Experiences with a Metropolitan Multiradio Wireless Mesh Network: Design, Performance, and Application

Vasilios A. Siris, Foundation for Research and Technology — Hellas (FORTH) and Athens University of Economics and Business
Elias Z. Tragos and Nikolaos E. Petroulakis, FORTH

ABSTRACT

Wireless mesh networks comprise nodes with multiple radio interfaces, and can provide low-cost high-speed Internet access or connectivity for data transfer. In this article we report our experiences and investigations with an experimental metropolitan multiradio mesh network that covers an area of approximately 60 km^2 in the city of Heraklion, Crete. We present the design and deployment of the network, experiments to quantify the network's performance, and an application that runs on top of it and exploits it's low-cost wide-area connectivity. The metropolitan network consists of 16 nodes, among which six are core nodes, each with up to four 802.11a wireless interfaces and an additional wireless interface for management and monitoring. The distance between core mesh nodes varies from 1.6 to 5 km, and the mesh network contains two gateways that connect it to a wired network. Our performance experiments involve rate, power, and channel control for long-distance metropolitan links, and include investigations of the timescales for the operation for these mechanisms. Finally, we present a system for continuous online electromagnetic field monitoring and spectrum sensing, which utilizes the metropolitan mesh network for collecting wide-area measurements from low-cost EMF measurement devices.

INTRODUCTION

Wireless multiradio multichannel mesh networks have the potential to provide ubiquitous and 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 performance of a multiradio mesh network in an actual metropolitan environment, we have deployed an experimental multiradio 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 testbed to:

- Investigate the performance of a multiradio mesh network, built from commodity components and containing 1 to 5 km links with directional antennas
- Evaluate channel assignment procedures for efficient wireless spectrum utilization
- Investigate medium access control (MAC)/network layer mechanisms and routing metrics for supporting performance guarantees in multiradio multichannel multirate mesh networks
- Investigate innovative applications that require pervasive, wide-area, and high-speed data transfer

The metropolitan testbed is built from commodity IEEE 802.11 components, which leads to significantly lower costs compared to other technologies, such as 802.16. Several other mesh and/or long-distance 802.11 networks have been developed worldwide. The 802.11b-based Digital Gangetic Plains rural area testbed contains 1–23 km links [1] with directional antennas. The WiLDNet network has 50–100 km links [2], which use a modified IEEE 802.11 MAC protocol to operate over such very long distances. The Roofnet network [3] considers only single-radio mesh nodes, in contrast to the multiradio nodes in our mesh testbed. The Quail Ridge wireless mesh network [4] contains 34 mesh nodes and 802.11g (2.4 GHz) links, most with distance smaller than 1 km. Rice University has deployed a metropolitan mesh network in East-End Houston with 21 mesh nodes in an area of approximately 3 km^2 [5], using uni- or omnidirectional antennas and 802.11g. Another mesh network deployed by Waikato University in New Zealand has 17 links with distances from 300 m to 17 km,[1] and has the goal of offering broadband access to remote schools. All the above wireless mesh networks operate in the 2.4 GHz band and use point-to-point links, whereas our metropolitan mesh network uses 802.11a (5 GHz) and contains a point-to-multipoint link. Moreover, an objective in this article is to use the wireless mesh testbed to investigate the interference, rate and power adaptation, and channel assignment for metropolitan wireless links with distances 1.6 to 5 km.

An important goal in the design of a multira-

[1] *http://www.crc.net.nz/crc-net.php*

0163-6804/12/$25.00 © 2012 IEEE

dio wireless mesh network is the efficient utilization of the limited wireless resources and the radio network infrastructure, by controlling key parameters of a wireless communications system such as the assigned channels, and the transmission rate and power. Interference is a key factor that can lead to reduced capacity and performance of wireless networks operating in unlicensed spectrum bands. Indeed, interference can exist between links belonging to the same network, or originate from external sources employing the same or different wireless technologies.

One method of reducing the interference is to appropriately select the channels of wireless interfaces, which also affects the connectivity of wireless mesh networks since two interfaces with omnidirectional antennas that are within transmission range of each other can communicate only if they operate on the same channel. The objective of rate control is to adapt the transmission rate to the channel characteristics in order to improve performance in terms of throughput and packet transmission delay.

Throughput depends on both the transmission rate and the packet loss ratio, and the maximum throughput is not necessarily achieved by the maximum transmission rate or the transmission rate with the lowest packet loss ratio. The power control problem in wireless networks is that of selecting the transmission power at each radio interface in the network, in order to balance energy consumption and performance. We have investigated the operation and performance of the above mechanisms in the metropolitan test-bed, in addition to the time-scales in which they should operate.

A key question we address is whether the aforementioned mechanisms should be performed on a small timescale (on the order of packets or hours) or a much larger timescale (on the order of days or weeks).

In this article, we also present an innovative system that utilizes the low-cost wide-area connectivity provided by the metropolitan wireless mesh network for collecting, processing, and presenting electromagnetic field (EMF) measurements. The rapid growth of wireless technology has brought to the forefront of public interest and concern the issue of increasing EMF radiation, especially from mobile telephony systems; thus, monitoring of EMF radiation is becoming increasingly important, as is the necessity to verify conformance to national and international thresholds. Additionally, the system is able to collect measurements for a wide range of frequencies, and hence can be used for low-cost metropolitan area spectrum sensing in cognitive radio networks. Such a system can create frequency usage maps, which can help cognitive users reduce spectrum sensing by avoiding frequencies that are already used.

The rest of the article is organized as follows. First, we describe the design and deployment of the metropolitan mesh network, and the interference and performance monitoring. Then we discuss experiments on channel assignment, rate and power adaptation, and investigation of the timescales of operation for these mechanisms. Next, we describe a system for real-time collection of EMF measurements and spectrum sensing that utilizes the low-cost wide-area connectivity provided by the metropolitan mesh network. Finally, we conclude the article.



Figure 1. *Topology of the Heraklion metropolitan wireless mesh network.*

Metropolitan Multiradio Mesh Network Design

Topology

The metropolitan mesh network covers an area of approximately 60 km^2 and currently contains 16 nodes, Fig. 1, among which six are core multi-radio mesh nodes (labeled with a leading "K" in the figure). The other ten mesh nodes (labeled with a leading "M" in Fig. 1) are management nodes, that are used to remotely monitor the state and the performance of the network and ensure its normal operation. Management nodes are connected with 802.11a/g links that operate in parallel to the experimental mesh network links. Core multi-radio mesh nodes are connected to the management network through a radio interface that is independent of the interfaces used for the experimental multiradio mesh network connectivity. The above design ensures that we can remotely and continuously monitor the multi-radio mesh nodes. The management network is assigned a static channel, which is different from the channels used in the experimental wireless mesh network, and uses directional antennas that are placed at a distance from the antennas used in the experimental network.

The network contains both point-to-point and point-to-multi-point links, using directional (dish and panel) antennas with gains ranging from 19 to 36 dBi. Each wireless interface is assigned a static IP address, and the Optimized Link State Routing (OLSR)[2] protocol is used for routing traffic in the network, whereas communication of the management nodes use static multihop routing. The mesh network is connected to a fixed network via two nodes (M1 — FORTH and K4 — University of Crete/UoC).

[2] *http://www.olsr.org/*

Together with the independent management and monitoring network, the remote power switch helps to provide continuous remote connectivity to all mesh nodes, enabling remote recovery from crashes and deadlocks, which are very common when experimenting with wireless devices.

Experiment 1 — receiver and transmitter is in same mesh node (K2) and antennas on same mast at distance ≈ 0.75 m		
Channel distance	TX K2→K3	RX K2→K4
0 (36–36)	2.970	2.358
1 (40–36)	2.995	2.976
2 (44–36)	2.997	2.997
Experiment 2 — receiver and transmitter is in same mesh node (K2) and on different mast at distance ≈ 2.5 m		
Channel distance	TX K2→K1	RX K2→K3
0 (36–36)	3	2.75
1 (40–36)	3	3
2 (44–36)	3	3
Experiment 3 — two receivers are in same mesh node (K2), and antennas on same mast at distance ≈ 0.75 m		
Channel distance	RX K2→K4	RX K2→K4
0 (36–36)	2.996	2.996
1 (40–36)	2.996	3
2 (44–36)	3	3

Table 1. *Throughput (Mb/s) measurements of two flows in three different experiments.*

Multiradio Mesh Node

Each multiradio mesh node consists of a mini-ITX board (EPIA SP 13000, 1.3 GHz C3 CPU, 512 Mbytes DDR400 memory) and a 40-Gbyte 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 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 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/g radio interface, which connects the node to the management and monitoring network that operates in parallel to the experimental mesh network. The independent management and monitoring network has proved to be indispensable, allowing us to continuously monitor the performance of the network, and conduct remote experiments with channel, power, and rate control mechanisms without loosing connectivity to the mesh nodes. Additionally, each mesh node contains an intelligent remote power switch (Dataprobe iBoot), which supports on/off power switching through a web interface, and timed power reboots based on the results when the power switch pings other devices, such as the mini-ITX board or some remote device to verify the wireless connectivity. Together with the independent management and monitoring network, the remote power switch helps to provide continuous remote connectivity to all mesh nodes, enabling remote recovery from crashes and deadlocks, which are very common when experimenting with wireless devices.

Interference and Performance Monitoring

Interference Measurements — Most prior measurement studies of outdoor 802.11 links focused on measuring the path loss and the time correlation of losses, and how loss is affected by factors such as received signal strength, link distance, interference, weather conditions, and technology (802.11b/g) [1, 6]. Work on the impact of adjacent channel interference for 802.11b/g is contained in [1, 7] and for 802.11a in [8]; the latter focuses on measuring the impact on the signal-to-noise ratio, and considers link distances of 60 m. The unique feature of our testbed and the results presented in this section is that they consider metropolitan 802.11a links with longer distances (1.6–5 km), and focus on the impact that interference has on the throughput above the MAC layer.

Next we investigate the interference between metropolitan links, when one of the two interfaces of each link under investigation is located in the same mesh node (K2). In particular, we consider the link pair K2–K3 and K2–K4, and the link pair K2–K3 and K2–K1. For the first pair, the two interfaces in node K2 are connected to two 21 dBi panel antennas, which are both on the same mast with a distance of approximately 0.75 m, and have a relative angle of approximately 150°. For the second pair, the two interfaces in node K2 are again connected to two 21 dBi panel antennas, which, however, are on a different mast with a distance of approximately 2.5 m, and have a relative angle of approximately 90°. Each experiment we present below shows the average from 10 runs, each run lasting for 100 s. For all results, the 95 percent confidence interval is less than 6 percent. Finally, the experiments involved UDP traffic with rate 3 Mb/s, generated using the *iperf* tool. Although the maximum link throughput in 802.11a is higher than 3 Mb/s, our goal was to measure the effect of adjacent channel interference, rather than the maximum link throughput. For this reason, we selected a low traffic rate to keep the CPU utilization low and avoid any impact it might have on the results.

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. Two 3 Mb/s UDP streams are transmitted over the links K2 → K3 (2 km distance) and K4 → K2 (1.6 km distance). Note that the two streams are independent, and the iperf sender for the first stream is located in a workstation connected to our internal laboratory network. Table 1 — Experiment 1 shows the throughput

achieved by each UDP flow for three different channel assignments. Observe that when both links are assigned the same channel (channel distance = 0, channels 36 – 36 in Table 1, the transmitter significantly affects the receiver (see throughput for Rx K2 → K4), as they are both located in the same mesh node (K2). 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.

We performed the same experiment as the one described above, but with two 3 Mb/s UDP stream transmitted over the links K2 → K1 (5.1 km distance) and K3 → K2 (2 km distance). Table 1 — Experiment 2 shows that the interference in this case is lower than in the previous experiment. This is due to the larger distance (approximately 2.5 m) between the antennas for node K2 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. Two 3 Mb/s UDP streams are transmitted over links K3 → K2 and K4 → K2. In this experiment the achieved throughput for both flows is identical, and equal to the UDP sending rate (as shown in Table 1 — Experiment 3), even when the links are assigned the same channel. Hence, when the receive interfaces are located in the same node, the interference between the two links is not significant.

The above experiments show that, depending on the distance between antennas, there can be significant interference between metropolitan links, when the transmitting and receiving ends of the two links are located in the same node, and even when the two links are assigned different but adjacent 802.11a channels. The interference can be avoided if the links are assigned channels with a one channel separation, or by placing antennas at some distance, which for our network hardware and configuration was typically more than 2 m.

Continuous Online Performance Monitoring — Continuous monitoring of core mesh network links allows quick detection and identification of anomalous link behavior. Indeed, as we will see in the following sections, for metropolitan links to achieve high performance under normal operation, it is sufficient to fix the transmission rate, power, and channel; the performance should be continuously monitored, and only when anomalies are observed the rate, power, or channel needs to be adapted. 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 min. The collected data is stored in an Round Robin Database (RRD), and the corresponding daily and weekly graphs are made available through an http server[3] using the RRD Tool.

Performance of Rate, Power, and Channel Control over Metropolitan Links

Next we investigate the performance of rate, power, and channel control in the metropolitan mesh network, including the time-scales for the operation of these mechanisms. Note that although the numerical results reported depend on the specific network hardware and configuration of our test-bed, and the degree of external interference, the more general/qualitative conclusions are applicable to any metropolitan-scale network.

Rate Control

We compare the throughput that is achieved using a fixed transmission rate scheme, with the throughput achieved with MadWifi's *SampleRate* algorithm, which we refer to as an auto-rate scheme, that adjusts the transmission rate on a per packet basis. The traffic was generated using the iperf tool. The graphs in this section report values averaged over intervals of 2 min. Figure 2a shows the throughput achieved on a specific link (K2–K4), with the auto-rate scheme (horizontal straight lines) and fixed transmission rate for different transmission powers. Observe that a fixed transmission rate scheme can achieve higher throughput than an auto-rate scheme if the transmission rate is appropriately selected. Moreover, the figure shows that higher improvement is achieved in the case of lower SNR values: when the transmission power is 15 dBm the improvement is approximately 30 percent, whereas when the transmission power is 12 dBm the improvement increases to approximately 42 percent.

Figure 2b shows that the throughput improvements of a fixed transmission rate scheme over an auto-rate scheme remain relatively the same during the course of a day.

The conclusions drawn from the above experimental results are the following:

- For high-quality (high signal-to-noise ratio [SNR]) links, fixing the rate to the highest transmission rate achieves similar throughput as SampleRate.
- For low-quality (low SNR) links, an appropriately selected fixed transmission rate can achieve significantly higher throughput than SampleRate.
- Finally, a fixed transmission rate scheme shows similar performance over long timescales (day). Hence, for a metropolitan link with directional antennas, the transmission rate does not necessarily need to be adapted over smaller timescales.

The above suggests that the adaptation of the transmission rate in small timescales (e.g., on the order of packet arrivals), which is typically followed by all auto-rate algorithms, is not necessary, and can even reduce performance for long distance metropolitan wireless links, which due to attenuation links with low SNR would be more commonplace.

This conclusion highlights the difference between long distance metropolitan links with directional antennas, compared to short distance links where throughput degradation of traditional rate control schemes for 802.11 networks can

The interference can be avoided if the links are assigned channels with a one-channel separation, or by placing antennas at some distance, which for our network hardware and configuration was typically more than 2 m.

[3] *Online access to some of the measurements is available at http://www.ics.forth.gr/HMESH/*

be solved only using sophisticated per-packet rate control mechanisms [11].

Joint Power and Rate Control

Next we investigate power control in metropolitan wireless links. The advantage of reducing the transmission power is the consequent reduction of the interference that is produced. Our objective is to identify the minimum transmission power that achieves close to maximum performance, and investigate how this minimum transmission power varies with time. Figure 3a shows the throughput for different transmission powers for both auto-rate and fixed transmission rates. The key observation from this figure is that the same high throughput is achieved with transmission power higher than 8 dBm with the auto-rate scheme and a 36 Mb/s fixed transmission rate, and higher than 11 dBm for a 48 Mb/s fixed transmission rate. These transmission power values are significantly lower than the maximum transmission power of 15 dBm and achieve the same throughput as seen in the figure. Figure 3b shows, for different times of day, the transmission power required to achieve throughput above 94 percent of the maximum throughput. This figure shows that the necessary minimum transmission power to achieve high throughput does not change significantly throughout the day. The conclusions drawn from the power control experiments for metropolitan links with directional antennas are the following:

- The transmission power can be significantly reduced, without a large impact on the achieved throughput.
- The minimum transmission power to achieve some minimum performance does not significantly change throughout the course of a day; hence, adjustment of the transmission power can occur on longer timescales.

Of course, the above results and the appropriate timescales for power control depend on the level of external interference and the timescales over which this interference changes.

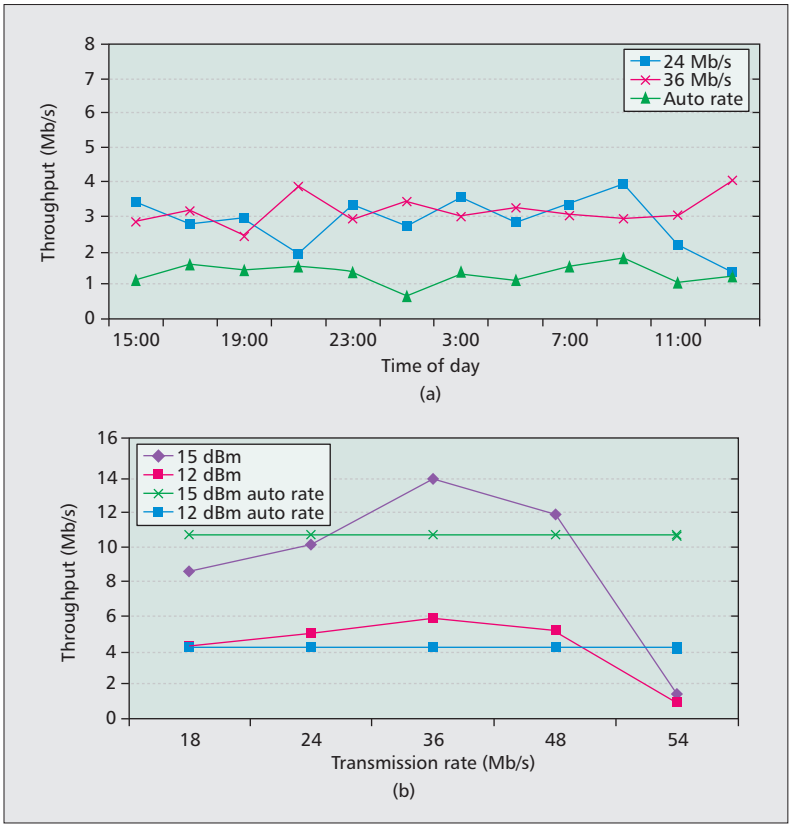


Figure 2. *Rate control measurements: a) throughput for auto-rate and fixed transmission rate; link K2–K4, distance 1.6 km; b) throughput for auto-rate for different periods of the day; link K2–K4, distance 1.6 km.*

Channel Assignment

Next we present experiments related to channel assignment in a metropolitan mesh network with directional antennas. A key requirement for channel assignment, when links operate in an unlicensed band, is to account for both internal (or intra-network) interference between links belonging to the mesh network, and interference from sources external to the network (external interference). In this section we present results of an approach that captures both types of interference, and show the influence on the overall performance when internal or external interference is not taken into account.

One approach for capturing intra-network interference is the Multi-Point Link Conflict Graph (MPLCG) presented in [9]. A vertex in the MPLCG represents a multipoint communication link, which is a set of interfaces that communicate with each other; all interfaces belonging to the same multipoint link should be assigned the same channel. An edge between two vertices in the MPLCG indicates that the two corresponding links interfere, and hence cannot be assigned the same or neighboring channels; the latter is enforced because interference between adjacent channels can exist, even in the case of IEEE 802.11a. In addition to capturing interference, another important component of channel assignment is the actual selection of channels. Possible metrics for selecting a channel are the one-way SNR, two-way SNR (which is taken to be the average SNR at the two interfaces belonging to the same link), or round-trip delay. These metrics can be measured online, and can capture the level of interference from external sources, both 802.11 and non-802.11; other approaches to channel assignment capture only interference between internal links, or external interference solely from 802.11 sources. The two SNR metrics capture adjacent channel interference, but do not capture the contention at the MAC layer between interfaces operating on the same channel. On the other hand, the round-trip delay metric can capture interference due to both adjacent and co-channel interference, since it is influenced by MAC layer contention.

Figure 4a compares the average packet delay achieved with a channel assignment procedure using the MPLCG for capturing intra-network interference and the packet delay metric for channel selection [9], with the Measurement-based Directional Channel Assignment (M-DCA) scheme presented in [10], and channel selection based only on the packet delay metric and channel selection that accounts only for internal interference. The M-DCA scheme relies on an antenna overhearing transmissions from other antennas inside its neighborhood, to identify which channels are already used by its neighbors,

and avoid selecting them. Figure 4a shows that the M-DCA scheme achieves an average delay which is approximately 19 percent higher than average delay achieved by the MPLCG scheme (channel assignment using the multi-link conflict graph for interference modeling, and the delay metric for channel selection). The above result shows that considering only interfering 802.11 networks in the channel assignment process is not enough, and that it is important to also consider the quality of the links, in terms of the packet delay or the SNR; this is especially important in networks with long distance links (typicall, above 1 km as in our metropolitan testbed), compared to networks with smaller distances (as in the network of [10] that contains links up to 60 m).

Figure 4a also shows the average packet delay when only the link quality (in terms of packet delay) is used, without modeling intranetwork interference. Observe that the average delay is 29 percent higher than that achieved when the multilink conflict graph is used, and 10 percent higher than the case when interference from 802.11 networks is considered. This shows that, in addition to the external interference and link quality (captured using the packet delay metric), it is important to consider the intranetwork interference in order to achieve high performance. Finally, the fourth column in Fig. 4a shows the average delay when only intranetwork interference is taken into account. Observe that the average delay is 46 percent higher than when channel selection takes into account both external interference and link quality, in addition to intranetwork interference; hence, many of the schemes appearing in the literature, which focus exclusively on capturing interference among links inside the network, would yield very low performance.

Figure 4b compares, for an interval of 22 days, the average packet delay when the channels are selected once at the beginning of the 22-day period, with an adaptive approach where new channels are selected every day. The MPLCG-based procedure previously proposed in [9] was used in both cases. The depicted delay is the average delay per link for all the links of the mesh network. The results show that the fixed approach achieves an average packet delay within 11.5 percent of the adaptive approach. This suggests that, for the environment where the experiment was performed, there are no significant gains in performing channel assignment on a timescale smaller than 1–2 weeks. This conclusion depends on the fact that in our network the traffic remained constant, and the external interference remained relatively the same, since our network used directional antennas and links with 802.11a, which is not as widespread as 802.11b/g.

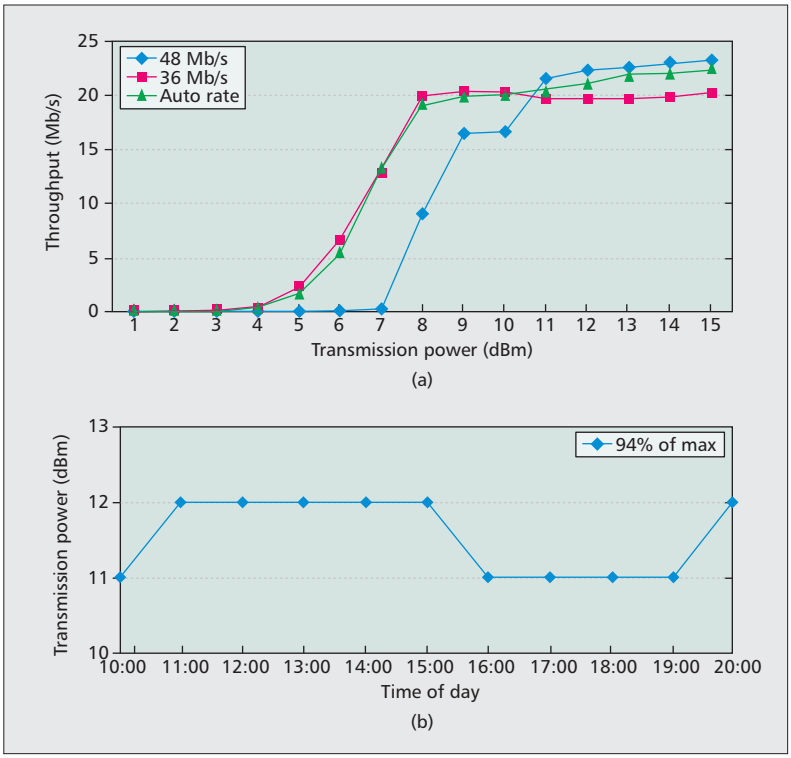


Figure 3. *Joint power and rate control measurements: a) Throughput for different transmission powers. Link K2-K3, distance 2 km; b) Transmission power for achieving throughput above 94 percent of the maximum. Link K2-K3, distance 2 km.*

EMF Monitoring and Spectrum Sensing Using a Metropolitan Mesh Network

Metropolitan mesh networks have many potential uses, such as providing wireless broadband connectivity in a large-scale environment for traffic monitoring, acting as sensors, and so on. In this section we present a novel application we have developed on top of our metropolitan wireless mesh network, which uses mesh nodes for performing measurements of EMF radiation. There are several potential uses for such an application, such as measuring the radiation at buildings when there are multiple antennas nearby, something that is always a major concern of citizens, especially when the measurements are taken near schools or hospitals. Another novel application of EMF monitoring is related to cognitive radio networks (CRNs). In CRNs, secondary users have the ability to access the free spectrum holes via a function called "spectrum sensing," which aims to avoid creating interference to licensed users. A longer spectrum sensing period results in less available time for accessing the spectrum and hence lower network throughput. Using EMF monitoring devices located at the mesh nodes in a metropolitan network, we can perform several measurements of the spectrum utilization at multiple bands, and from these measurements we can extract a map/database of the available spectrum holes that can be used by cognitive users in order to avoid long periods of spectrum sensing. Depending on the time and frequency granularity of the measurements, their collection in some central server can be demanding in terms of both throughput and delay.

The key idea of the monitoring system is to use a low-cost EMF measurement device located in a node connected to the metropolitan mesh network. The EMF monitoring node consists of the following: a low-cost EMF measurement device, a mini PC for controlling the EMF mea-

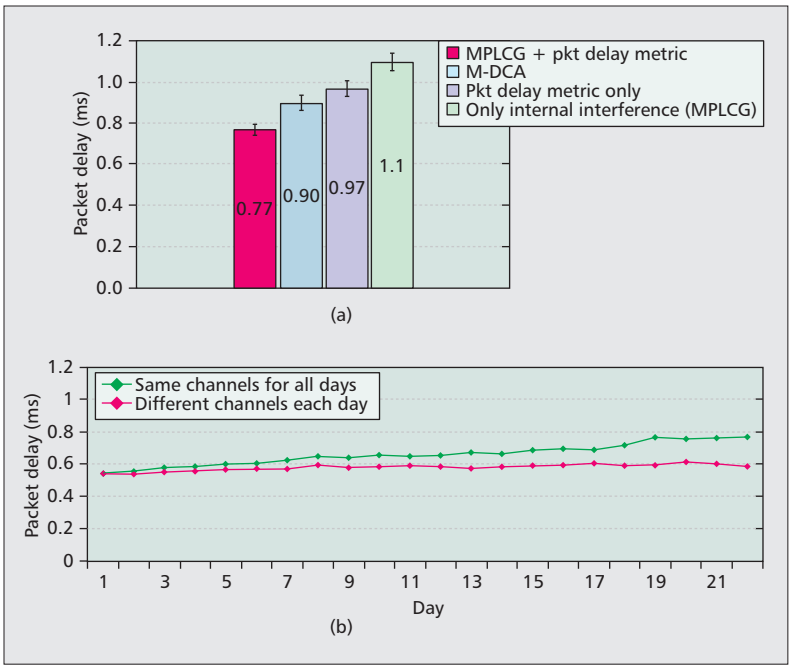


Figure 4. *Channel assignment measurements: a) comparison of channel assignment algorithms; b) comparison of performance when channels are assigned once at the beginning of a 22-day period with performance when channels are adjusted each day.*

surement device, and software modules for collecting, processing, and presenting the measurements.

The EMF measurement device is a low-cost Aaronia Spectran Analyzer. Our current system uses the HF 6060, which has a measurement range of 10–6000 MHz. Additional monitoring nodes we are currently deploying will use the HF 6080, which has a range of 10–7000 MHz. The HF60X0 spectrum analyzer can be connected to two types of antennas: a directional antenna (HyperLOG 6080) and an omnidirectional dipole antenna (BicoLOG 20300). The HyperLOG 6080 directional antenna has a range of 700 MHz–7 GHz. The BicoLOG 20300 omnidirectional dipole antenna has a range of 40–3000 MHz. The BicoLOG antenna is more costly than the HyperLOG antenna, but can measure lower frequencies, which include the FM and TV bands. On the other hand, the lower-cost HyperLOG directional antenna can measure higher frequencies, up to 7 GHz, but requires pointing it toward the area we are interested in measuring. The currently deployed EMF monitoring node includes the omnidirectional dipole antenna BicoLOG 20300.

The EMF monitoring node contains a small PC, which is based on a mini-ITX board (EPIA SP 13000, 1.3 GHz C3, CPU) with 512 Mbytes DDR400 memory, and an 80 GB 2.5" HDD. The mini-ITX PC runs the Aaronia software for controlling and collecting measurements from the HF60X0 analyzer, which also allows external configuration of the spectrum analyzer for parameters that include frequency range, sample time, and resolution bandwidth. An Apache server running on the mini-ITX allows remote access of the measurement graphs through a web interface. There are four Perl scripts executed every 5 min, which take the data from the HF60X0 analyzer and create graphs presenting the EMF measurements in a different manner —per band, per operator, and time series — and can also be executed on demand through a web interface.

Advantages

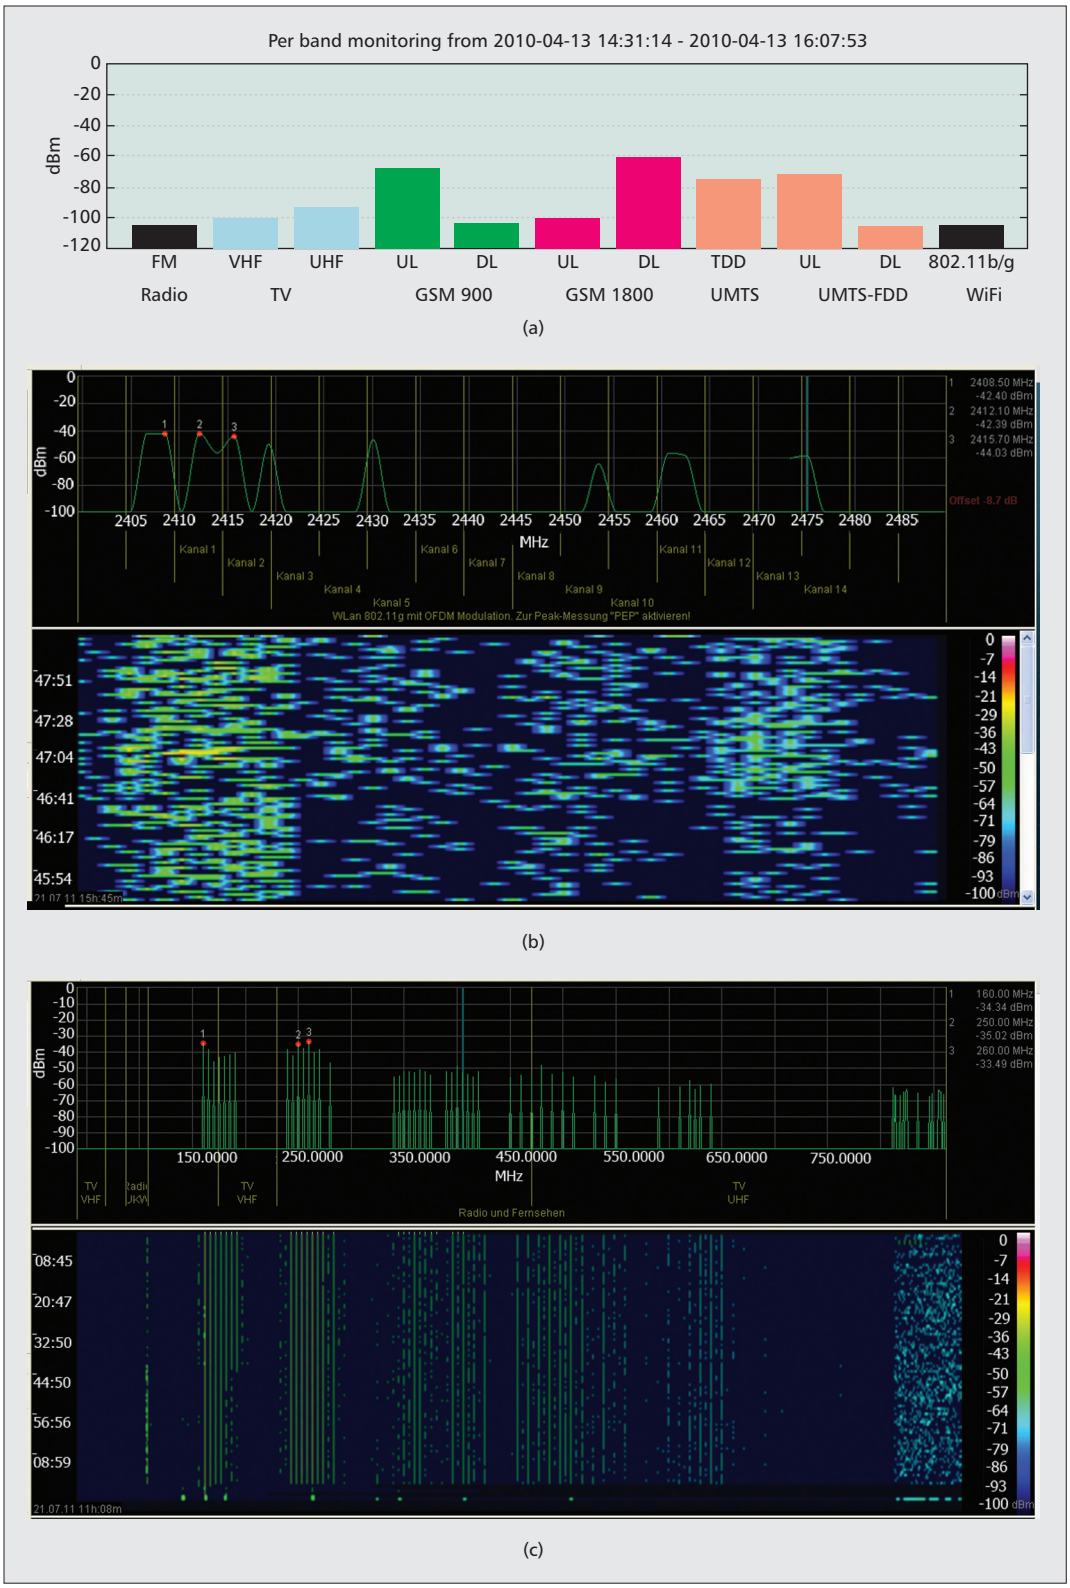
The advantages of using an EMF measurement device connected to a metropolitan mesh network for collecting EMF measurements include the following:

- Higher range: An EMF measurement device can be used to monitor frequencies up to 7 GHz (with the appropriate antenna), which is higher than the capabilities from specialized standalone EMF monitors, whose range is typically limited to 3 GHz.
- Real-time remote measurement collection: EMF measurement devices with real-time monitoring capabilities together with a metropolitan coverage mesh network allow real-time remote collection of EMF measurement data.
- Low cost and high speed: Small (handheld) EMF measurement devices with advanced spectrum analyzer capabilities are significantly cheaper than standalone EMF monitoring devices with remote communication (Global System for Mobile Communications [GSM], third generation [3G], etc.) capabilities, and wireless mesh networks provide higher speeds at lower costs compared to mobile technologies.
- Advanced flexibility: Together with the mini PC, the EMF monitor can be controlled remotely to collect measurements in different frequency ranges (bands) and different time windows.

Monitoring Capabilities

The collected measurement data is in dBm units. In addition to the default units, the measurements can be presented in dBmV, mV/m, mA/m^2, and finally mW/m^2. We note here that electromagnetic radiation limits are usually represented in mV/m or mW/m^2. Figure 5 shows examples of EMF per-band and spectrum sensing measurements.

Per-band and Per-Operator Monitoring — This option allows the presentation of EMF levels in frequencies of various well-known bands. Figure 5a shows an example of per-band monitoring. The displayed values are the average of measurements taken at time intervals of approximately 6 min. The graph is refreshed periodically, based on the last stored values in the database. Each column corresponds to a different band. Note that the values for the GSM and 3G/Universal Mobile Telecommunications System (UMTS) bands are higher than approximately –80 dBm, which is due to the mobile telephony antennas located opposite the EMF monitor node. In addition to per-band monitoring, our system also supports EMF measurements, per-operator frequencies, and time series measurements, which allows the daily, weekly, monthly, and yearly presentation of EMF measurements.



The EMF monitoring node consists of the following: a low-cost EMF measurement device, a mini PC for controlling the EMF measurement device, and software modules for collecting, processing, and presenting the measurements.

Figure 5. *Presentation of EMF and spectrum sensing measurements: a) per band EMF monitoring; b) spectrum sensing in the 2.4 GHz band; c) spectrum sensing in the TV radio band.*

Spectrum Sensing — This option displays, in addition to the aggregate transmission power in different bands, the spectrum holes as a function of time. Specifically, the top of Figs. 5b and 5c shows the received aggregate power in different frequencies in the 2.4 GHz and TV radio bands at a specific time. The bottom part of Figs. 5b and 5c show the received power in the frequencies as a function of time, where time is depicted on the vertical axis. As expected, in the TV radio band only a small portion of the spectrum is being utilized, whereas in the 2.4 GHz band

A more appropriate approach would be to select a static rate, power, and channel, and continuously monitor the network to detect any anomalies or performance drops; once such events are detected, the rate, power, and channel can be adjusted.

almost all the channels of the (unlicensed) spectrum are utilized most of the time. This shows the necessity of using cognitive radios in mesh networks, as they are capable of accessing the free spectrum and achieving much higher performance than operation in the crowded WiFi bands. From these figures, we can observe the utilization of the spectrum in a long time period and extract the spectrum white spaces in each area. This information can be used to create databases of white spaces in order to assist cognitive radio devices in their spectrum sensing function. By installing the EMF devices in the mesh nodes of our metropolitan mesh network, we can extract white space databases for a very large geographical area. The cognitive devices will consult these databases in order to avoid sensing the overcrowded frequencies and focus only on the long-term unused frequencies. This will help minimizing the time dedicated to spectrum sensing and, as a result, maximize the time for accessing the spectrum, which leads to increased network performance.

Conclusions

We present some of our experiences and investigations with an experimental metropolitan multiradio mesh network. In particular, we investigate how interference between metropolitan 802.11a links operating on adjacent channels affects throughput. For rate control, our results show that the adaptation of the transmission rate on very small timescales (on a per packet basis as is commonly the case with widely used auto-rate algorithms) not only does not improve performance, but can also result in significantly lower performance compared to an appropriately selected fixed transmission rate scheme. For channel assignment, we show results demonstrating the importance of taking into account both internal and external interference. Our experiments for the timescales of rate, power, and channel control in metropolitan links with directional antennas show that there are no significant improvements for performing adaptation on short timescales on the order of minutes or hours. Hence, a more appropriate approach would be to select a static rate, power, and channel, and continuously monitor the network to detect any anomalies or performance drops; once such events are detected, the rate, power, and channel can be adjusted. Finally, we present a system for continuous online EMF monitoring and spectrum sensing, which utilizes the low-cost wide-area connectivity provided by the metropolitan mesh network for collecting, processing, and presenting data from low-cost EMF measurement devices.

Acknowledgment

This work was supported in part by the European Commission in the 7th Framework Program through project Enhanced, Ubiquitous, and Dependable Broadband Access using MESH Networks (EU-MESH), ICT-215320, http://www.eu-mesh.eu. The authors would like to thank Vangelis Angelakis, Stefanos Papadakis, Kostas Mathioudakis, and Manolis Delakis for contributing to the design and deployment of the metropolitan wireless mesh network.

References

[1] K. Chebrolu, B. Raman, and S. Sen, "Long-Distance 802.11b Links: Performance Measurements and Experience," *Proc. ACM MOBICOM*, 2006.
[2] R. Patra *et al.*, "WiLDNet: Design and Implementation of High Performance WiFi Based Long Distance Networks," *Proc. USENIX Symp. Networked Sys. Design and Implementation*, 2006.
[3] K. N. Ramachandran *et al.*, "Interference-Aware Channel Assignment in Wulti-Radio Wireless Mesh Networks," *Proc. IEEE INFOCOM*, 2006.
[4] J. Bicket *et al.*, "Architecture and Evaluation of an Unplanned 802.11b Mesh Network," *Proc. ACM MOBICOM*, 2005.
[5] D. Wu, D. Gupta, and P. Mohapatra, "Quail Ridge Reserve Wireless Mesh Network: Experiences, Challenges and Findings," *Proc. IEEE TRIDENTCOM*, 2007.
[6] J. Camp *et al.*," Measurement Driven Deployment of a Two-Tier Urban Mesh Access Network," *Proc. ACM MobiSys*, 2006.
[7] G. Bianchi, F. Formisano, and D. Giustiniano, "802.11b/g Link Level Measurements for an Outdoor Wireless Campus Network," *Proc. IEEE WoWMoM*, 2006.
[8] T. Ireland *et al.*, "The Impact of Directional Antenna Orientation, Spacing, and Channel Separation on Long-distance Multi-hop 802.11g Networks: A Measurement Study," *Proc. 3rd Int'l. Wksp. Wireless Network Measurement*, 2007.
[9] C.-M. Cheng *et al.*, "Adjacent Channel Interference in Dual-Radio 802.11a Nodes and Its Impact on Multi-hop Networking," *Proc. IEEE GLOBECOM*, 2006.
[10] M. Delakis and V. A. Siris, "Channel Assignment in a Metropolitan Wireless Multi-Radio Mesh Network," *Proc. 5th Int'l. Conf. Broadband Commun., Networks and Sys.*, London, Sept. 2008.
[11] S. M. Das *et al.*, "DMesh: Incorporating Practical Directional Antennas in Multichannel Wireless Mesh Networks," *IEEE JSAC*, vol. 24, no. 11, 2006, pp. 2028–39.
[12] E. Ancillotti, R. Bruno, and M. Conti, "Design and Performance Evaluation of Throughput-Aware Rate Adaptation Protocols for IEEE 802.11 Wireless Networks," *Perf. Eval.*, vol. 66, issue 12, Dec. 2009, pp. 811–25.

Biographies

Vasilios A. Siris [M'98] (vsiris@aueb.gr) is an assistant professor at the Department of Informatics of Athens University of Economics and Business, and a research associate at the Institute of Computer Science of FORTH. He received a degree in physics (1990) from the National and Kapodistrian University of Athens, Greece, his M.S. (1992) in computer science from Northeastern University, Boston, Massachusetts, and his Ph.D. (1998) in computer science from the University of Crete, Greece. In Spring 2001 he was a visiting researcher at the Statistical Laboratory of the University of Cambridge, and in Summer 2001 and 2006 he was a research fellow at the research laboratories of British Telecommunications (BT), UK. His research interests include resource management in wired and wireless networks, future Internet architectures, and traffic measurement and analysis.

Elias Z. Tragos (etragos@ics.forth.gr) is a research associate in the Telecommunications and Networks Laboratory of the Institute of Computer Science at FORTH. He received his diploma in electrical and computer engineering (2003), his M.B.A. in technoeconomics (2006), and his Ph.D. (2008) from the School of Electrical and Computer Engineering of the National Technical University of Athens, Greece. His research interests are in the area of mobile and wireless networks, mesh and ad hoc networks, radio resource management, mobility and policy based management, P2P networks, and cognitive networks. He is a member of the Technical Chamber of Greece.

Nikolaos E. Petroulakis [M'09] (npetro@ics.forth.gr) has been a research scientist in the Telecommunications and Networks Laboratory of the Institute of Computer Science at FORTH since 2007. He received his degree in mathematics (2004) from the National and Kapodistrian University of Athens, Greece, and his M.Sc. (2006) in digital communications from the Department of Engineering and Design of the University of Sussex, United Kingdom (2006). He is also a part-time member of the Computer Emergency Response Team of FORTH (FORTHcert) since 2009. His research interests include wireless networks, cognitive radios, and network security.