs they are connected to. We present the design of our mechanisms, discuss some of their properties and investigate their expected performance overhead, especially on delay-sensitive applications like VoIP.**

I. Introduction

Wi-Fi equipment has become standard for laptops and handheld devices and appears as the predominant technology for local wireless connectivity. The success of this technology is largely attributable to its low cost and ease of deployment. More importantly, though, it is the fact that this family of standards operates in unlicensed spectrum that has facilitated its deployment, since it was straightforward for commercial operators, academic institutions, or even plain radio communications enthusiasts and tech-savvy users to build wireless service architectures on top of it, without the need for acquiring a license to operate on these particular spectrum bands.

In recent years, a trend towards open wireless access is observable. The properties of this wireless technology have made it particularly attractive for the development of such schemes and have resulted in municipality or state initiatives which aim at offering citizens low-cost wireless connectivity. A more obvious expression of this trend, though, are Wireless Community Networks (WCNs) [1]–[5]; individual WLAN enthusiasts use inexpensive and sometimes hand-crafted equipment (e.g. directional antennas) to build autonomous wireless networks, providing free connectivity to their participants and sometimes offering broadband wireless access to roaming users.

On the other hand, home WLAN owners typically refrain from configuring their APs for open access, raising security concerns and failing to identify the proper incentives. Residential Wi-Fi networks are typically attached to flat-rate fixed broadband lines, which may be unused for many hours a day. This excess capacity could be offered to mobile users roaming around these APs, if the appropriate incentives for sharing existed. Considering the increased coverage of Wi-Fi networks in metropolitan areas, if such a scheme existed, a WLAN-based access solution might offer a low-cost alternative to traditional mobile services offered by 2.5G and 3G systems.

However, the high density of WLANs in urban areas also has a dark side; in many cases, the number of WLANs that are within the communication range of a user is so high that we should be more concerned about interference than coverage. Therefore, in a chaotic environment where WLAN APs are deployed in an unplanned and unmanaged manner, self-organizing schemes for optimally configuring their operation parameters are necessary. Information on spectrum usage is vital input for such schemes, and we believe that wireless clients themselves should actively participate in collecting them.

It is our position that the issues of open wireless access and interference mitigation could be jointly considered. In prior work, we have proposed a decentralized scheme for reciprocal WLAN sharing, called the Peer-to-Peer Wireless Network Confederation (P2PWNC) [6]. In this paper, we extend the basic P2PWNC mechanisms making it QoS aware. Mobile users that join P2PWNC are given the opportunity to enjoy better QoS, provided that they monitor spectrum usage at their location and report this information to the AP they are attached to when requested. The AP can then utilize this information for sophisticated interference mitigation, for instance by means of frequency selection or power control. To the best of our knowledge, this is the first work that couples interference reporting with WLAN sharing.

The remainder of this paper is structured as follows. In Section II we study the fundamental issues of community-based wireless access provision and interference mitigation, with

pointers to related work. In Section III we give an overview of P2PWNC, before we propose QoS extensions which provide users the incentives to contribute spectrum usage information, in Section IV. In the discussion that follows in Section V, we refer to implementation issues, study the effects of the spectrum sensing procedure on VoIP performance and deal with the potential for attacks on the proposed mechanisms. We conclude the paper in Section VI.

II. Related work

A. Open wireless access provision

In recent years, the trend towards open wireless access is more obvious than ever. Below we characterize such schemes based on the initiative behind their emergence.

Community-initiated wireless networks are the result of collective efforts of individual volunteers and function on an altruistic, not-for-profit basis. Such *Wireless Community Networks* [1]–[5] aim at providing free interconnection among their members and a variety of services exclusively offered to the community. They typically have a wireless mesh architecture, with directional wireless links connecting nodes. Sometimes, WCNs operate public hotspots to offer Internet access to passers-by.

In a similar fashion, an individual WLAN owner may open his private hotspot for public access, in order to increase wireless coverage in his city, without anticipating monetary compensation. Instead, he is either driven by pure altruism, or expects that his offering will be reciprocated in the same way by other community members when he is mobile [6]. The fact that usually access over residential WLANs is limited shows that we cannot rely on pure altruism for a large scale open wireless access scheme.

Following the above trend, commercial players have entered the scene, offering mediation services for the development of wireless communities and trying to make a profit out of this service. FON [7] is an example of such a scheme.

In many cases, a municipality initiative is behind wireless AP deployment in public spaces. Municipalities may get into agreements with private companies, permitting them to deploy their wireless solutions, in order to achieve citywide wireless coverage and inexpensive access for their citizens. This model has been adopted by the City of London [8], as well as the municipality of Philadelphia [9].

In this work, we focus on WLAN hotspot sharing communities that operate on a reciprocal basis and not pure altruism.

B. Interference mitigation

The problem of interference among neighbor IEEE 802.11-based WLANs has received a lot of research attention recently. It has its roots in their unplanned deployment and the fact that, at the same location, a limited number of APs can operate in an interference-free fashion. The IEEE 802.11 standard specifies that the available spectrum is divided in channels. In its 11b and 11g variants, only 3 of these channels are non-overlapping. More than 3 IEEE 802.11b/g APs in the same area means that some of them are assigned (at least partially) overlapping channels and thus suffer from interference, which in turn results in degraded performance. In typical urban deployments, the probability of coexistence of more than 3 WLANs at the same spot is high.

Among others, Akella et al. [10] have observed the severe interference problems that appear in present-day chaotic WLAN deployments and have proposed mechanisms based on power control and transmission rate adaptation to tackle them. Power control is also used in a work [11] with similar motivation and mainly for the case of high density enterprise or campus-wide WLANs.

A key issue in tackling with interference is the information used as input for mitigation strategies. AP-centric schemes [12] rely only on information collected locally at the AP sites. Although sometimes simpler, these schemes fail to capture spectrum usage conditions at the client locations. To limit this *Hidden Interference Problem*, information about client-perceived interference is crucial [13], [14].

In this work we do not study the development of new interference mitigation algorithms. Rather, we are interested in how the necessary input for them will be collected. Although their nature and internal workings affect interference information collection and representation, our mechanisms are more generic and can be adapted to the needs of these algorithms.

III. Peer-to-Peer WLAN sharing

In prior work [6], we have proposed that wireless Internet bandwidth be exchanged in a reciprocal manner; one shares his Internet connection with anonymous passers-by over WLANs with the anticipation that he will enjoy the same benefit from another peer when mobile. That is, our approach to the problem is based on the peer-to-peer paradigm. Thus, private WLAN owners have an incentive to contribute Internet bandwidth, given that they value much the mobile network access that they will enjoy as good contributors. We call our scheme the *Peer-to-Peer Wireless Network Confederation (P2PWNC)*.

To lower the entry barrier to the system, no registration with central authorities is required, nor any strong user identification scheme. Participants are identified by self-issued, uncertified public-secret key pairs. To join P2PWNC, users simply configure their access points for open access and install the necessary software. To facilitate our scheme's deployment, we have designed and implemented it to run on top of typical off-the-shelf WLAN hardware (e.g. Linux-based Linksys WRT54G wireless routers).

Accounting is based on digital proofs of service (*receipts*) that mobile users provide to visited APs. Receipts are stored in repositories that are distributed across the system. Each peer maintains its own repository, which represents his (partial) view of the system's history of service provisions.

Receipt repositories are the input to the *reciprocity algorithm*, which identifies good contributors and detects *free riders*. Each time a mobile user requests service, the reciprocity algorithm is invoked by the visited peer to decide whether the visitor is a good contributor and deserves to be reciprocated.

The visited peer, however, uses only his own view of the system to come to a decision. To assist in giving potential service providers a better view of their overall contribution and have better chances of getting access, visitors also supply parts of their own receipt repositories via a *gossiping protocol*. This is how the system's history gets distributed over the peers.

The output of the reciprocity algorithm is a *Subjective Reputation Metric (SRM)*. It encodes the subjective view of a prospective consumer's contribution to the community in the eyes of the prospective provider. SRM's values fall within the [0, 1] range. The higher the SRM, the more service the provider "owes" to the consumer. Details on how the SRM is calculated can be found in [6], [15].

IV. QoS as an incentive for interference reporting

Here, we extend the basic P2PWNC scheme with QoS awareness. Our purpose is on the one hand to reward good contributors with better service and, on the other hand, to provide some additional bandwidth to mobile users in exchange for spectrum usage information. We assume that the AP operator values such information and will use them as input to interference mitigation algorithms, but such algorithms are outside the scope of this work.

When more than one P2PWNC clients are connected to an AP, they get proportional bandwidth to their SRMs. We introduced such a QoS scheme in [16]. Also, if they wish to increase the bandwidth available to them, they need to perform a channel scan when requested and report their findings to the AP. Request frequency is a parameter defined by the AP operator and may depend on the exact interference mitigation algorithm used. Very frequent requests incur high overhead for wireless clients, as described in Section V-B, but, for reasonable request frequencies, the overhead is negligible. A client may or may not reply to such a request.

Also, a client may have a very low or zero SRM value. In this case, the only service he can get is due to interference reporting. The service such clients enjoy as a reward for their reports is adequate only for low-bandwidth web browsing or e-mail.

We assume that the hotspot owner has partitioned his Internet bandwidth into three parts. First, a fixed portion of it (B_o) is guaranteed for personal use. Then, another portion (B_p) is offered to visitors and is distributed proportianlly to each one's Subjective Reputation Metric (SRM_i, for visitor i), as returned by the execution of the reciprocity algorithm. The remaining B_b portion of the bandwidth[1] is a bonus for users who assist in interference reporting. Therefore, each one of an AP's n attached visitors is guaranteed a V_i portion of the bandwidth as follows:

$$V_i = B_p^{[i]} + B_b^{[i]} \tag{1}$$

where

$$B_p^{[i]} = \frac{SRM_i}{\epsilon + \sum_{j=1}^{n} SRM_j} \cdot B_p \tag{2}$$

and

$$B_b^{[i]} = r_i \cdot \frac{B_b}{n} \tag{3}$$

In (2), we use $0 < \epsilon << 1$ to avoid a zero denominator, in case all clients have zero SRMs. The r_i factor in (3) is the percentage of successful interference reports from client i. If a client always refuses to reply to spectrum sensing requests, then $r_i = 0$ and the client gets no extra bandwidth. A user can get at most $\frac{B_b}{n}$ bonus bandwidth.

The bandwidth an AP operator allocates as a reward for interference reporting depends on how much he values such information. In practice, it is expected that the AP owner will make sure that $B_b << B_p$. This is because a high B_b value may give peers the incentive not to contribute resources of their home WLANs and expect to consume enough bandwidth only in exchange for interference reports. It should be noted that $B_p^{[i]}$ and $B_b^{[i]}$ are adjusted by the AP when necessary, that is when a visitor joins or leaves the network and on each spectrum sensing request (to recompute r_i). Also, the AP reclaims unused bandwidth when no users are attached (B_p portion) or when some users do not contribute all requested interference reports and, thus, do not enjoy the full bonus.

V. Discussion

A. Implementation issues

The interference reporting scheme that we propose is straightforward to implement on current commercial WLAN hardware. Clients can collect spectrum usage information by performing channel scans when requested by the AP. Then, for each WLAN they detect, they report at least the following information:

{SSID, RSSI, channel number}

The above tuple can be extended to include positioning information where scanning took place, feasible in outdoor settings, provided that the reporting client is GPS-enabled. In the future, the use of the upcoming IEEE 802.11k [17] standard, which is designed for radio resource management, should also be considered for interference reporting and will facilitate the collection of more sophisticated information, such as channel load. Spectrum usage information can be reported to the AP using ASCII-based messages, in accordance with the P2PWNC protocol format.

As to the proposed QoS mechanisms, readily available tools [18] can be used to implement them, even on top of common Linux-based APs, such as Linksys WRT54GS wireless routers[2].

[1] Bandwidth values are expressed in bits/sec

[2] http://www.linksys.com

B. Performance overhead

1) Protocol overhead: The overhead of the reporting protocol is low. A request for spectrum sensing is expected to be performed infrequently, and the amount of the reported information is proportional to the number of APs that a station can sense. Usually, a client will detect no more than a few tens of nearby APs. For each one of them, it will have to report the AP's SSID, RSSI and channel number, as well as some additional information (e.g. client's GPS coordinates, if available). Assuming a text-based protocol for the representation of this information, less than a hundred bytes will be necessary for each detected network. Considering that scanning will not be performed very often (e.g. one scan per minute) the necessary bandwidth for the reporting process is expected to be low and not significantly affect application performance. However, this is one of the issues that are to be studied in future work.

2) Spectrum sensing effects: Scanning the available spectrum bands does not come without a cost for clients. A typical IEEE 802.11 active scan on all the available channels may require more than 0.3 seconds [19], [20]. While scanning, a station is unable to transmit/receive packets for its applications. Thus, application performance will degrade as scanning occurs more frequently. Quality degradation is expected to be more noticeable in delay-sensitive interactive applications, such as VoIP. For this purpose, we have performed some simple experiments to test the effects of spectrum sensing on user-perceived VoIP quality.

We emulated voice conversations by setting up bidirectional UDP flows between two laptop PCs, which were attached to a Linksys WRT54GS wireless router. One of them is connected to the router using IEEE 802.11b at 11Mbps at channel 11, with RTS/CTS disabled. This station is equipped with an Intel Pro/Wireless 2200BG card managed by the ipw2200 Linux driver. We measured that scanning took approximately 0.25 seconds on average. The other station is attached to the router's 100Mbps Ethernet port. Both PCs run Linux with a 2.6 kernel. We verified that our experiments took place in an interference free environment, where all collocated WLANs were operating at channels 1 to 6.

For each VoIP flow, we sent 50 packets per second with 20 bytes of audio payload each and 12 bytes for the RTP header. This traffic pattern corresponds to the G.729 codec, which is used by many available Wi-Fi VoIP phones. The 20 bytes of packet payload contain 20 msec of voice. Each voice call lasted for approximately 3 minutes. We have assumed that at the receiver end there is a dejitter buffer which introduces a 60 msec delay in the playout process.

We conducted a set of experiments where the wireless client periodically performed a channel scan while participating in a VoIP session, for various scanning frequencies. For each experiment, we collected delay and loss information for each packet for both the uplink and the downlink traffic. One-way end-to-end per packet delay was calculated comparing transmission and reception timestamps. The necessary synchronization between participating nodes was achieved using the Network Time Protocol (NTP) over the stations' Ethernet interfaces.

To estimate user perceived voice quality, we have used the evaluation scheme proposed in [21]. This scheme is a reduction of ITU-T's E-model [22] to transport level metrics, which are directly measurable in our testbed.

Using the proposed methodology, we can derive a score that represents the subjective quality of a voice call based only on network delay, jitter and packet loss measurements. For the codec configuration described above, this score (*R-score*) is given by the following formula:

$$\begin{aligned} R = {} & 94.2 - 0.024 \cdot (d_{network} + 85) \\ & - 0.11 \cdot (d_{network} - 92.3) \cdot H(d_{network} - 92.3) - 11 \\ & - 40 \cdot \ln[1 + 10 \cdot (e_{network} + (1 - e_{network}) \cdot e_{dejitter})] \end{aligned}$$

where:

- $d_{network}$ is the end to end network delay
- $e_{network}$ represents network loss
- $e_{dejitter}$ represents loss in the dejitter buffer
- $H(x) = 1 \quad if \quad x > 0; \quad 0 \quad otherwise$

For a call of acceptable quality, average R-score should be over 70.

Our results are summarized in Table I. We present the achieved *R-scores* for various scanning frequencies. In our experiments, the time period between two successive scans ranges from 1 second to 1 minute. Performing frequent spectrum sensing incurs significant overhead, which can result in unacceptable voice quality for both the upstream (wireless client to AP) and the downstream flows.

As to the factors that contribute to performance degradation, we have omitted network packet loss, because in all cases it was zero or negligible. On the other hand, although average end-to-end delay was not significant, periodical scans introduced jitter, causing the ratio of packets rejected at the receiver end due to excessive delay to increase ("Dejitter buffer loss" columns). In the extreme case when a site scan is requested once per second, more than 7% of the packets sent may be dropped at the dejitter buffer. It should be noticed that in all our experiments that involved scanning, uplink traffic suffered more than the downlink one.

It is reasonable to assume that an AP will not request for a channel scan more often than twice per minute, since spectrum usage conditions are not expected to change that fast. In this case, user-perceived performance degradation is minimal. This is an encouraging observation, since it implies that clients suffer little from interference reporting and, considering the QoS benefits they will enjoy, have the incentive to cooperate and assist in the interference mitigation process.

C. Truthful interference reporting

An issue that needs to be carefully considered is the incentives-compatibility of *truthful* interference reporting. Do clients have an incentive to cheat and fake their reports? Is there a means of verifying their validity?

TABLE I
USER-PERCEIVED VOIP QUALITY

Time between scans	**Uplink**			**Downlink**		
(sec)	Delay (msec)	Dejitter buffer loss	R-score	Delay (msec)	Dejitter buffer loss	R-score
1	13.58	0.0734	**57.83**	10.11	0.0391	**66.88**
2	8.78	0.0548	**62.37**	6.35	0.0287	**69.71**
3	6.18	0.0371	**68.05**	4.75	0.0211	**73.36**
4	4.97	0.0291	**70.22**	3.94	0.0169	**75.29**
5	3.82	0.0210	**73.76**	3.24	0.0123	**76.79**
10	2.54	0.0121	**75.87**	2.32	0.0068	**78.51**
30	1.50	0.0037	**79.95**	1.44	0.0022	**80.43**
60	1.28	0.0021	**80.46**	1.26	0.0011	**80.80**
no scanning	1.04	0.0000	**81.14**	1.07	0.0000	**81.13**

As evident from the results of Section V-B, frequent interference reports incur significant application performance degradation. In such cases, clients may be tempted to avoid scanning and supply counterfeit reports, in order to receive their $B_b^{[i]}$ bandwidth bonus. On the other hand, if they do so, and given that a significant number of truthful clients is associated with the same AP, the AP may apply a majority rule and detect, with some error probability, the misbehaving node, comparing its reports with the reports of the other clients. In such cases, the AP may apply punishment mechanisms on the peer. If the maximum possible $B_b^{[i]}$ for a client is sufficiently low, he may not risk providing fake reports and will simply not reply to spectrum sensing requests.

Another potential attack is for a client to perform a single scan and keep replaying the same (or similar) reports on each subsequent sensing request, thus keeping the value of r_i and, consequently, $B_b^{[i]}$ (see Eq. 3), high. This attack is harder to detect, but its effects are less harmful, since the client provides an amount of true information.

In all above cases, though, considering that sensing requests are not performed frequently and do not significantly affect application performance, incentives for performing such attacks are not clear. Part of our future research efforts will focus on investigating the space of potential attacks to our mechanisms and devising methods for effectively combating them.

VI. CONCLUSION

In this work, we proposed that reciprocal WLAN sharing among communities of users can be jointly considered with interference mitigation. To this end, based on our prior work on peer-to-peer Wi-Fi sharing, we proposed that QoS benefits can act as an incentive for roaming users to report spectrum usage information to the APs they attach to. This information can be utilized by sophisticated interference mitigation schemes. To verify that the overhead of spectrum sensing, under reasonable assumptions, is not high enough to discourage user participation, we carried out experiments on a real WLAN testbed to investigate its effects on user-perceived VoIP quality.

REFERENCES

[1] "Athens Wireless Metropolitan Network," http://www.awmn.net.
[2] "Seattle Wireless," http://www.seattlewireless.net.
[3] "CUWiN - Community Wireless," http://www.cuwireless.net/.
[4] B. Milic and M. Malek, "Analyzing large scale real-world wireless multihop network," *IEEE Communications Letters*, vol. 11, pp. 580–582, July 2007.
[5] R. D. J. Kramer, A. Lopez, and A. M. J. Koonen, "Municipal broadband access networks in the Netherlands - three successful cases, and how New Europe may benefit," in *Proc. AccessNets '06*, Athens, Greece, September 2006.
[6] E. C. Efstathiou, P. A. Frangoudis, and G. C. Polyzos, "Stimulating participation in wireless community networks," in *Proc. IEEE INFOCOM 2006*, Barcelona, Spain, April 2006.
[7] "FON," http://en.fon.com.
[8] "City of London goes wireless: Launch of Europe's most advanced WiFi network," http://www.cityoflondon.gov.uk/Corporation/media_centre/files2007/73_07.htm.
[9] "Wireless Philadelphia Executive Committee," http://www.phila.gov/wireless.
[10] A. Akella, G. Judd, S. Seshan, and P. Steenkiste, "Self-management in chaotic wireless deployments," in *Proc. ACM MOBICOM 2005*, Cologne, Germany, August 2005.
[11] V. Mhatre, K. Papagiannaki, and F. Baccelli, "Interference mitigation through power control in high density 802.11 WLANs," in *Proc. IEEE INFOCOM 2007*, Anchorage, Alaska, USA, May 2007.
[12] K. K. Leung and B.-J. Kim, "Frequency assignment for multi-cell IEEE 802.11 wireless networks," in *Proc. VTC 2003*, Orlando, FL, USA, October 2003.
[13] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan, and W. A. Arbaugh, "A Client-Driven Approach for Channel Management in Wireless LANs," in *Proc. IEEE INFOCOM 2006*, Barcelona, Spain, April 2006.
[14] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot, "Measurement-based self organization of interfering 802.11 wireless access networks," in *Proc. IEEE INFOCOM 2007*, Anchorage, Alaska, USA, May 2007.
[15] E. C. Efstathiou, "A Peer-to-Peer Approach to Sharing Wireless Local Area Networks," Ph.D. dissertation, Athens University of Economics and Business, 2006.
[16] E. C. Efstathiou, F. A. Elianos, P. A. Frangoudis, V. P. Kemerlis, D. C. Paraskevaidis, E. C. Stefanis, and G. C. Polyzos, "Public infrastructures for Internet access in metropolitan areas," in *Proc. AccessNets '06*, Athens, Greece, September 2006.
[17] IEEE 802.11 WG, *Draft Supplement to Standard for Telecommunications and Information Exchange Between Systems - LAN/MAN Specific Requirements - Part 11: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Specification for Radio Resource Measurement, IEEE 802.11k/D0.7*, The Institute of Electrical and Electronics Engineers, Inc., New York, USA, 2003.
[18] "Linux iproute2/tc," http://linux-net.osdl.org/index.php/Iproute2.
[19] A. Mishra, M. Shin, and W. A. Arbaugh, "An empirical analysis of the IEEE 802.11 MAC layer handoff process," *Computer Communication Review*, vol. 33, no. 2, pp. 93–102, 2003.
[20] I. Ramani and S. Savage, "SyncScan: practical fast handoff for 802.11 infrastructure networks," in *Proc. IEEE INFOCOM 2005*, Miami, FL, USA, March 2005.
[21] R. G. Cole and J. H. Rosenbluth, "Voice over IP performance monitoring," *Computer Communication Review*, vol. 31, no. 2, pp. 9–24, 2001.
[22] ITU-T Recommendation G.107, "The E-model, a computational model for use in transmission planning," December 1998.