

A CARMEN mesh experience: deployment and results

Pablo Serrano, Antonio de la Oliva, Carlos J. Bernardos, Ignacio Soto, Albert Banchs
Universidad Carlos III de Madrid
Email: {pablo, aoliva, cjbc, isoto, abanchs}@it.uc3m.es

Arturo Azcorra
Universidad Carlos III de Madrid and IMDEA Networks
Email: azcorra@it.uc3m.es, arturo.azcorra@imdea.org

Abstract

When there is no wired connectivity, wireless mesh networks (WMNs) can provide Internet access with lower cost and greater flexibility than traditional approaches. This has motivated the design of new protocols and algorithms for WMNs, and recently the deployment of experimental prototypes. In this paper we add to these previous works with the performance evaluation of a first CARMEN deployment, with the following distinguishing features: i) it is an indoor deployment, ii) it is used by real users to connect to the Internet, and iii) it is built using off-the-shelf hardware. The results show that mesh technology can provide users with a satisfactory Internet experience, and motives further research along this line.

1 Introduction

Wireless mesh networks represent a very promising technology for providing Internet access to areas that are far from the wired access to the network [1], [2]. In particular, by means of multiple wireless hops, mesh networks allow to bring connectivity to these areas. This represents a major advantage over wired networks since the cost of setting up a wireless infrastructure is much lower than that of deploying new wired links. Additional key advantages of mesh networks are flexibility and speed of deployment. Among other initiatives in the area of mesh networks, the CARMEN project [3] is currently investigating the provisioning of high quality Internet access to users by means of mesh networks.

A considerable effort has been paid until now to the design of algorithms and protocols for mesh networks. In addition to these efforts, experimental testbeds are fundamental to understand the issues behind real deployments. A number of experiments have been performed to the date,

including [4], [5], [6], [7], [8]. The present paper adds to these previous experiments by reporting the experiences of a mesh testbed that was used to offer Internet access during the ACM CoNEXT 2008 conference, which was held at University Carlos III in December 2008. Some of the distinguishing features of the reported experiment are:

- Our experiment was run with real users who were connected to the Internet via mesh and others that used a traditional WLAN connection with an Access Point (AP) connected to the wired network. A questionnaire was distributed among the users to assess the performance of a mesh network as compared to a traditional setting.
- Our deployment was an indoor one and therefore was affected by a number of impairments (e.g., walls) that do not typically affect outdoor deployments. This allowed us to evaluate the impact of these impairments onto mesh networks.
- The experiment was performed with off-the-shelf hardware, which allowed us to analyze the suitability and issues of this type of hardware when use to form a mesh network.

This paper is devoted to reporting of the implementation experiences and the results obtained from the above experience. The key contributions of the paper are:

- We describe the deployment of the experimental testbed and design of the different aspects of the testbed, including the equipment setting, channel assignment and load distribution among others.
- We conduct measurements to understand the limitations of off-the-shelf equipment when used to build mesh networks and propose a network design to tackle such limitations.

- We perform measurements of our running mesh network and assess its performance under realistic traffic conditions with real users. We also report on the quality of our mesh network experienced by these users, as compared to a single-hop one.
- We report our implementation experiences and propose a number of design guidelines resulting from these experiences.

The rest of the paper is organized as follows. Section 2 describes the wireless mesh deployment, Section 3 presents some performance measurements and analyze the obtained results. Finally, Section 4 is devoted to the conclusions and future work.

2 Mesh deployment description

In this section we describe the wireless mesh scenario that we deployed during the ACM CoNEXT 2008 conference¹, as well as the heuristic used to find the frequency planning that maximized the throughput performance of our mesh.

2.1 Experimental setup

The mesh network was deployed to provide Internet access in three different locations of the CoNEXT conference, namely the main conference room, the registration & coffee break area, and the Internet room. The attachment point to the university's wired network that provided Internet access were located in a different building than where the main conference took place.

We first describe the logical topology of our setup (see Figure 1). Two Access Points were set up in the main conference room, configured to operate in IEEE 802.11bg mode in two different non-overlapping channels. A Linux PC configured as a router (GW) was connected to the department's 100 Mbps Ethernet local area network, providing Internet access through the University Internet connection (1 Gbps). Eight Mesh Routers (MeRs) were deployed, configured to provide two different paths between the Internet Gateway and the main conference area. Another Linux router (R) was set up in the main conference room. This router (R) and the Internet Gateway (GW) were configured to perform equal cost multipath routing (using the Linux kernel advanced routing capabilities²), over the two available paths on a per-flow base (i.e. the next-hop used to forward a packet is the same for a given flow, but different next-hops may be used for the routing of packets belonging to different flows).

An MeR are equipped with up to three wireless interfaces. IEEE 802.11a was used for the interfaces of the mesh links, since it provides more non-overlapping channels and a better throughput than 802.11bg. MeR2 had a third 802.11bg interface, operating in access point to also provide Internet access in the Internet room. A different approach was followed to extend the Internet connectivity to the registration & coffee break area. In this case, two Wireless Distribution Points (WDPs) – configured to form a Wireless Distribution System (WDS) – were deployed, because we wanted all the APs of the main conference area (i.e. AP2, AP3 and AP4) to belong to the same layer-2 network. This way, a user terminal can seamlessly handover among all the APs deployed in the main conference area, without requiring to change its IP address, and therefore, without restarting any transport-layer connection.

Public addresses were provided to the users' terminals through DHCP. We had a /24 address block (`PrefMesh.0/24`) available for use in our mesh. This block of IP addresses was divided into three pieces: two /26 blocks (`PrefMesh.0/26` and `PrefMesh.64/26`) and one /25 block (`PrefMesh.128/25`). A DHCP server was installed in MeR2, to serve IP addresses from `PrefMesh.0/26`, and another DHCP server was installed in the router R, configured to serve addresses from `PrefMesh.64/26` and `PrefMesh.128/25`. Private addressing was used for the internal links of the mesh.

We now describe the physical setup, paying attention to the different hardware we used. Since cost is a key factor that determines the feasibility of mesh deployments, we decided to use off-the-self routers to show that even with non-specialized hardware we are able to fulfill the requirements of a real-life scenario. In particular, we chose the Asus WL-500GP router model. This small residential router is equipped with a 266 Mhz processor, an IEEE 802.11bg WLAN interface and an IEEE 802.3 Ethernet interface connected to a VLAN capable 5-port switch. This is a popular and cheap router, that exhibits two interesting and very useful features: its firmware is based on Linux and can be easily modified, and it has a mini-PCI slot that allows to change the original wireless card. We installed a new, open source firmware, *OpenWRT*³ Kamikaze 7.09 distribution with a Linux-2.4 kernel, bringing us more flexibility in the use and configuration of the router than with the original firmware. We also removed the original Broadcom 802.11bg mini-PCI card and inserted instead an Atheros based 802.11bg (Alfa Networks AWPCI085S) one. This card is supported by the Madwifi⁴ driver. Because of the different frequency band used by 802.11a, we also had to change the routers' antennae. We used low gain (8 dBi) external antennae for all the routers, except for links MeR3-MeR4 and MeR5-MeR6, in

¹<http://conferences.sigcomm.org/co-next/2008/>

²<http://lartc.org/>

³<http://www.openwrt.org/>

⁴<http://www.madwifi.org/>

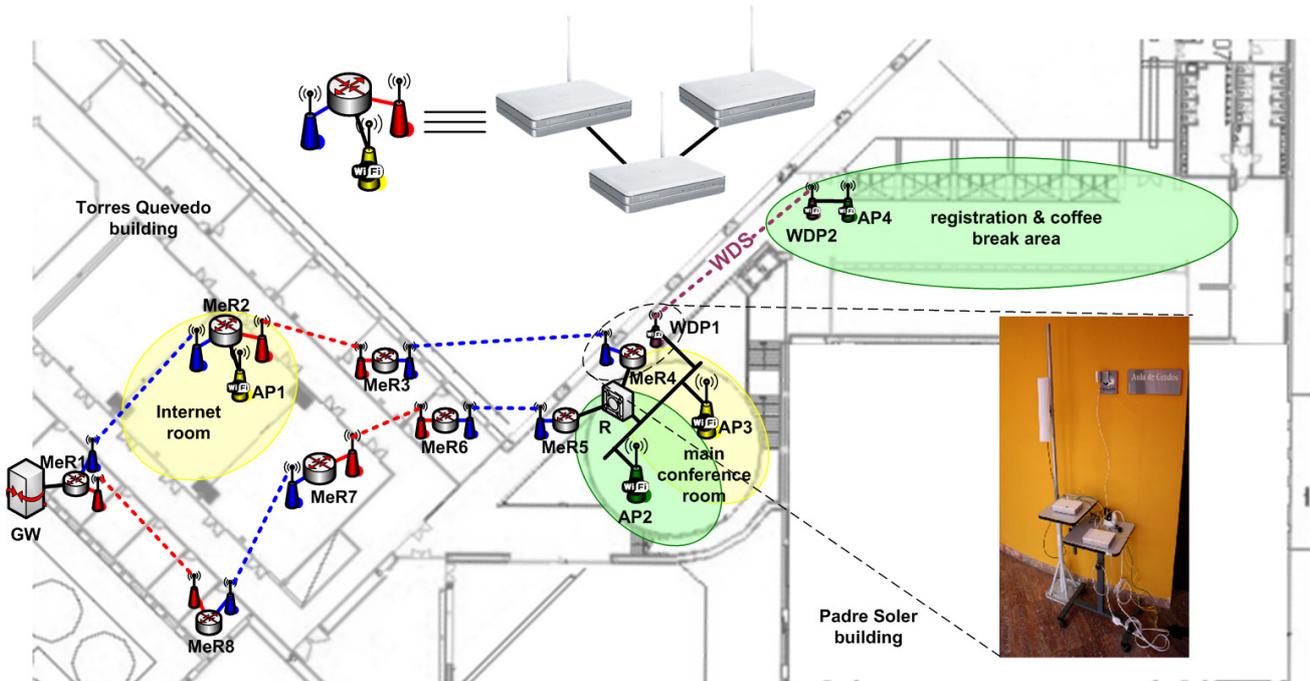


Figure 1. Mesh deployment

which we used high gain (17 dBi) sectoral antennae (see Figure 1), because of the distance and traversed obstacles (mainly walls) of those links. Since the Asus router has only one single wireless interface, our logical MeRs were physically implemented by using one or more Asus routers (as many as MeR wireless interfaces), interconnected through their wired interfaces, as shown in Figure 1.

An additional wired interface of each of the routers was used to perform several control and management plane operations, such as the global synchronization of the local time of all the routers, the remote execution of tests and the retrieval of the results for off-line processing. A central node was used to control and monitor all the routers of our deployment. This prevented the impact of all these operations on the network interfaces being used for the packet forwarding (data plane).

2.2 Frequency Planning

The use of 802.11a provides our deployment with several non-overlapping channels and less interferences from near devices (there were many 802.11bg networks, but just a few 802.11a WLANs). More specifically, our mesh prototype consists of seven wireless links, while the 802.11a devices used support up to seven non-overlapping channels, namely: 36, 44, 52, 60, 149, 152 and 165. However, we still had to run several experiments to maximize throughput performance of the mesh network. This was motivated by

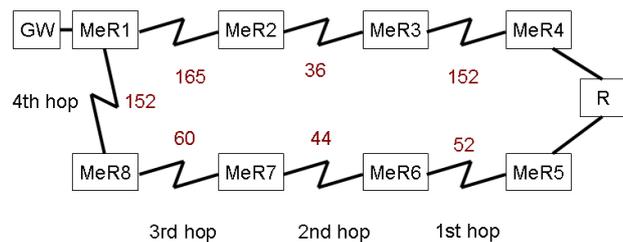


Figure 2. 802.11a channels used

the following results we obtained during a preliminary measurement phase (we omit the description of the experiments for the sake of brevity):

- Frequency has a large impact on performance. By means of UDP-based unidirectional communications, we found that not only some radio channels were asymmetrical, but also that performance could vary significantly when moving from one channel to the next non-overlapping one.
- There is a noticeable interference between non-overlapping channels. Despite we took great care to place devices far from each other⁵, we found wireless

⁵That is, farther than the *far region* field of the antennae $2D^2/\lambda$, D and λ being the maximum overall dimension of the antenna and the wavelength, respectively.

links operating in different frequencies could interfere with each other, this leading to the result that some frequencies pairs were better suited than others when links were relatively close.

Because of the above, performance of the mesh network heavily depends on the channel mapping used (e.g., during the frequency planning phase we had variations of one order of magnitude). The configuration used, illustrated in the simplified deployment of Fig. 2, is obtained through the following heuristic:

- First, we find the pair of frequencies that provides the best performance for the 1st hop. To that aim, we used iperf to set two pairs of TCP flows between the Router (R) and the mesh routers MeR3 and MeR6, respectively, and five 1-minute experiments for each of the fourteen possible configuration of two frequencies of links (MeR3, MeR4) and (MeR6, MeR5). In each experiment we measured the bandwidth r_i obtained by each of the four TCP flows. Out of all the values obtained we choose the configuration that provides the maximum $\min\{r_i\}$, i.e. the best minimum of the four TCP data rates. From this step on, the frequencies of the 1st hop are set to (152, 52), as depicted in the figure.
- Next we perform a similar experiment to set the frequencies of the 2nd hop: we run four TCP flows between R and the routers MeR2 and MeR7, and perform another sweep on the frequencies used in the links (MeR2, MeR3) and (MeR7, MeR6). We take again the configuration that provides the best $\min\{r_i\}$, i.e. the frequency pair (36, 44).
- We proceed similarly for the 3rd and 4th hop, although in the last case the sweep is one only for a single wireless link (MeR1, MeR8), and the four TCP connections are between the same two nodes, GW and R. The final configuration is depicted in Fig. 2. A quite noticeable result is that channel 152 is used for two links, while channel 149 is never used.

The rationale for the above heuristic is the following. Given our physical deployment of Fig. 1, the links more likely to interfere with each other are the “parallel” links of Fig. 2, as these communication may suffer from similar radio conditions (they transverse the same obstacles). Therefore the heuristic proceeds step by step tackling the configuration of the links most likely to interfere, based on the assumption that these are the most critical ones. Note that an exhaustive search on the whole configuration space for 7 frequencies would require a prohibitive time. With the obtained configuration the minimum r_i obtained was approximately 6.5 Mbps, and a total TCP throughput of $\sum r_i = 42$ Mbps.

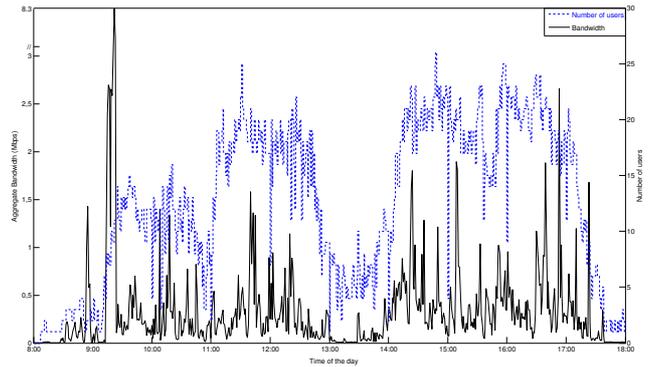


Figure 3. Bandwidth and unique users over time

3 Results/Performance evaluation

All the 195 CoNEXT attendees were provided with instructions to connect to the Internet through the regular Wi-Fi access. Out of these 195, we randomly selected 75 and provided them also with instructions to use the Mesh access. In this paper we present the results for the second day of the conference, with 168 attendees. For this date, the conference started at 9:00 and ended at 17:05, with two coffee breaks and a lunch break distributed as follows: from 10:30 to 10:50 (coffee), from 12:30 to 14:15 (lunch), and from 15:30 to 15:50 (coffee).

In Fig. 3 we plot the average traffic carried over the wireless mesh (black full line) as well as the number of unique MACs served (blue dashed line) every minute. As expected, there is a high correlation between the number of unique devices identified and the traffic generated, as well as with the conference agenda. There are several peaks of 1.5 Mbps, with a quite large and unexpected peak of 8.3 Mbps in the early morning – probably from a curious user assessing the performance of the mesh access. Note that, despite the number of unique MACs per minute varies between 1 and 25, the total number of unique MACs identified is 57. There is always at least one MAC using the mesh, something expected given the registration desk is connected to the Internet through the WDS, and during the lunch break despite there are around 6 devices connected, very little traffic is carried.

We next analyze the traffic balancing mechanism. To that aim, we consider the most loaded traffic hour⁶ and analyze the amount of traffic carried over each wireless path, measured every minute. The result is plot in Fig. 4 for the hour under study in which, according to Fig. 3, there are

⁶We did not consider the peak in the early morning, as we are interested in the mesh performance under regular operations, and that such large peak due to very few users is quite uncommon.

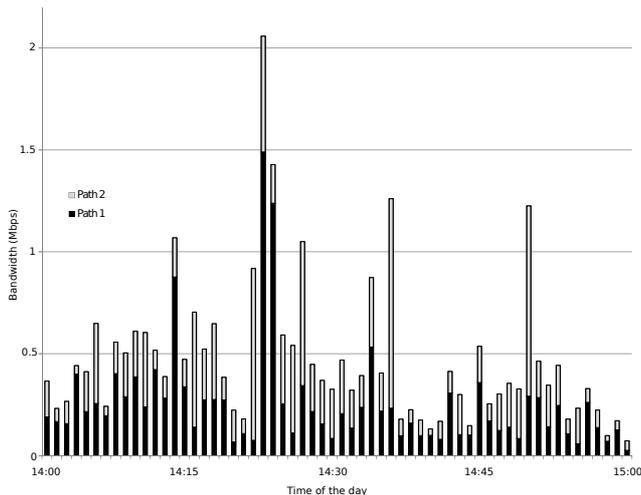


Figure 4. Traffic per radio path conditioned to the number of devices

around 20 different MAC addresses. Given the load balancing mechanism is per-flow based, the results can be summarized as follows:

- Most of the time the flow based scheme achieves a proper balance between paths. Despite these schema do not provide optimal performance, for the case considered both paths carry similar amounts of traffic. This can be attributed to the relatively large number of users and, therefore, flows.
- The asymmetries in path balancing typically happen when there are spikes of traffic. Given that the number of users is relatively constant, these spikes are probably not caused by new user arrivals with new flows, but rather because for a limited time some bandwidth-demanding flow appears.

We next analyze the traffic the type of traffic the mesh is transporting throughout the day. To that aim, we provide in Table 1 the total number of frames captured classified by transport and application layer protocol. For ease of comparison, we also provide the relative weight grouped by protocol layer. Clearly, most of the time the mesh is used to transport TCP segments, as these constitute more than 90% of the total traffic. Out of these, again more than 90% of the packets are due to the HTTP/HTTPS protocol, that includes both web surfing and webmail. With respect to UDP, most of the packets sent are voice over IP traffic, which gives a first idea about the good performance of the mesh network (we present next results from user feedback). It is worth remarking also that the number of IPsec packets is quite noticeable, being larger than the total DNS traffic.

Table 1. Type of traffic

Frames (%)	Protocol	Frames (%)	Protocol
2107983920 (91.3)	TCP	1684943977 (79.9)	HTTP
		264344310 (12.5)	HTTPS
		41994371 (2.0)	SSH
		116701262 (5.6)	Other
46718370 (2.0)	UDP	34434763 (73.7)	VoIP
		5384817 (11.5)	DNS
		6898790 (14.8)	Other
154774466 (6.7)	Other	72089282 (46.6)	IPsec
		82685184 (53.4)	Other

Table 2. User Feedback

How do you grade the performance of the mesh?		
	Responses	%
It was quite good	8	66.7
It was reasonable	2	16.7
It was quite bad	2	16.7
TOTAL	12	100
As compared to the Wi-Fi access?		
Better than the Wi-Fi access	2	18.2
Similar to the Wi-Fi access	6	54.5
Worse than the Wi-Fi access	3	27.3
TOTAL	11	100

Finally, we provide some numerical values for user feedback in Table 2. Not only we wanted to obtain information in terms of the *absolute* performance of the mesh, but also as compared to the regular Wi-Fi access in case this was used. To that aim we asked users through an online polling tool⁷ two different questions, the first about the overall mesh experience, and the second compared to the regular access. Only 12 users filled in the survey which, given that for the date considered we had more than 50 different WLAN devices connected through the mesh, cannot be considered a large number. Still, it is around 20% of all users of that date, and therefore constitutes a reasonable amount of feedback. We consider that the results are quite promising for the deployment of mesh networks: more than two thirds of users qualify the mesh performance as “reasonable” or better, and about 75% of replies rate the mesh access as similar as or better than the Wi-Fi access⁸.

⁷<http://www.surveymonkey.com/>

⁸Note that there is even one user who did not reply to the second question because he did not bother to try the regular Wi-Fi access, as the quality of the mesh experience was good enough for his needs – he/she explicitly explained that through a text box.

4 Conclusions and Future Work

Mesh networks are getting increasing attention from the research community, although deployment experiences are still not so common and mainly are focused in outdoor environments. These outdoor deployments are typically gradual and provide a best-effort service to users with no other means to connect to the Internet. In this paper we have presented the results of a real indoor mesh deployment, used as an alternative Internet access for the ACM CoNEXT 2008 conference. Note that, contrary to previous deployments, in our mesh prototype we wanted to assess if an off-the-shelf mesh prototype can provide users with a similar Internet experience than the one obtained through the Wi-Fi access. Ours is therefore a temporary deployment in parallel to the regular one, that actually “competes” with it, as in case we had a worse performance the mesh would have carried no traffic. However, this was not the case and actually we received quite positive feedback.

Several lessons have been learnt during the experience. First, if we are aiming to offer a good service level, some attention has to be paid to the deployment. Several issues to consider have been identified in the paper (e.g., channel asymmetries, interference between channels), and have been tackled through a simple but efficient heuristic algorithm for frequency assignment, one of the key configuration steps.

Second, we can provide a service with a quality equivalent to a standard Wi-Fi access using a mesh created with relatively inexpensive equipment (commercial off-the-shelf Wi-Fi routers). This was confirmed by surveys with user satisfaction measurements, throughput analysis, and even the protocol analysis that showed a noticeable amount of VoIP traffic had been carried through the mesh.

Third, we can achieve this service level avoiding complex configuration solutions that are not appropriate for a temporal deployment. For example, we used a simple load balancing mechanism readily available in the equipment that made up the mesh, and the results show that, as expected, even if it did not provide perfect load distribution between the available paths, it was reasonable enough to achieve good performance.

In future work we will further analyze the data gathered during the experience. For example we will review layer 2 performance parameters in the different links to better understand the behavior of the mesh and the radio links. Also we will test the mesh under synthetic traffic to test the maximum achievable performance.

Acknowledgment

The authors would like to thank ACM CoNEXT 2008 for allowing us to perform the experiment. We would also like

to thank the users that kindly participated in the experiment and filled the online survey in.

The research leading to these results has received funding from the European Community’s Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 214994. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the CARMEN project or the European Commission.

References

- [1] I. Akyildiz, X. Wang, and W. Wang, “Wireless mesh networks: a survey,” *Computer Networks*, vol. 47, no. 4, pp. 445–487, September 2005.
- [2] S. Faccin, C. Wijting, J. Knecht, and A. Damle, “Mesh WLAN Networks: concept and System Design,” *IEEE Wireless Communications*, vol. 13, no. 2, pp. 10–17, April 2006.
- [3] A. Azcorra, T. Banniza, D. Chieng, J. Fitzpatrick, D. Von-Hugo, M. Natkaniec, S. Robitzsch, and F. Zdarsky, “Supporting carrier grade services over wireless mesh networks - the approach of the european fp-7 strep carmen,” *IEEE Communications Magazine*, accepted.
- [4] “Microsoft self-organizing neighborhood wireless mesh networks.” [Online]. Available: <https://research.microsoft.com/mesh/>
- [5] “Champaign-urbana community wireless network.” [Online]. Available: <http://www.cuwireless.net/>
- [6] “roofnet.” [Online]. Available: <http://pdos.csail.mit.edu/roofnet/doku.php>
- [7] “Meraki.” [Online]. Available: <http://meraki.com/oursolution/mesh/>
- [8] “Open-mesh.” [Online]. Available: <http://open-mesh.com/>