Experiences and Investigations with Heraklion *MESH*: An Experimental Metropolitan Multi-Radio Mesh Network[*]

Manolis Delakis, Kostas Mathioudakis, Nikolaos Petroulakis, Vasilios A. Siris[†]

Institute of Computer Science, FORTH
P. O. Box 1385, GR 711 10 Heraklion, Crete, Greece

ABSTRACT

We present our experiences and investigations with an experimental metropolitan multi-radio mesh network that covers an area of approximately 60 Km^2 in the city of Heraklion, Crete. The network consists of 14 nodes, among which six are core nodes with up to four 802.11a wireless interfaces each, and an additional wireless interface for management and monitoring. The distance between core mesh nodes varies from 1.6 to 5 Km, and the wireless mesh network contains two gateways that connect it to a fixed network. Our investigations have considered the online monitoring of the links between core nodes, the interference between the wireless interfaces located in the same mesh node, channel assignment procedures, and the use of the metropolitan mesh network for real-time collection of electromagnetic field (EMF) measurements.

Categories and Subject Descriptors: C.2.1 [Computer-Communication Networks]: Network Architecture and Design–*wireless communication*

General Terms: Design, Experimentation.

Keywords: long-distance 802.11a links, inter-link interference, channel assignment, electromagnetic field monitoring.

1. INTRODUCTION

Wireless multi-radio multi-channel mesh networks have the potential to provide ubiquitous and ultra high-speed broadband access in urban and rural areas, to both fixed and mobile users, with low operation and management costs. To investigate issues related to the management and perfor-

mance of a multi-radio mesh network in an actual metropolitan environment, we have deployed an experimental multi-radio mesh network that covers an area of approximately 60 Km^2 in the city of Heraklion, Crete, Greece. Our objective is to use the network as a metropolitan scale test-bed to investigate the performance of a multi-radio mesh network built from commodity components with 1 to 5 Km links, to evaluate channel assignment procedures for efficiently utilizing the wireless spectrum, to investigate MAC/network layer mechanisms and routing metrics for supporting performance guarantees in multi-radio, multi-channel, multi-rate mesh networks, and to investigate innovative applications that require pervasive and high-speed broadband access. Use of commodity IEEE 802.11 technology can lead to significantly lower costs compared to other technologies, such as IEEE 802.16.

Other mesh and/or long-distance 802.11 networks include the 802.11b-based Digital Gangetic Plains rural area test-bed with $1 - 23$ Km links [6], the WiLDNet network with $50 - 100$ Km links [8], the Roofnet network which considers single-radio mesh nodes [3], the Quail Ridge wireless mesh network [13], and other city-wide mesh networks[1]. Unlike the above, our test-bed consists of 1 to 5 Km city links, and utilizes multi-radio mesh nodes with directional antennas.

In this paper we report our experiences and investigations related to the design of the metropolitan mesh network (Section 2), the interference between different wireless interfaces in the same mesh node (Section 3.1), online monitoring of the links between core mesh nodes (Section 3.2), channel assignment procedures (Section 4), and the use of the metropolitan mesh network for real-time collection of electromagnetic field (EMF) measurements (Section 5). Ongoing and future research directions are presented in Section 6.

2. METROPOLITAN MULTI-RADIO MESH NETWORK DESIGN

2.1 Multi-radio mesh node

Each multi-radio mesh node consists of a mini-ITX board (EPIA SP 13000, 1.3 GHz C3 CPU, 512 MB DDR400 memory) and a 40 GB 2.5" HDD. A four slot mini PCI to PCI adapter (MikroTik RouterBOARD 14) holds four 802.11a/g mini PCI adapters (NL-5354 MP PLUS Aries 2, Atheros-based High Power Super A/G dual Band 802.11a/b/g). The
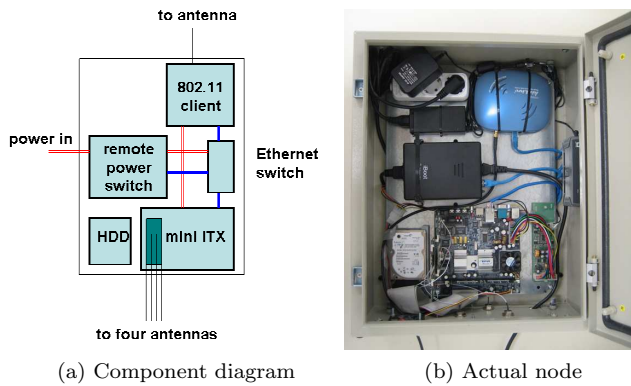
(a) Component diagram   (b) Actual node

Figure 1: Multi-radio mesh node.

mini-ITX runs Gentoo 2006 i686 Linux (2.6.18 kernel) with the MadWiFi driver version 0.9.2. Finally, the nodes run OLSR daemon version 0.4.10 (by olsr.org), which implements the Optimized Link State Routing (OLSR) protocol.

One of our design requirements was to allow remote management, monitoring, and recovery of the mesh nodes, even in situations when a node's mini-ITX board crashes or its wireless interfaces are down. To address this requirement we added to each mesh node an additional 802.11a client, Figure 1, which connects to a management and monitoring network that operates in parallel to the experimental mesh network. Additionally, to enable remote recovery of the mesh node's mini-ITX board, each node contains an intelligent remote power switch (Dataprobe iBoot), Figure 1; this supports off/on power switching through a web interface, and timed power reboots based on the results from the power switch pinging other devices (the mini-ITX board or some remote device).

2.2 Topology

The metropolitan mesh network covers an area of approximately 60 Km^2 and currently contains 14 nodes, Figure 2, among which six are core mesh nodes, whose design was discussed in the previous section. The distance and antennas used for the links between core mesh nodes[2] are shown in Table 1. Each wireless interface is assigned a static IP address, and the OLSR protocol is used for routing traffic in the network. The mesh test-bed is connected to a fixed network through two nodes (FORTH and UoC).

3. INTERFERENCE AND PERFORMANCE MONITORING

Most prior measurement studies of outdoor 802.11 links focus on measuring the path loss and the time correlation of losses, and how the loss is affected by factors such as received signal strength, link distance, interference, weather

---

[2]Two core mesh nodes are under deployment, and are not shown in Table 1.

Table 1: Links between core mesh nodes

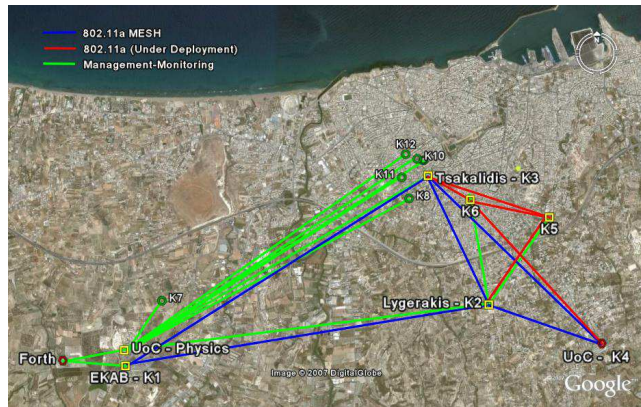| Link	Distance (Km)	Antennas
Ekab-Lygerakis	5.0	29 dBi grid-21 dBi panel
Ekab-Tsakalidis	4.9	29 dBi grid-21 dBi panel
Lygerakis-Tsakalidis	2.0	21 dBi-19 dBi panel
UoC-Lygerakis	1.6	21 dBi-21 dBi panel
UoC-Tsakalidis	3.3	21 dBi-19 dBi panel

Figure 2: Heraklion MESH.

conditions, and technology (802.11b/g) [1, 6, 2, 11]. Work on the impact of adjacent channel interference for 802.11b/g is contained in [6, 5] and for 802.11a in [4]; the latter focuses on measuring the impact on the signal-to-noise ratio, and considers link distances of 60 meters.

One the other hand, our work considers metropolitan links with distances 1.6 - 5 Km, and focuses on the impact interference has on the throughput above the MAC layer. To the best of our knowledge, this is the first throughput and interference study that investigates long-distance 802.11a links.

3.1 Interference experiments

3.1.1 Interference between wireless interfaces in a laboratory environment

We first investigate, in a laboratory environment, the interference between 802.11a wireless interfaces and the existence of cross-talk between the four wireless adapters in a single mesh node. The experiment scenario involved two flows entering two wireless interfaces of a mesh node and exiting through the other two interfaces. Each of the four mesh node interfaces were connected to a panel antenna, and the clients where laptops. We consider two channel assignment cases: in the first, the interfaces are assigned neighboring channels, and in the second case the interfaces are assigned channels with a one channel separation (i.e., leaving one unassigned channel between the assigned channels), which corresponds to channel distance 2 in Table 2. The throughput achieved by the two flows and the 95% confidence interval is shown in Table 2. Observe that there is significant interference, which results in very low throughput, when the wireless interfaces are assigned neighboring channels (channel distance 1). On the other hand, when there is a one channel separation the throughput achieved by each flow is close[3] to the average throughput achieved when each flow is alone (with a one channel separation, 24.85 ±0.05 Mbps) and the throughput over one wireless link (24.85 ±0.11 Mbps). Finally, unlike [10], our results do not show any significant crosstalk among the four mini PCI wireless adapters (which recall are all located on the same mini PCI to PCI adapter), nor any interference due to the small distance between the wireless adapters, as reported in [6]. We attribute this to the newer hardware we have used, which apparently has improved electrical shielding properties.

---

[3]This depends on the type and distance of the antennas.

Table 2: Throughput (Mbps) for different channel assignments. Each flow traversed two interfaces (one incoming and one outgoing) of a single mesh node in a laboratory environment.

Channel dist.	Flow 1		Flow 2	
	avg	conf. int.	avg	conf. int.
1 (36-40,44-48)	10.61	±0.51	9.65	±0.36
2 (36-44,52-60)	24.58	±0.04	24.47	±0.11

3.1.2 Inter-link interference in a metropolitan environment

Next we investigate the interference between metropolitan distance links. One of the two interfaces of each link under investigation is located in the same mesh node (node *Lygerakis*). In particular, we consider the link pair *Lygerakis - Tsakalidis* and *Lygerakis - UoC*, and the link pair *Lygerakis - Tsakalidis* and *Lygerakis - Ekab*. For the first pair, the two interfaces in node *Lygerakis* are connected to two 21 dBi panel antennas (Table 1), which are both on the same mast with a distance of approximately 0.75 meters, and have a relative angle of approximately 150 degrees. For the second pair, the two interfaces in node *Lygerakis* are again connected to two 21 dBi panel antennas (Table 1), which are however on a different mast with a distance of approximately 2.5 meters, and have a relative angle of approximately 90 degrees.

Each experiment we present below shows the average from 10 runs, each run lasting for 100 seconds. Also shown is the 95% confidence interval. Finally, the experiments involved UDP traffic with rate 3 Mbps, generated using the iperf tool.

Tx/Rx in the same node: We first consider the interference between links when a receive and transmit interface exists in the same mesh node. In particular, two 3 Mbps UDP stream are transmitted over the links *Lygerakis → Tsakalidis* (2 Km distance) and *UoC → Lygerakis* (1.6 Km distance). Note that the two streams are independent, and the iperf sender for the first stream is located in a workstation[4] connected to our internal laboratory network.

Table 3 shows the achieved throughput (measured at the iperf receiver) by each UDP flow, for three different channel assignments. Observe that when both links are assigned the same channel (36-36 in Table 3), the transmitter significantly affects the receiver (see throughput for Rx *Lygerakis ← UoC*), both located in the same mesh node (*Lygerakis*). When the two links are assigned neighboring channels (40-36), the interference is significantly reduced, but still appears to exist. On the other hand, when there is a one channel separation (44-36), there is no interference and the throughput is essentially equal to the UDP sending rate. In comparison, in a laboratory environment (Table 2) the interference between neighboring channels is much more significant than in a metropolitan environment (Table 3); this is because the metropolitan experiments used directional (panel) antennas with different orientations, whereas the laboratory experiments considered laptops and panel antennas with the same orientation.

We performed the same experiment as the one described above, but with two 3 Mbps UDP stream transmitted over

[4]This was done to avoid loading the CPU of the mesh node (*Lygerakis*) where the two interfaces under examination resided, since an iperf sender consumes a significant amount of CPU resources.

Table 3: Throughput (Mbps) of two flows when receiver and transmitter is in same mesh node (*Lygerakis*), and antennas on same mast at distance ≈ 0.75 meters.

Channel dist.	Tx *Lyger. → Tsakal.*		Rx *Lyger. ← UoC*	
	avg	conf. int.	avg	conf. int.
0 (36-36)	2.970	±0.004	2.358	±0.125
1 (40-36)	2.995	±0.002	2.976	±0.006
2 (44-36)	2.997	±0.002	2.997	±0.002

Table 4: Throughput (Mbps) of two flows when receiver and transmitter is in same mesh node (*Lygerakis*), and antennas on different mast at distance ≈ 2.5 meters.

Channel dist.	Tx *Lyger. → Ekab*		Rx *Lyger. ← Tsakal.*	
	avg	conf. int.	avg	conf. int.
0 (36-36)	3	≤ ±0.001	2.75	±0.1
1 (40-36)	3	≤ ±0.001	3	≤ ±0.001
2 (44-36)	3	≤ ±0.001	3	≤ ±0.001

the links *Lygerakis → Ekab* (5 Km distance) and *Tsakalidis → Lygerakis* (2 Km distance). Table 4 shows that the interference in this case is lower than in the previous experiment (Table 3). This is due to the larger distance (approximately 2.5 meters) between the antennas for node *Lygerakis* that correspond to the above two links.

Rx/Rx in the same node: Next we investigate the interference between links when two receive interfaces are located in the same mesh node. In particular, two 3 Mbps UDP streams are transmitted over the links *Tsakalidis → Lygerakis* and *UoC → Lygerakis*. Table 5 shows that the achieved throughput for both flows is the same, and essentially equal to the UDP sending rate, even when the links are assigned the same channel. Hence, the interference between the two links in this case is not significant.

3.2 Online performance monitoring

Our experience has shown that it is important to continuously monitor core mesh network links, since this allows us to quickly identify and to help understand anomalous link behavior. For this reason, we have developed a set of perl and shell scripts that continuously monitor important performance metrics for all links between core nodes. The metrics include the signal-to-noise ratio (SNR), transmission rate, MAC and physical layer errors, two-way delay, and throughput. The scripts are executed every five minutes, except the scripts for measuring the throughput, which are executed every 30 minutes. The collected data is stored in an RRD database, and the corresponding daily and weekly graphs are made available through an http server[5] using the RRDTool, Figure 3.

[5]See http://www.ics.forth.gr/HMESH

Table 5: Throughput (Mbps) of two flows when two receivers are in same mesh node (*Lygerakis*), and antennas on same mast at distance ≈ 0.75 meters.

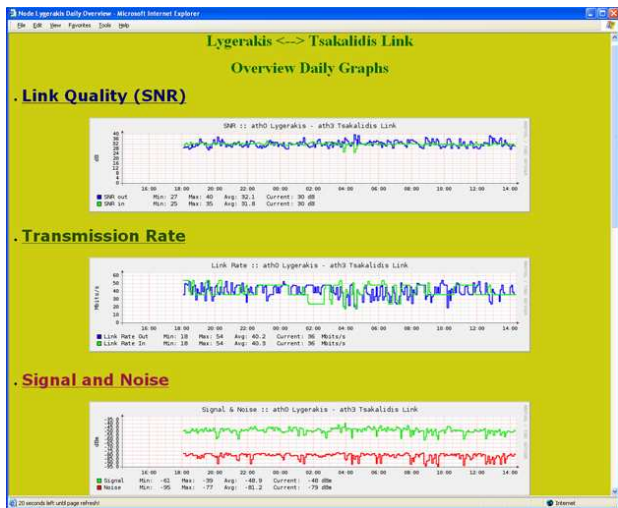| Channel dist.	Rx *Lyger. ← Tsakal.*		Rx *Lyger. ← UoC*	
	avg	conf. int.	avg	conf. int.
0 (36-36)	2.996	±0.002	2.996	±0.002
1 (40-36)	2.996	±0.002	3	≤ ±0.001
2 (44-36)	3	≤ ±0.001	3	≤ ±0.001

Figure 3: Web display of online monitoring tool.

Table 6 shows typical values we have observed for the signal-to-noise ratio (SNR) and transmission rate of the core links in an interval of 24 hours. These results show that the links are asymmetric, and both the link quality and transmission rate varies throughout the day. Moreover, the variation of the link quality is different for different links.

4. CHANNEL ASSIGNMENT

Next we consider the problem of channel assignment in the metropolitan mesh network. Our objective is to identify issues and investigate various procedures for channel assignment that can be applied to an actual metropolitan mesh network. These procedures should consider the fact that mesh nodes have multiple radio interfaces and use directional antennas.[6]

Next we briefly summarize some related work on channel assignment in mesh networks. The work of [9] considers a distributed load-aware channel assignment algorithm. The order (priority) in which links are assigned channels is based on their distance from a single gateway. A link selects the channel with the least channel load, where the channel load is a weighted combination of the aggregate traffic load and the number of nodes using the channel. The work of [8] considers a centralized interference-aware channel assignment algorithm, where the link ordering is based on the distance

---

[6]The current mesh deployment contains only point-to-point links, but in the future we also expect to have multipoint links. The procedures we describe can be extended to multipoint links, e.g. by considering average SNR and delay values over all multipoint links.

Table 6: Link SNR and transmission rate

Link	SNR min,max,avg	Rate min,max,avg (Mbps)
Ekab-Lygerakis	18,21,19	6,36,26
Lygerakis-Ekab	20,23,21	6,36,24
Ekab-Tsakalidis	16,20,18	6,54,23
Tsakalidis-Ekab	16,23,17	6,48,16
Lygerakis-Tsakalidis	25,35,31	18,54,40
Tsakalidis-Lygerakis	27,40,32	18,54,40
Lygerakis-UoC	18,23,20	12,36,22
UoC-Lygerakis	10,22,19	12,48,23
Tsakalidis-UoC	13,32,22	6,48,21
UoC-Tsakalidis	18,22,20	12,36,23

to a single gateway, and each link selects the channel with the best rank, which is the average rank based on channel utilization and the number of interfering radios. The work of [7] considers a distributed algorithm where each interface greedily selects the channel with the smallest aggregate interference cost for all interfaces within its range; the interference cost function is a linear function of the channel separation. Finally, the work in [12] proposes a centralized (based on the Tabu search) and distributed greedy algorithm for channel assignment that minimizes the aggregate interference, which assumes a priori knowledge of the interference between two channels and considers the traffic traversing the mesh network links.

The procedure for channel assignment that we consider greedily selects for each link the best channel that can be assigned to it, based on local information. The following three issues related to the above procedure can be identified:

- Link order: This is the order that links are assigned a channel. Alternatives are fixed ordering (e.g., based on the distance to a fixed network gateway and the average utilization of links), ordering based on increasing SNR values, and random ordering.

- Consideration of inter-link interference: Two alternatives are the following: 1) generate test-traffic on links with assigned channels in order to account for the interference these links create to others, and 2) consider an a priori model for inter-link interference that avoids the assignment of the same or neighboring channels to two interfaces belonging to the same mesh node or within the range of each other; such an exclusion of channels is motivated from the results of the interference experiments in Section 3.1.

- Channel assignment metric: This is the metric used to select the channel to assign to a particular link. We consider three metrics: 1) one-way SNR, 2) two-way SNR (this is the average SNR on the two interfaces belonging to the same link), and 3) round-trip delay.

4.1 Experimental results

4.1.1 Comparison of channel assignment metrics

We performed 10 runs with each of the three metrics identified above. The results indicate that 83% of the time, the one-way and two-way SNR yield the same three best channels. This suggests that despite the asymmetry of the links in the two directions (Table 6), the SNR in both directions is correlated. Additionally, for the two-way SNR metric, 80% of the time the SNR difference between the 1st and 2nd best channels is less than 1.5 dB, whereas in 55% of the time the SNR difference between the 1st and 3rd best channels is less than 1.5 dB. This difference is smaller than the typical variations of the SNR during a day, Table 6.

Table 7 shows the channel assignment and the average delay on all core mesh links with one application of the different metrics. Observe that there are links for which the different metrics assign the same channel. Moreover, the average delay across all core mesh links, when channels are assigned according to the three metrics, is very close.

4.1.2 Consideration of inter-link interference

Above we mentioned one approach for considering inter-link interference by generating test-traffic on links with an

Table 7: Channel assignment for links *Ekab-Lyg*, *Ekab-Tsak*, *UoC-Tsak*, *UoC-Lyg*, *Tsak-Lyg*).

Metric	Channels	Avg delay (ms)
one-way SNR	{ 112,48,124,56,132}	0.71
two-way SNR	{ 116,48,124,60,140}	0.72
two-way delay	{ 116,60,124,48,140}	0.74

Table 8: Channel assignment for links *Ekab-Lyg*, *Ekab-Tsak*, *UoC-Tsak*, *UoC-Lyg*, *Tsak-Lyg*.

Type	Channels	Avg delay (ms)
same channel	{120,60,128,60,124}	1.04
neighboring channel	{124,52,132,56,140}	0.83
1-channel separation	{116,48,124,60,140}	0.72

assigned channel. Test-traffic can be intermittently generated using ping or continuously generated using iperf. With the first alternative, we have observed that approximately 60% of the assignments had identical or neighboring channels for different links. With continuous traffic generation using iperf, there were no same channel assignments and only approximately 20% neighboring channel assignments.

To illustrate the effects of identical and neighboring channel assignments, Table 8 shows the total average delay for different channel assignments. Observe that when there are links that are assigned the same or neighboring channels, the delay is higher than when links are assigned channels with a one channel separation.

5. REAL-TIME COLLECTION OF EMF MEASUREMENTS

Next we discuss the use of the metropolitan network for real-time collection of electromagnetic field (EMF) measurements. This is important given the growing concern about electromagnetic radiation, especially from mobile telephony systems. Even though the exact implications of electromagnetic radiation on public health are not known, it is important to continuously monitor the EMF levels and check conformance to national and international threshold levels.

5.1 EMF monitoring node

EMF monitoring can be achieved by using low-cost EMF monitoring devices placed in nodes that contain, similar to a mesh node, a mini-ITX board (EPIA SP 13000, 1.3 GHz C3 CPU, 512 MB DDR400 memory) and a 80 GB 2.5" HDD, Figure 4. The EMF monitoring node can connect to a mesh node either through a wireless or wired interface. The mini-ITX board runs Windows XP Pro, since the software for controlling and collecting measurements from the EMF monitoring device was available only in this operating system. The mini-ITX board also runs an http server, which allows publishing and accessing the collected EMF measurements from a remote station over the metropolitan mesh network.

The advantages from using a metropolitan mesh network for collecting EMF measurements include the following:

- Higher range: EMF meters can be used to monitor frequencies up to 7 GHz, which is higher than the capabilities from specialized stand-alone EMF monitors, whose range is typically limited to 3 GHz.
- Real-time remote measurement collection: EMF meters with real-time monitoring capabilities together with a metropolitan coverage mesh network allow real-time remote collection of EFM measurement data.

Figure 4: The EMF monitoring node contains a mini-ITX and a low cost EMF monitor (Spectran 60X0 analyzer, which can measure up to 7 GHz).

Table 9: Frequency bands

Bands	Frequencies (MHz)			
FM	75-108			
TV	VHF		UHF	
	174-230		470-862	
GSM	900		1800	
	Uplink	Downlink	Uplink	Downlink
	880-915	930-960	1720-1785	1825-1880
802.11	b / g		a	
	2400-2483.5		5150-5350	5470-5725

- Low cost: Small (handheld) EMF meters with advanced spectrum analyzer capabilities are significantly cheaper than stand-alone EMF monitoring devices with remote communication (GSM) capabilities.
- Advanced flexibility: Together with a general purpose CPU board, the EMF monitor can be controlled remotely to collect measurements in different frequency ranges (bands) and different time windows.

5.2 EMF monitoring capabilities

Next we discuss three types of measurement capabilities that we have implemented, which demonstrate the system's flexibility in terms of measurement granularity and aggregation: per-band, per-operator, and time-series monitoring.

5.2.1 Per-band monitoring

This displays the EMF radiation in different bands, such as radio (FM), TV, GSM, and 802.11 in the 2.4 GHz and 5 GHz range, Table 9. Measurements are in units of mV/m and intervals of approximately 6 minutes[7], Figure 5.

5.2.2 Per-operator monitoring

Per-operator monitoring displays the EMF radiation in various frequencies assigned to the different (Greek) mobile operators, Figure 6. This can be used to verify conformance to national exposure thresholds and to detect undeclared or rogue base station antennas.

5.2.3 Time-series monitoring

The third type of measurement displays per-band EMF measurements in the form of a time series. This allows real-time monitoring of the fluctuations of EMF radiation levels in different bands and different time intervals during the course of a day, Figure 7.

---

[7]Six minutes is the time interval recommended by Greece's National Telecommunications Committee (EETT).
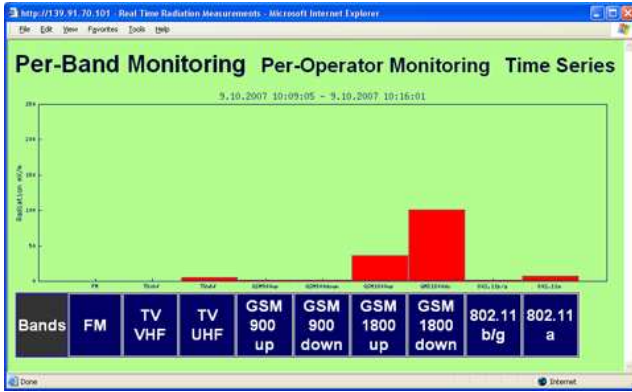
Figure 5: Per-band monitoring. The measurements are average values in 6 minute intervals.
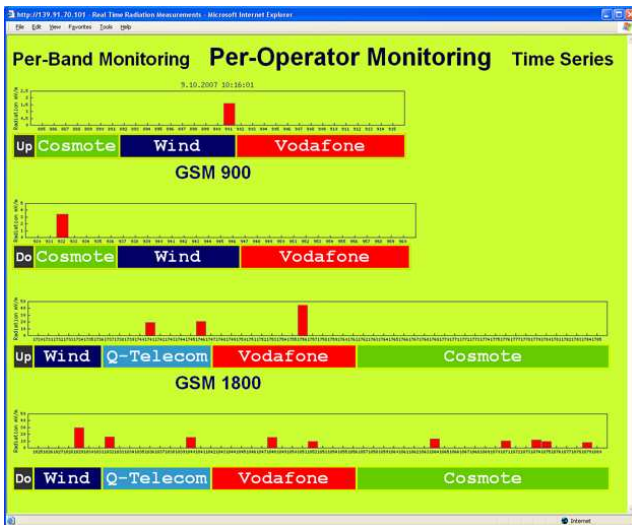

Figure 6: Per-operator monitoring. The measurements represent EMF levels in frequencies assigned to the different Greek mobile operators.
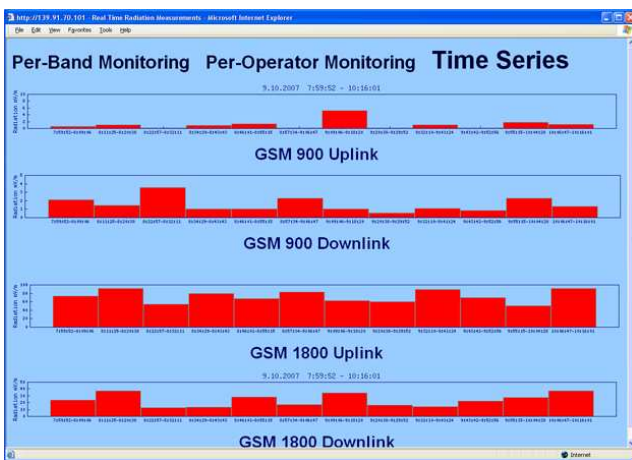

Figure 7: Time-series per-band monitoring.

6. CONCLUSION AND ONGOING WORK

We presented some of our experiences and investigations with an experimental metropolitan multi-radio mesh network, and in particular how interference between metropolitan 802.11 links affects throughput, an online performance monitoring system for the mesh network, a channel assignment procedure with initial experimental results, and the use of the metropolitan network for real-time collection of EMF measurements.

Ongoing work is investigating how the performance of metropolitan links is influenced by distance, transmission power, and rate control algorithm. Related to channel assignment, we are investigating the construction of a more refined interference model based on actual measurements, while also considering the actual link utilization; this would be more accurate than the simple model presented in this paper, and is faster and less cumbersome than generating test-traffic on all links with assigned channels. Other related work is investigating contention-aware routing metrics.

7. REFERENCES

[1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-Level Measurement from an 802.11b Mesh Network. In *Proc. of ACM Sigmetrics*, 2004.

[2] G. Bianchi, F. Formisano, and D. Giustiniano. 802.11b/g Link Level Measurements for an Outdoor Wireless Campus Network. In *Proc. of IEEE WoWMoM*, 2006.

[3] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and Evaluation of an Unplanned 802.11b Mesh Network. In *Proc. of ACM MOBICOM*, 2005.

[4] C.-M. Cheng, P.-H. Hsiao, H. T. Kung, and D. Vlah. Adjacent Channel Interference in Dual-Radio 802.11a Nodes and Its Impact on Multi-hop Networking. In *Proc. of IEEE Globecom*, 2006.

[5] T. Ireland, A. Nyzio, M. Zink, and J. Kurose. The Impact of Directional Antenna Orientation, Spacing, and Channel Separation on Long-distance Multi-hop 802.11g Networks: A Measurement Study. In *Proc. of 3rd Int'l Workshop on Wireless Network Measurement (WiNMee)*, 2007.

[6] K.Chebrolu, B. Raman, and S. Sen. Long-Distance 802.11b Links: Performance Measurements and Experience. In *Proc. of ACM MOBICOM*, 2006.

[7] B. Ko, V. Misra, J. Padhye, and D. Rubenstein. Distributed Channel Assignment in Multi-Radio 802.11 Mesh Networks. In *Proc. of IEEE WCNC*, 2007.

[8] K. N. Ramachandran, E. M. Belding, K. C. Almeroth, and M. M. Buddhikot. Interference-Aware Channel Assignment in Wulti-Radio Wireless Mesh Networks. In *Proc. of IEEE INFOCOM*, 2006.

[9] A. Raniwala and T. Chiuch. Architecture and Algorithms for an IEEE 802.11-Based Multi-Channel Wireless Mesh Network. In *Proc. of IEEE INFOCOM*, 2005.

[10] J. Robinson, K. Papagiannaki, C. Diot, X. Guo, and L. Krishnamurthy. Experimenting with a Multi-Radio Mesh Networking Testbed. In *Proc. of WinMee*, 2005.

[11] A. Sheth, S. Nedevschi, R. Patra, S. Surana, L. Subramanian, and E. Brewer. Loss Characterization in WiFi-based Long Distance Networks. In *Proc. of IEEE INFOCOM*, 2007.

[12] A. P. Subramanian, H. Gupta, and S. R. Das. Minimum-Interference Channel Assignment in Multi-Radio Wireless Mesh Networks. In *Proc. of IEEE SECON*, 2007.

[13] D. Wu, D. Gupta, and P. Mohapatra. Quail Ridge Reserve Wireless Mesh Network: Experiences, Challenges and Findings. In *Proc. IEEE TRIDENTCOM*, 2007.