Wirefree in Patmos

M Bletsas

This paper describes the author's involvement in an effort aimed at providing broadband Internet connectivity to the island of Patmos, a small island in the southeast Aegean sea. Different approaches to deploying WiFi are explored for the purpose of providing the convenience of always-on networking to all the physical spaces people frequent, both indoor and outdoor, thus providing cheap ubiquitous broadband Internet connectivity in places that very often have no other way to get that kind of service.

1. Introduction

Patmos has been enjoying first class Internet dial-up access via a direct E1 (2 Mbit/s) circuit to Athens since 1999. A local dial-up ISP (12Net) was founded by Nicholas Negroponte to facilitate access to the Internet.

Greece has been slow in making available consumer broadband-connectivity options. Consumer DSL became available only in late 2003 and only in the large cities. It is safe to assume that places like Patmos will not be enabled for DSL connections for the next few years (at least). In early 2002 there were no consumer broadband options available; we considered a variety of possible scenarios that would allow us to implement a broadband-access network in Patmos.

Rolling out an xDSL infrastructure was one of the seriously considered options: given the small size of the island, local-loop lengths were short enough for DSL; however, the price of the required equipment, the monthly recurring costs for the lease of the copper circuits, as well as the limited number of circuits that we could terminate at the local ISP's premises (collocation of the terminating equipment at the telco's premises was not really an option) quickly ruled out DSL as a viable option.

At the same time, commoditisation of wireless LAN equipment had already resulted in low prices and availability. Couple that with the fact that Patmos is a small island with most of its population concentrated in two towns, Hora and Skala, 2 km apart from each other in direct line and you have a situation that begs for the application of wireless local loop (WLL) technology. But given the scale of the project and the anticipated friction of the bureaucracy, we steered away from WLL solutions that required spectrum licences.

Many companies have been making carrier-class WLL equipment utilising the industrial-scientific-medical (ISM) 2.4 GHz licence-free band. Their equipment is weatherised for installation outdoors, and has all the necessary management functionality to support link monitoring and remote diagnostics, tools that are necessary to maintain reliable connectivity to the network's users. Such equipment has been in use for years with various rural (and urban) ISPs in the United States, where the regulatory framework makes them very attractive.

Another important attribute of the US environment is the presence of the 900 MHz and 5 GHz licence-free bands. The consequence of that is increased flexibility that provides wireless Internet service providers with many tools to work around environmental problems like foliage for example (utilising 900 MHz equipment) or congestion and/or bandwidth issues (utilising the wider and less populated 5 GHz band). The 900 MHz band is only available to the USA since everywhere else it overlaps with the GSM cellular telephony band. The 5 GHz band is currently being cleared for use in most places that allow licence-free operations; however, the details vary considerably from country to country.

In sharp contrast with the FCC, the European Telecommunications Standards Institute (ETSI) recommends that the maximum effective radiated power (ERP) in the 2.4 GHz ISM

band in Europe should not exceed 100 mW (20 dBm) as opposed to the more liberal 1 W (30 dBm)[1] US standard.

Most European governments (including Greece) have adopted these recommendations in their regulatory frameworks. In our case, this precluded the use of FCC-certified equipment, including some interesting ones such as Nokia's Rooftop[2] system. Furthermore, Greek regulations only allow the use of the 2.4 GHz band for licence-free use. The 5 GHz band has not been cleared for such use in Greece yet.

The Rooftop system looked very attractive because of its *ad hoc* mesh architecture, its weatherised one-piece packaging, that did not need any external antenna cables, and the resulting ease of installation of the whole system. Unfortunately the 1 W ERP of transmitted power made the system illegal for deployment in Greece.

When cost considerations were also accounted for, IEEE802.11b systems started to look like the only way to go, given how cheap the radios had become. So the question became: 'Can we build a access network using IEEE802.11 radios for all the necessary[3] links?' The answer was 'yes'. The remainder of this paper provides a detailed overview of the Patmos implementation — what decisions were made and why, what worked, what did not, and a retrospective assessment of the project.

2. Implementation

We begin with some expectations management. With 20 dBm one can cover distances up to 3 km quite comfortably in open space. In making this assumption we use typical receiver-sensitivity characteristics[4] from popular IEEE802.11b wireless cards. For point-to-point links one can play games with separate transmit and receive antennas (placing all the antenna gain on the receive side); however, such tricks increase the cost dramatically and do not work in point-to-multipoint scenarios.

When one compares the ETSI and FCC regulations for the 2.4 GHz band, the conclusion is that ETSI regulations are effectively discouraging the use the licence-free ISM band for outdoor applications. Under the FCC rules distances of up to 40 km are possible using high-gain antennas. That difference in ERP limits makes the comparison with similar community network projects in the USA [2] meaningless, since there is an order-of-magnitude difference in the radius of the network possible with licence-free equipment.

Unless somebody lives really close to an IEEE802.11 access point and there is line of sight through a window or door, indoor connectivity is not possible without some kind of external antenna. 2.4 GHz RF waves do not penetrate walls, foliage or solid structures.

[1] With the FCC regulations, one can actually go beyond 1W ERP by actually lowering the transmitter's power and increasing the antenna gain, allowing for long-range, but spatially very narrow, links.

[2] Currently sold as the WaveWireless SpeedLan 9000 systems.

[3] Client access and backhaul.

[4] They range from –82 dBm to –94 dBm, depending on the signalling speed [1].

most IEEE802.11 implementations employ a DCF for contention-based channel access

In general, an access network consists of access points, where end users connect, and backhaul links, that connect these access points to the upstream network provider. For cost minimisation and ease of deployment it is desirable that the same equipment facilitated both access point and backhaul functionality.

The main determining trade-off here was the cost of the equipment that the end user would have in order to connect to the network versus the reliability of the backhaul links.

Given the desired topology of the network, point-to-multipoint backhaul links were necessary. Implementing them with plain WiFi access points and clients was not feasible. Some digging into the IEEE802.11 medium access control (MAC) layer [3] is necessary in order to explain why.

For most purposes it is sufficient to say that IEEE802.11 wireless local-area networks look a lot like Ethernet networks, and they look almost identical to higher-level protocols. A more in-depth look is very important in our case. IEEE802.11, much like its 802.3 counterpart (Ethernet), assumes a medium shared among all the stations that participate in the network. When that medium is a wire, it is easy to detect when the medium is busy and when it is not, since everybody is connected to the same wire. When the medium is an RF channel though, the assumption that everybody can hear everybody else's transmissions (holding back theirs) is not always safe to make, and the larger the distance between the nodes the greater the chance that they will not be able to hear each other (because of the higher probability that there are going to be obstacles between any two of them).

Given that everybody has to co-operate in order to facilitate access to the radio channel, IEEE802.11 employs a distributed co-ordination function (DCF) for its contention-based channel access. Every node can get access and transmit to the radio channel provided that it has found the channel to be free from other transmissions at that point in time, as long as all nodes use the same rules for determining the state of the channel.

Based on the above, it becomes obvious that the presence of hidden nodes in wireless LANs (nodes that cannot be seen from all other nodes) increases the probability of collisions (concurrent transmissions) and decreases the effective throughput. There are mechanisms in the IEEE802.11 DCF to alleviate the effects of hidden nodes ('request-to-send'/ 'clear-to-send' packets); however, these cannot mitigate the effects of large numbers of hidden nodes.

Networks have physical dimensions, and signals do need time to propagate through them. Propagation delay forces the existence of interframe spaces in contention-based MAC

schemes. Each node, before assuming that the channel is free, has to find the channel quiet for the period of time it takes for the signal to propagate from a node at the other end of the network's diameter. These intervals are defined in the MAC specification and have two effects:

- they essentially define the diameter of the network (since placing a node further away than the interframe space mandates, will break that node's clear channel assessment functionality by delaying the transmissions from other nodes reaching it),
- they force idle periods on the network, lowering the effective bit rate that the network can supply.

The designers of IEEE802.11, being aware of DCF's operational limitations, also defined a point co-ordination function (PCF) [4] in which a central node, visible to all the rest, referees access to the medium, by polling each of the nodes consecutively[5]. This resembles token-based medium access schemes with the central station distributing the token.

The PCF has not been widely implemented; however, it was absolutely necessary for our wireless backbone, since nodes were going to be up to 2 km apart from each other and generally visible only to a central one and not to each other.

3. Deployment

Our initial deployment in Patmos was scheduled for Spring 2002. In the beginning of that year, therefore, we set out to find equipment with point-to-multipoint bridging as well as plain IEEE802.11 access point functionality.

After testing a variety of equipment that offered point-to-multipoint functionality, we ended up buying used dual-radio Orinoco access points from eBay that could support PCF-like functionality via third-party add-on software, thus lowering our capital expenditures dramatically.

External antennas were connected to the radios using LMR-400 low-loss coaxial cable, lightning arrestors and short 'pigtail' cables that had the miniature RF connectors necessary to attach an external antenna to the PCMCIA radios.

With the topology shown on the map (see Fig 1), a client computer can connect to the network in any of the following ways.

- Direct connection to one of the IEEE802.11b public access points (red dots on the map)

  Line of sight to the antenna is required as well as a small portable external antenna (like the Orinoco range-extender antenna). To get connectivity inside a house the use of an external antenna is required. No antenna is needed in the vicinity (~200 m radius) of the access points.

- Direct connection to one of the backbone (PCF) base-stations (green circles on the map)

  The hardware is exactly the same as in the previous case (a modified wireless-card driver is required). However, one can expect more reliable service this way, since hidden-node issues are eliminated.

- Connection via a bridge/router (satellite) configured as a PCF client

  This is the suggested solution if multiple computers have to be connected using one radio, or in the case where an IEEE802.11b public access point needs to be installed and linked to the backbone.

As far as routing to the public Internet was concerned, we had to hide all the wireless nodes behind a network address translation (NAT) router Linux box, which also implemented access control using the popular open source NoCat captive portal[6]. That decision was mandated by the small IP address space allocated to 12Net. The same box runs all of the (open source) SNMP monitoring software (MRTG and RRD by Tobias Oetiker) for the wireless links, as well as VPN software (PopTop) used for remote monitoring and configuration.

The total cost of equipment and materials for our initial deployment was $12K and the physical installation took two people just one week. At this point, two years after our initial installation and with the deployment shown on the map (Fig 1), the total cost has been around $25K and today a connection is feasible from 90% of the island's buildings, as well as from 75% of its land area. Even more interesting is the fact that connectivity is feasible from the sea surrounding the island, something that opens the possibility of very interesting applications. For example, we are working with an environmentalist group to provide live continuous monitoring of endangered species habitat in the area.

4. Lessons learned

The biggest impediment to our network's growth was the difficulty of installation. Since our electronics were not weatherised, they had to be installed either indoors or in a weatherproof outdoor enclosure and then connected to the antennas via pigtails, lightning arrestors and bulky low-loss LMR-400 cable (Fig 2). To make things even more complicated, a good ground connection had to be found for the lightning arrestors.

the biggest impediment to our network's growth was the difficulty of installation

Although these problems have been solved by many equipment manufacturers by placing all of the electronics and the antenna in the same weatherised enclosure, we could not find ETSI-compliant equipment that would operate in the

[5] One has to note here that such polling schemes have been around since the very early days of wireless networking. AlohaNet had probably the first implementation and WiMax is using a similar method in its MAC layer.

[6] A Captive Portal is an access control system where the user, upon connecting to the network, gets a log-in page on his/her Web browser the first time that a Web page access is attempted.

Fig 1 Current backbone topology (June 2004 — no client links are depicted).

licence-free 2.4 GHz band when we were planning our initial deployment. The situation has improved somewhat two years later; however, weatherised equipment that is compliant with the Greek regulations is still very hard to get.

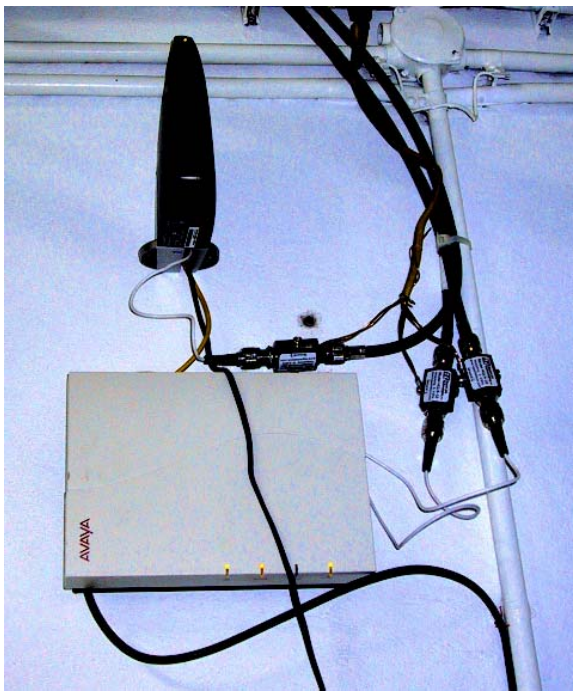

Fig 2 3-radio wireless node (two radios are used for backhaul connections and one for local WiFi access).

Equally important to the ease of deployment is reliability. In the Patmos case, the network has to run for months at a time without any physical maintenance beyond remote diagnostics. At this point the network has been up for two years and has survived two winters. We had to replace three antennas that were not waterproof, and two power supplies (the island's diesel-powered generators do not produce very high-quality power). We also had to power-cycle one of the backbone nodes on two occasions. The problem was alleviated by new firmware on the PCMCIA radio cards and the access points.

A factor that we have working for us is the relative low noise level in the ISM band in Patmos, which allows us to operate wireless links with relatively low (~5 dB) system margins (the difference between the received signal level and the receiver's sensitivity threshold[7]) with minimum packet loss. The relative low density of vegetation on the island is also helpful.

To further demonstrate this point, we did try an 18 km link to the neighbouring island of Arki using the same 32 mW radios on both ends, and using a 22 dBi grid antenna on the Arki end of the link, pointing to a 6 dBi omni-directional antenna on Patmos (Fig 3).

The Orinoco access points provide SNMP monitoring support, as well as very good remote configuration and monitoring tools. That support has proven to be indispensable in the remote troubleshooting of the network, which has to happen from Cambridge, Massachusetts in most cases. We use MRTG

[7] The minimum input signal level needed to receive.

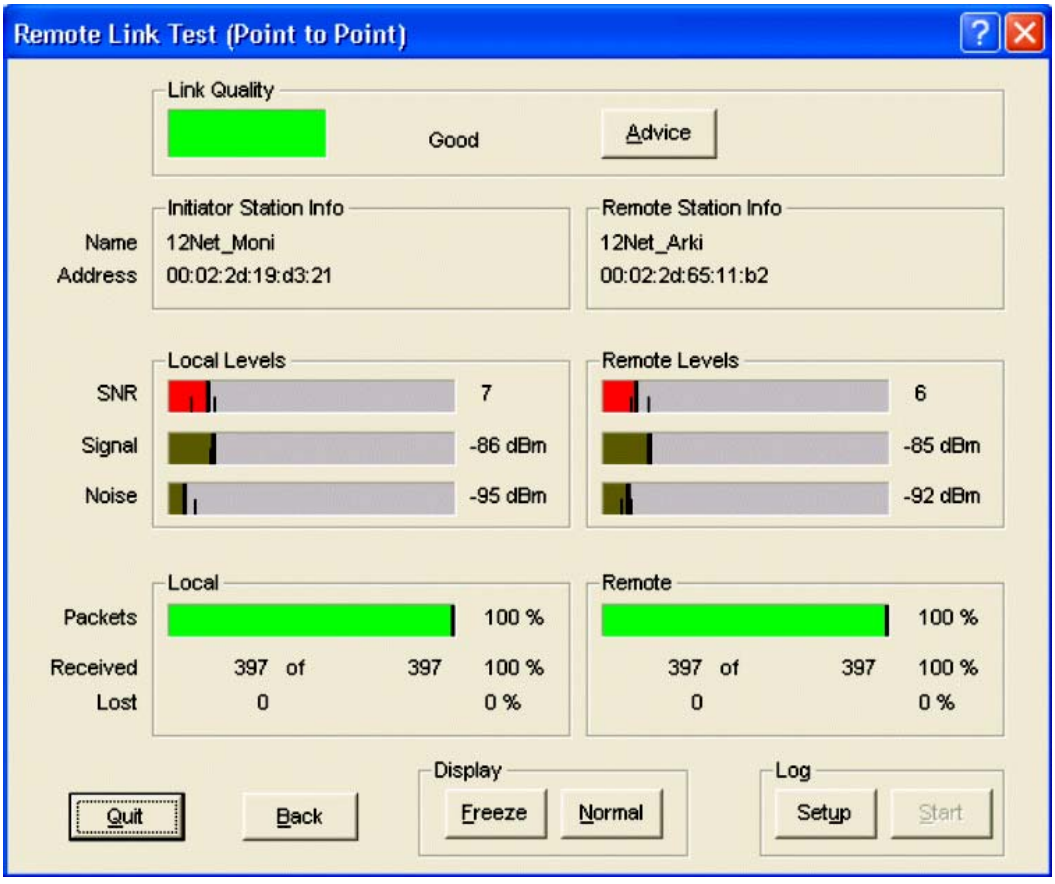

Fig 3 Link test measurements between Patmos and Arki.

to graph traffic and the RRD tool to graph signal and noise measurements of our wireless links and the graphs have been extremely helpful exposing trouble spots in the network.

The use of polled access on the backbone links has given us predictability in the backbone's performance. Such claims cannot be made for the IEEE802.11b end-user connections, since there users that are closer to the antenna end up experiencing better performance (at the expense of their more distant peers). We try to minimise the effects of this problem by adding public IEEE802.11b access points closer to the users (that is why there are three access points close to the network's base) and by connecting distant users using the PCF options.

Given the static configuration of the point-to-multipoint radio links, implementing redundant paths in the backbone cloud becomes costly since it requires extra hardware in hot-spare mode. As the number of users grows, such an expense can be justified and we plan to add one such link during 2004.

An alternative solution would be to use a mesh (or *ad hoc*) networking architecture. In such networks, nodes discover their neighbours automatically, maintain current link status information and select dynamically the best possible path for data in transit. Adding more nodes increases both the aggregate throughput of the network as well as the overall reliability.

5. Mesh weaving

During the past decade there was substantial research investment in the area of *ad hoc* routing protocols [5]. There is an expanding base of open source software components that can be used to put together mesh-enabled nodes.

That is the approach taken by the Locustworld [6] project, which uses a small x86 node running a Linux kernel coupled with a free implementation of the AODV *ad hoc* routing protocol and one IEEE802.11 radio interface. The Locustworld node uses that radio to communicate with all the other nodes that it can reach via its omnidirectional antenna. Client nodes (PCs, PDAs, etc) can be connected either via the node's Ethernet interface or via the IEEE802.11 interface.

A similar approach has been taken by MIT's Rooftop [7] project. The main difference in the MIT approach lies in the use of the far more sophisticated Click modular router architecture running a variation of the DSR routing algorithm. Roofnet's routing protocol (SrcRR) tries to address the challenges imposed by the wireless environment. Those include intermediate quality of most links, asymmetric link-loss rates, frequent changes in link-loss rates, and frequent losses of routing protocol packets due to interference from hidden terminals.

Given the almost identical hardware used by the two projects, they both tend to suffer from similar physical installation constraints, i.e. routing the LMR-400 cable from the antenna outside to the node's electronics box inside as well as finding a good grounding path for the lightning arrestor.

Far more important though are the limitations imposed by the single IEEE802.11 radio and its omni-directional antenna. All participating nodes have to use the same frequency channel to communicate, which severely impedes the system's throughput as well as its routing stability.

Nortel's Mesh Networks [8] product tries to solve both of these problems in a very elegant way. Nortel's node utilises an IEEE802.11a (5 GHz) radio for the backhaul (transit) links between the nodes and an IEEE802.11b/g radio for client (end-user node) access. The transit link radio is connected to a switched multi-beam antenna that automatically selects the best beam for directional communication. The transit link radio also selects a different channel for different destination nodes reducing this way the possibility of collisions and increasing overall system throughput.

All the hardware is enclosed in a weatherproof radome pod for quick and easy installation in both outdoor and indoor spaces. Furthermore, traffic on the transit links is encrypted transparently on a per-client basis facilitating communications privacy.

6. Conclusions

Two years have passed from our first installation in Patmos, and our network has grown beyond our initial expectations. Although there have been many technological developments since the initial design and choice of equipment was made in 2002, the regulatory framework has not changed and that (unfortunately) makes our choices still valid in 2004. Despite the strait-jacket imposed by that framework, we have shown that providing broadband connectivity to underserved areas can be done with minimal capital expenses using affordable off-the-self technology.

The social implications of this experiment will take some time to manifest; however, the author found extremely gratifying the bragging of some Patmos users to their friends in Athens regarding their Internet access speeds. For a place that very often in the winter becomes physically inaccessible (due to bad weather), it is very important for its people to feel tightly connected to the rest of the world.

WiFi still stands to reap a big boost in utility from the deployment of *ad hoc* routing protocols

Returning to the technical aspects, the buzz these days for wirelessly backhauling traffic revolves around another recently ratified IEEE802-series standard, namely 802.16 or WiMax as it is known in the trade press.

Specified with a variety of operating frequencies in mind (occupying both licensed and unlicensed spectrum), it tries to be a solution for both end users that deploy their own infrastructure as well as for established carriers. That might prove to be the biggest impediment to its popularity, since its

many variants might become an impediment to the economies of scale that made WiFi so cheap.

WiFi still stands to reap a big boost in utility from the deployment of *ad hoc* routing protocols in the operating systems of the popular mobile computing platforms. Those protocols, coupled with the extended addressing and autoconfiguration capabilities of IPv6 (the next generation of the Internet protocol) will make possible infrastructure-less communications networks that will be formed by and for their users.

Despite its popularity and ubiquity, we should not forget that WiFi is yet another frame transport technology (and as such it is difficult to solely base a business plan upon it). With its ubiquity comes low cost and accessibility, so it is quite safe to predict that it will eventually become as popular as Ethernet. Never quite a meal in itself, but always a must-have ingredient!

Acknowledgements

I would like to thank Nicholas Negroponte and Walter Bender for making it possible for me to pursue such a fun project, the Patmos crew (Dimitri Negroponte, Elias Kamaratos, Michalis Anastasiadis) for their help with the installation of the equipment, and the anonymous reviewers and Walter Bender for helping me improve this paper.

References

1 Proxim Networks, Orinoco Gold Classic Card Datasheet — http://www.proxim.com/learn/library/datasheets/gold_pccard.pdf

2 Barranca M: 'Unlicensed wireless broadband profiles: community, municipal and commercial success stories', NewAmerica Foundation (April 2004) — http://www.newamerica.net/Download_Docs/pdfs/Pub_File_1547_1.pdf

3 Gast M: '802.11 wireless networks, the definite guide', O'Reilly and Associates (April 2002).

4 Kopsel A, Ebert J-P and Wolisz A: 'A performance comparison of point and distributed coordination function of an 802.11 WLAN in the presence of real-time requirements', Proceedings of the 7th Intl Workshop on Mobile Multimedia Communications (MoMuC2000) (October 2000).

5 Perkins C: 'Ad-hoc networking', Addison-Wesley (January 2001).

6 LocustWorld — http://www.locustworld.net

7 Aguayo D, Bicket J, Biswas S, De Couto D S J and Morris R: 'MIT Roofnet implementation', (August 2003) — http://www.pdos.lcs.mit.edu/roofnet/design

8 Nortel Networks: Wireless Mesh Network Solution — http://www.nortelnetworks.com/solutions/wrlsmesh/

Michail Bletsas, a research scientist and director of computing at the MIT Media Lab, designed and deployed most of the Internet network infrastructure systems at the Lab.

Currently, he is experimenting with wireless networks that are implemented using off-the-self, low-cost components to provide broadband Internet access to underserved areas.

Before joining the Media Lab, he was a systems engineer at Aware Inc, where he designed and wrote high-performance software libraries for Intel's distributed-memory parallel supercomputers, and was involved in the development of one of the first ADSL Internet-access test beds.

He holds a Diploma in electrical engineering from the Aristotle University of Thessaloniki and an MSc degree in computer engineering from Boston University.