

**Finding A Reasonable Compromise for the
Use of 802.11 Equipment in Australia –**
How to solve the last mile problem without upsetting anyone too much

Greg Baker BSc.
The Institute for Open Systems Technologies Pty Ltd.

May 22, 2002

Abstract

This proposal describes a small extension to the existing spectrum-usage licensing. This extension is simple to implement and to police.

It will help uptake of wireless broadband in Australia, particularly in rural and regional Australia.

It preserves the investments made by existing players and also maximises the efficient use of available bandwidth.

This proposal introduces the concept of the “Mesh” – an open access network that may be used by commercial entities as long as certain conditions are met. These conditions are designed to promote co-operation between commercial entities to solve last-mile problems in Australia.

Contents

| | | |
|----------|--|-----------|
| 1 | What's wrong now? | 4 |
| 1.1 | Crime and the 2.4GHz spectrum | 5 |
| 1.2 | What are some bad solutions, and why are they bad? | 5 |
| 1.2.1 | Remove the need for carrier licenses | 5 |
| 1.2.2 | Keep the requirement for carrier licenses | 6 |
| 1.3 | What are the issues? | 7 |
| 1.4 | The Mesh compromise | 8 |
| 1.4.1 | Why is the Mesh good for carriers? | 9 |
| 1.4.2 | Why is the Mesh good for non-carriers? | 9 |
| 1.4.3 | Why is this good for the Australian consumer? | 10 |
| 2 | The solution | 11 |
| 3 | What the amendment does | 12 |
| 4 | Answering objections to the Mesh | 17 |
| 4.1 | Will the Mesh get congested? | 17 |
| 4.1.1 | Intra-cell congestion | 18 |

| | | |
|----------|---|-----------|
| 4.1.2 | Inter-cell congestion | 20 |
| 4.2 | Is the Mesh scalable? Is it reliable? | 22 |
| 4.3 | How does the Mesh protect itself? | 23 |
| 4.4 | Will people be concerned about base stations? | 24 |
| 4.5 | What is the economic model behind the Mesh? | 24 |
| 4.6 | Who will invest in Mesh points? | 25 |
| 5 | Conclusions | 27 |

List of Figures

| | | |
|-----|---|----|
| 3.1 | An ISP can use wireless easily to reach customers... | 13 |
| 3.2 | A second ISP sets up an overlapping cell, and both ISPs benefit | 14 |
| 3.3 | A community group puts an antenna on someone's roof | 15 |
| 3.4 | The Mesh has grown like topsy and it is a lot of retransmissions to get back to town now for some... | 16 |
| 3.5 | But the consumers can spontaneously fix problems themselves if they really want to | 16 |

Chapter 1

What's wrong now?

Several companies are flagrantly breaking the law, big business is unwilling to invest and well-meaning volunteer efforts to help are being thwarted by unfortunate combinations of legislation. Welcome to the world of 802.11 wide-area wireless access in Australia.

In this chapter I discuss what problems exist (section 1.1) and some things that we could try that are a really bad idea (section 1.2). Then section 1.3 summarises the mire of conflicting requirements and proposes some criteria to help us make a decision and section 1.4 outlines a compromise that keeps everyone at least moderately happy.

In the next chapter – chapter 2 – I will detail a tiny amendment to the Telecommunications Act 1997 to define the compromise. If you are just interested in the legal details, you might want to skip there.

Chapter 3 tries to give a picture of life after the amendment. If you aren't interested in the details or the reasons behind it all, chapter 3 gives a good overview.

Finally chapter 4 answers some of the major objections to what I am proposing. Since this is filled with formulas, technical discussions and other such eye-watering stuff, this would be of interest to engineers and members of the telecommunications industry.

1.1 Crime and the 2.4GHz spectrum

As of this writing, the general legal opinion on using 802.11 antennas in the 2.4GHz range for a large-scale network is:

- That you do not need a carrier license to set up and use such a network if you do not offer the service to the general public.
- That you do not need a carrier license if you are a not-for-profit entity and you do not provide access to the internet.
- That you do need an (expensive) carrier license if you want to do anything else.

Unfortunately, there are several organisations which cannot reasonably fit into any of these categories. The vast majority of ISPs (internet service providers) do offer services to the general public, are for-profit operations and cannot afford a carrier license, but would dearly love to be able to offer 802.11 services to their community. Small regional ISPs are particularly keen to see expanded options for 802.11 as they have no other two-way broadband options that they can offer their customers. And 802.11 is cheap and easy enough for them to deploy that if they can do so legally, they will.

The pressure is so high to run 802.11 networks that some ISPs are planning on doing it illegally anyway. I have been approached by two such companies in the last three months to help them do exactly that¹.

There is clearly a need for more open access to the spectrum; what can be done? Section 1.2 suggests what we should *not* do.

1.2 What are some bad solutions, and why are they bad?

1.2.1 Remove the need for carrier licenses

Australia could conceivably opt for the “anarchy solution”. Since we already have a regulatory framework that allows 802.11 equipment to be used in a

¹I declined, if you are interested.

variety of ways, we could open this up to allow anyone to use it, even to make a profit. There would be no carrier licenses for its use (or very cheap ones) and anyone wanting to build a wide area network would end up using it.

Unfortunately, there are some big problems with this:

- With many overlapping networks trying to use the same bandwidth, interference is going to be a constant problem. How will we declare that a particular region has enough transmitters in it, and stop anyone from placing any more? Whose responsibility is it to fix these problems?
- It destroys the investment of those few ISPs who have paid for a carrier license. It also defeats one of the main purposes of the carrier license, which is to ensure fair uptake in rural areas.

1.2.2 Keep the requirement for carrier licenses

Alternatively, Australia could opt to maintain the status quo.

If carrier licenses are kept expensive and hard to acquire, then it is likely that it would only ever be economically worthwhile to deploy base stations in highly built-up commercial zones and in rich residential areas. This is utterly pointless, since we already have both cable and ADSL capabilities to these regions, and all we would be adding is slightly better (but nowhere near universal) mobility.

Moreover, since most of the existing carriers have huge investments in 3G mobile, they are not going to touch anything that might lessen their revenue stream. Splitting the bets over two wireless technologies just isn't a good idea for them.

And this is what we have seen so far – there is an almost total lack of investment into 802.11 infrastructure in Australia. The level of 802.11 roll-out by the carriers has been dwarfed by the volunteer efforts of community wireless groups².

And none of this solves the problem of section 1.1 – there are companies who are attempting to run networks without a carrier license anyway, regardless

²e.g. The Meshes of Sydney, Melbourne, Brisbane, Bendigo, etc.

of whether it is legal or not.

1.3 What are the issues?

Notice that the two main options (“anarchy” in section 1.2.1 and “status quo” in section 1.2.2) both fail to encourage deployment of 802.11 in regional and rural areas. And that is particularly tragic, because it will waste a golden opportunity to bring these areas up to near city-levels for broadband internet access. 802.11 can have “cell” sizes of anything up to 30km radius (900km² – though 50-60km² is probably more realistic) using cheap off-the-shelf equipment.

Note also, that both options mean that some existing business will lose their investments in infrastructure and licenses. “Status quo” makes criminals of regional ISPs; “anarchy” wastes the investments of the carriers.

“Status quo” under-uses the spectrum (thus wasting a completely renewable resource). On the other hand, “anarchy” over-uses it (thus wasting a completely renewable resource).

Finally, a regulatory regime that people openly talk about flouting (“status quo”) is clearly inadequate, and a regulatory regime that produces chaos (“anarchy”) is possibly more so.

So I propose that there are four criteria by which we should measure any proposed change we might wish to make.

1. Does this proposal provide a business model to make investment in rural and regional Australian access possible? (i.e. is this just going to benefit the big cities?)
2. We have a business climate which requires simultaneous closed and open access to the spectrum. Does this proposal preserve the investments of all parties, and give them options for future business?
3. We have to strike a balance between over-using and under-using the spectrum. Does the proposal waste the spectrum?
4. Can this proposal be enforced easily, and are people likely to want to abide by it? Does the proposal facilitate the legal structures to promote

the above three points?

1.4 The Mesh compromise

Section 1.3 outlined the difficulties of the two main options (section 1.2) available.

I propose that we can legislate the existence of the “Mesh”. The actual legislative changes are outlined in chapter 2. The legal framework for this is quite straightforward and only requires a small amendment to the Telecommunications Act 1997.

The Mesh consists of open access wireless networks. It is deliberately designed to be:

- always available as a public resource,
- not requiring a carrier license to use,
- restricted so that it is very hard to monopolise.

The rules of the Mesh are:

1. You cannot charge anyone for transferring data over your part of the Mesh.
2. You are allowed to charge for electronic or data services which happened to be transferred across the Mesh. For example, you could run a mail server on behalf of a client and charge them for running the mail server, even though they only access it across the Mesh. You would not, however be able to charge them for the cost of transferring the data to them, since that is disallowed by the previous rule.
3. You must give access to everyone (including your competitors) to your part of the Mesh. You can prioritise traffic to give your customers faster response than your competitors, but you cannot prevent them from accessing your portion entirely.
4. Your competitors must give you access to their part of the Mesh (which is just rule 2 in reverse).

5. You are not forced to join the Mesh – if you have a carrier license, you are still free to build your own wireless networks just as before, with any access and charging scheme you choose. Carriers can completely join the Mesh if they want to. They could do a combination of both, and of course, they can provide services just like any non-carrier making use of the Mesh.
6. You still have to observe the relevant regulations with regard to radio-communications transmission levels and also any telecommunications legislation.
7. You do not need a carrier license if your wireless network is part of the Mesh.

The remaining sections of this chapter analyse this proposal in terms of the criteria of section 1.3 on page 7.

1.4.1 Why is the Mesh good for carriers?

The item numbers here refer to the item numbers from section 1.3.

1. Carriers can make use of infrastructure provided by others in rural areas, or deploy their own. They are able to leverage deployments in areas that they would consider unprofitable.
2. Carriers keep their full access to the spectrum.
3. There is no danger of under-utilised spectrum but at the same time they are not likely to be badly affected by cross-noise in long haul links.
4. The legal status of carriers continues to be clear, since it is unchanged.

1.4.2 Why is the Mesh good for non-carriers?

(By non-carriers, I am typically referring to ISPs, although other organisations are likely to get involved.)

1. Non-carriers can also make use of existing infrastructure, as they are not reliant on carriers or their budgets for deploying in rural and regional Australia.
2. They get access to the same spectrum as the carriers on a level playing field, and there are provisions for preventing competitors from anti-competitive behaviour.
3. It provides compromises between the anarchy that arises from a completely deregulated use of the spectrum and the lack of access from an over-regulated use.
4. It establishes a clear legal framework for their activities and innovation to continue.

1.4.3 Why is this good for the Australian consumer?

1. It gives rural and regional Australia control over their own broadband destiny, and puts them on a par with the capital cities.
2. The improved business climate will encourage competition and lower prices
3. The resulting network is available to a larger proportion of the population, who may have been otherwise unable to pay for higher speed data services.
4. It gives assurance to Australian consumers that their providers are operating within the freedom of the law.

Chapter 2

The solution

The Telecommunications Act 1997 SECTION 34 states:

- (3) For the purposes of paragraph (1)(h), a network is an *exempt network* if:
 - (a) the network is used, or for use, for the sole purpose of supplying carriage services on a non-commercial basis;
 - or
 - (b) the network is of a kind specified in the regulations.

This document suggests an amendment to add a third alternative:

- (c) the network
 - (i) uses radiocommunications equipment as defined in the Radiocommunications Act 1992
 - (ii) provides open and free public access to all end-users under non-discriminatory terms
 - (iii) provides open interoperability to other networks exempt under this clause (c), transparently forwarding communications between itself and those networks
 - (iv) that no charges, fees or other kind of commercial transactions are levied for the carriage of data over the network

Chapter 3

What the amendment does

The idea of the change in section 2 is to introduce to the Australian telecommunications market an open wireless network which for convenience I will call the “Mesh”. Data transferred across the Mesh cannot be charged for, and the Mesh itself cannot be monopolised by any one party.

This now frees up ISPs in Australia to offer broadband access to customers via broadband wireless. While they cannot charge customers for transferring data over the Mesh, they most certainly *can charge* customers for:

- accessing the internet via the Mesh – here the ISP may charge for data transferal between the Internet and the ISP’s Mesh access point
- storing the customer’s email on a computer which can be reached via the Mesh
- running programs on a server in the ISP’s office that displays remotely to the customer’s computer where that data is sent across the Mesh
- equipment to access the Mesh – naturally they can sell wireless LAN cards to customers, antennas and services for setting a customer up to do this

Under the current legislation, these options are not available to ISPs.

How might this work if the amendment in chapter 2 were made? An ISP sets up a wireless base station to cover its local community, looking something




Figure 3.1: An ISP can use wireless easily to reach customers. . .

like figure 3.1. They cannot afford to pay for a carrier license, so they elect to establish the wireless network under the Mesh clause.

This now allows anyone with a wireless card in the region to transmit data to other members of the local community. The ISP will probably set up some kind of authenticating proxy server for internet access, so only customers who have paid for accounts with the ISP will be able to access the Internet (e.g. surfing the web).

Things get interesting when a second ISP sets up a base station nearby as in figure 3.2. They also do not want to pay for carrier license, so they opt to join the Mesh as well. The second ISP can use the first ISP's wireless coverage to reach customers. This sounds unfair, but of course conversely




Figure 3.2: A second ISP sets up an overlapping cell, and both ISPs benefit

the first ISP can use the second's coverage. Suddenly the first ISP can pick up customers from a much larger coverage area. Both ISPs now have the same coverage (expanded) area.

Possibly another nearby community group may become interested in high-speed internet access. They spend a few hundred dollars (in 2002 terms) setting up a base station to join the Mesh formed by the first two ISPs as in figure 3.3. This third community now has a choice of two high-speed ISPs. The ISPs now have access to an ever larger market without them having to expend any further investment in equipment.

Some months later there are several more cells and one of the outer networks




Figure 3.3: A community group puts an antenna on someone's roof

has slow access because it has to run through congested neighbour networks (figure 3.4). Some of the local businesses jointly pay for a long-distance point-to-point wireless connection from the worst affected area back into the centre of town. This improves the resilience of the network as well, since there are now multiple paths for data to travel along to reach the two ISPs (figure 3.5).

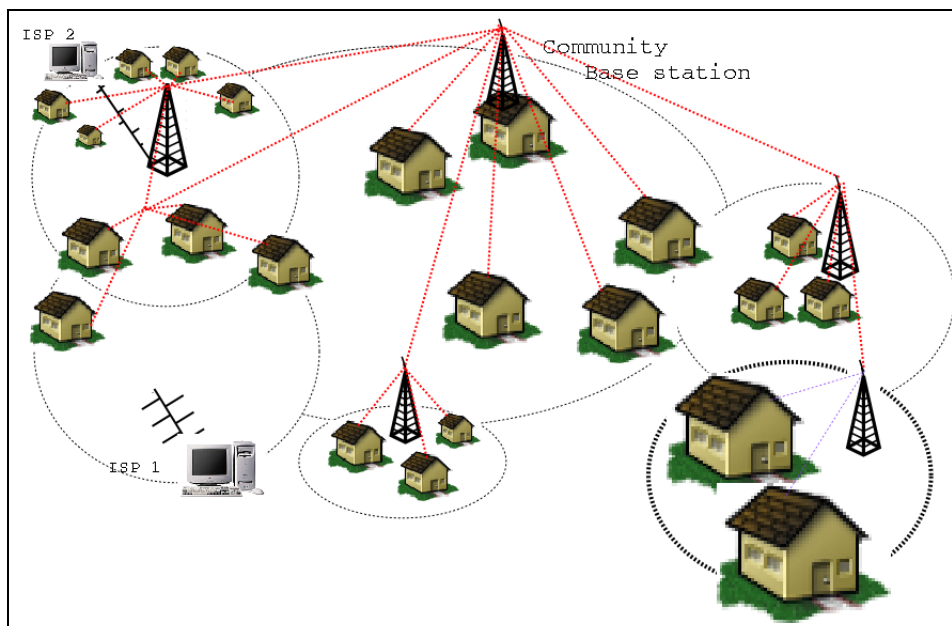


Figure 3.4: The Mesh has grown like topsy and it is a lot of retransmissions to get back to town now for some...




Figure 3.5: But the consumers can spontaneously fix problems themselves if they really want to

Chapter 4

Answering objections to the Mesh

4.1 Will the Mesh get congested?

The Mesh offers free traffic transfer, so it is almost guaranteed to face some kind of congestion as everyone wants to get on board. This congestion is not a bad thing – this means that the available bandwidth is being used to 100% capacity, and will be close to optimal efficiency. But will that be enough?

The Mesh offers an interesting principle that does not occur often in telecommunications networks – that disaffected users have direct control to fix their own problems; for example, if there are too many users in a cell, the cell size can be shrunk. An ISP interested in customers in the area, a local council, a community group, a body corporate or even an individual can easily afford the few hundred dollars required to set up the equipment to form an additional cell, thus doubling the effective available bandwidth. Or they may find themselves reliant on traffic through a neighbouring congested cell, in which case they may set up a directional antenna to long-haul their data closer to where they want it¹. If any kind of congestion is enough to cause disrup-

¹ Wide area 802.11 networks can be done with omni-directional antennas (providing access for large numbers of people in a relatively small area) or directional antennas (providing a link between two points along a line). The advantage of directional antennas is that there is not a lot of interference – two directional links can cross over each other as if they are not there. The only interference is if you are on the far side of a directional antenna and some signal is leaking past. Directional antennas will be relatively rare in the

tion and economic loss, those affected by the disruption are in a position to resolve the issue themselves.

Having said that, it is nonetheless worth discussing the kinds of congestion the Mesh may experience.

Intra-cell congestion where the number of users in one cell is so large that they do not get acceptable bandwidth.

Inter-cell congestion where the amount of traffic going from one cell to another is too large.

4.1.1 Intra-cell congestion

Intra-cell congestion can be resolved by reducing the cell size, so that there are fewer users in each cell. Some rough figures may be informative here. Sydney, with the densest population, is likely to suffer the most from inter-cell congestion, so how bad is it going to be?

Currently the volunteer-effort Sydney Wireless community (www.sydneywireless.com) has approximately 1000 interested users – people who would be happy to be hosting a base station if they are not already². The numbers are still growing, and it does not seem unlikely that they could reach 2500 potential stations across Sydney. For this generation of wireless technology, 10% of Sydney's population (approximately 5 million) might well come on line wanting to use it, of which maybe 50% might be active during peak time (i.e. 250000 concurrent users) and the base stations should be able to handle 11Mb/s (11000 kb/s).

This can be approximated by the following highly simplified equation³.

Mesh, since a lot of traffic may hop from one (omnidirectional antenna) cell to the next.

Of course, there was no particular reason for a directional antenna here. Perhaps an enterprising ISP might put a wireline or microwave link in and charge customers for transferring that traffic – perfectly allowed within the Mesh rules

²Actually, I have taken some liberties here. Most of the 1000 people have only registered that they want to be part of the network, not necessarily look after any part of the infrastructure. However, from my dealings with the kinds of people involved, I think they would be happy enough with a slightly expanded role.

³A few assumptions have been overlooked. Firstly, network collisions reduce the effective available bandwidth down somewhat. This can be as high as 20% in extreme geographic cases. Secondly, I have assumed a uniform distribution of population around

$$\begin{aligned}\text{Max bandwidth per person} &= \frac{\text{Bandwidth per cell}}{\text{Concurrent users per cell}} \\ &= \frac{\text{Bandwidth per cell}}{\frac{\text{Number of concurrent users}}{\text{Number of cells}}}\end{aligned}$$

Filling in the numbers, we calculate:

$$\begin{aligned}\text{Max bandwidth per person} &= \frac{11\text{Mb/s}}{\frac{250000 \text{ people}}{2500 \text{ cells}}} \\ &= 0.11\text{Mb/s} \\ &= 110\text{kb/s}\end{aligned}$$

So congestion will happen when everyone is pushing around 110kb/s. This puts it well ahead of dial-up and single-channel ISDN, but a little behind ADSL or cable. But this is a *worst case* analysis, and most of the time it should do better. This rules out video conferencing and video on demand, but certainly does quite well for a lot of other broadband applications.

Notice that 250000 people sharing 2500 cells is 100 concurrent users per cell. The more people sharing a cell, the lower the bandwidth.

So rural areas will do vastly better than city regions. Population densities in rural Australia are often below 1 person per square kilometre. An 802.11 cell can reasonably cover 60km² a region, suggesting less than 60 people per cell. Of course, not all of those 60 users will be active at any given time – perhaps only half at worst? Let us run the same calculation for a country region:

$$\begin{aligned}\text{Max bandwidth per person} &= \frac{11\text{Mb/s}}{\frac{30 \text{ people}}{1 \text{ cell}}} \\ &= 0.367\text{Mb/s} \\ &= 367\text{kb/s}\end{aligned}$$

base stations.

This is quite acceptable; it is on par with satellite, better than low-speed ADSL and not far behind cable.

I think then, it is safe to conclude that if a cell is not having to forward traffic from another cell, that no unmanageable congestion will occur. So, at the very least, **the Mesh will solve Australia's last-mile problem.**

Why is this? The last-mile problem is that it is easy for ISPs and telecommunications companies to put large numbers of points of presence around a city, but difficult for them to run the last mile to the customer premises. The worst case for the Mesh is that it is only able to do the last mile; the best case is that it does the last mile and also acts as a metropolitan area trunk. Whatever happens, the last-mile problem is solved.

Note that only with the provisions outlined in section 2 is this the case – if we continue with carrier control, there is no incentive; and if we release all regulatory control, then ISPs and telecommunications will have to make individual agreements with hundreds of different wireless cell operators.

4.1.2 Inter-cell congestion

This is extremely hard to predict, and will depend heavily on the kinds of applications that people deploy, and to some extent on how the Mesh grows. What are the obvious applications?

- Sending email
- Playing computer games
- Voice-over-IP
- Surfing the web

All these are detailed in the next sections below and cover worst-case scenarios. It may transpire that the network and applications develop in ways such that inter-cell congestion does not turn out to be even as much of a problem as outlined. In particular, all the above calculations were based on 11Mb/s connections, when 150Mb/s equipment is already in the foreseeable future.

Email

Email is not a big problem at all. An email that takes 1 minute to arrive at its destination instead of 1 second is rarely a problem. Traffic prioritisation will handle SMTP quite effectively even through absurdly high congestion.

Computer games

Computer games are very latency sensitive, but not particularly heavy on bandwidth. Also, gamers seem to play at non-peak times (2:00am in the morning seems to be normal gamer behaviour!). This means that even if gaming takes over the resources of the whole network, it will probably not be at a time that others will be inconvenienced. To give a numeric example, Quake⁴ for instance, rarely if ever asks for more than 20kb/s. Even if every byte of this data were relayed to every point in the network, this would still allow up to 550 people ($550 \times 20\text{kb/s} = 11\text{Mb/s}$) to be playing simultaneously from each server. While it would be nice to be able to support more than this, it is probably enough. What this means is that gaming communities will have to form around geographic regions⁵.

Voice over IP

Voice over IP is an interesting one. I suspect that it won't work reliably, although I can't pin down numbers to show this. If most phone calls are to other people within one cell, then there should be few problems. However, this is unlikely, and at a guess, most telephone calls would surely hop through more than 25 (city-based) cells. Equivalently, a cell that initiates a Voice-over-IP call would have to support forwarding another 25 at the same time. How many concurrent calls can a cell handle then, if each call uses (at 9.6kb/s?)

$$\text{Number of initiated calls} = \frac{\text{Bandwidth available}}{\text{Number of forwarded calls} \times \text{Bandwidth per call}}$$

⁴Presumably there should be a (TM) here somewhere.

⁵This is actually a highly desirable result since it may well introduce a sense of greater community if there is a chance of gamers meeting easily off-line.

$$\begin{aligned} &= \frac{11\text{Mb/s}}{25 \times 9.6\text{kb/s}} \\ &= 45 \end{aligned}$$

45 concurrent calls is not very many. It is probably insufficient for reliable usage, which is good. This means that there is still a market for a carrier-run network to handle voice – carrier networks are not under the same restrictions as to provide open wireless access and have much better hope of handling congestion intelligently.

Web surfing

Web surfing is the final most interesting application. In the worst case scenario, of a Mesh that is so congested that no inter-cell bandwidth is available, ISPs would need to put an access point for their network into every congested cell. This is exactly the same as all other existing broadband technologies, only without the expense of solving the last-mile problem. On the other hand, if inter-cell bandwidth is adequate to provide a reasonable broadband experience, then the end result cost is much cheaper for an ISP than rolling out any other kind of broadband service to their customers.

4.2 Is the Mesh scalable? Is it reliable?

Although there is really no fundamental requirement for it, the mostly likely traffic to be sent over the Mesh is internet (TCP/IP or UDP/IP) traffic. The internet protocols already include dynamic routing capabilities.

The routing structure of a mesh of more than a few wireless networks is probably going to be too complicated for any human can understand. The only way it will be useable is through BGP, EGP or some other kind of dynamic route table update mechanism.

The update times on these protocols are somewhat slow (often several minutes) because they were designed for a somewhat more rigidly structured internet, but they are certainly capable of being adapted to cope with the faster topology changes. It is probably possible to adapt them to cope with load balancing across congested interconnects, spreading traffic across as many adjacent networks as possible.

The Mesh would be characterised by vastly redundant interconnections; each cell overlapping most of its neighbours. The “route around damage” model of the internet’s dynamic routing tools would have plenty of alternate paths to work with. Once there is a good number of base stations in a geographical area, reliability of the network should be among the best ever seen for any kind of network infrastructure.

The Mesh’s scalability is likely to be quite good as well. While the number of interconnections will be higher than the greater internet, then actual number of networks may be quite low. It is unlikely that there would be a need for more than 50000 networks to cover Australia. This is several orders of magnitude smaller than the Internet, which manages with quite minimal central control. If the Mesh uses the current or next generation internet protocol (IPv4 or IPv6), there is no reason to expect it to scale less efficiently.

4.3 How does the Mesh protect itself?

One of the great dangers in this kind of network is subversive activity. The greater internet in the last decade has been a story of denial-of-service attacks and hacker vandalism. With an open access policy, how can the Mesh eject misusers?

There are two main techniques:

Prioritisation Nothing in the terms of the amendment requires an operator to give equal priority to all traffic. The operator has a right to make sure that traffic from a neighbouring rogue cell only gets transferred if there is absolutely nothing else needing to be sent. And this is sufficient for many purposes – if a greedy user wants to use up all available network bandwidth, then they can have whatever is left over after everyone else has had their share. Routing traffic is essentially cost free, so the operator incurs no imposition for letting a greedy user do this.

Terms An operator may give access to a network under non-discriminatory terms. Although no charge can be levied, access can be restricted to only those who agree to certain non-onerous network behaviours. For example, an operator may only allow traffic of a certain quality of service (e.g. deny all real-time traffic). Or they may condition on the

traffic being “non objectionable”, giving operators an excuse to refuse traffic from known misusers.

Of course, the Mesh does not guarantee anonymity. In fact, it almost completely precludes it, since you are easily traceable to the base station you are coming from (which puts you in a circle of a few kilometres radius). Moreover, you are running a radio transmitter, which can be triangulated even further to find out exactly where you are. This in itself makes misbehaviour rather less worthwhile.

4.4 Will people be concerned about base stations?

With much furore over mobile phone base-stations, will people be happy to see thousands of 802.11 base stations dotting the landscape? I believe the answer is yes, since the community will exert a much greater degree of control over their placement.

For mobile phone stations, it is a large national or international corporation that chooses where it will deploy. The deployment is placed to maximise the revenue that that company can earn. Residents may object, but there is little chance of their objections being heard or acted upon.

With base stations for the Mesh, the deployment will often be as a result of a few individuals having an idea and wanting better internet access. The base station will be there clearly for the benefit of local residents, and the person responsible will probably be local. If there is a concern about its location, it becomes much more practical to find a good alternative if the instigator is available and present.

4.5 What is the economic model behind the Mesh?

Unlike most other broadband transmission, radio bandwidth can be viewed as close to a pure public commodity.

non-rival consumption The lines of point-to-point links can happily cross over each other without interference (in nearly every case). Once the safety and equipment-usage procedures are in place, there is almost no extra cost associated with allowing additional users.

non-excludable consumption Short of equipping a policing force with signal strength meters and establishing a large-scale licensing scheme, the equipment required is too affordable and easy to set up to enable any kind of exclusioning.

Economic theory teaches us that the commercial market cannot efficiently allocate resources to a pure public commodity. The best way to make use of radio bandwidth is therefore to empower community groups, individuals, charities and governments to establish the infrastructure.

4.6 Who will invest in Mesh points?

Given that section 4.5 has ruled out wide-spread large-scale commercial investment, a related question is “is this whole idea an idealistic pipe dream?”. The answer to this is an emphatic **no**. As I write this there are wireless broadband communities in every capital city and many large regional centres. These communities are unable to connect to the internet to provide a vastly better service to their members purely because the legal infrastructure does not allow them to do it in any reasonable way. At the same time, there are several ISPs who would be delighted to offer access services to these communities, and who can do so profitably. Simply announcing the legality of the connection would allow the demand for broadband wireless services to be unleashed within *days*.

Back to answer the question:

- Individuals have already shown an interest in subsidising the development of these networks, as discussed in the previous paragraphs.
- Rural shire councils could be pushed by their constituents to put up some omnidirectional antennas on a few tall buildings or mountains to allow their whole shire to get high-speed internet access.

- Charities and computer user groups may choose to help impoverished areas by establishing base stations and connectivity. Such activities have already been seen in a few places around the world.
- Politicians in marginal electorates may well get highly-developed wireless infrastructures⁶.

⁶Of course, I'm not condoning this behaviour, merely acknowledging that these things happen. Ironically, unlike most pork-barelling it would probably work since knowledge industries would be attracted to a geographic region with extremely high bandwidth connectivity.

Chapter 5

Conclusions

If we allow companies and individuals to use consumer-grade wireless network equipment to form an open and public network, we can:

- Alleviate last mile problems in Australia.
- Build a state-of-the-art national broadband infrastructure extraordinarily cheaply.
- Avoid wasting the available bandwidth resources.
- Preserve investments of existing companies, both carriers and non-carriers.
- Improve the ability of consumers to control their access to telecommunications.
- Improve the quality of technology available to all Australians.

This is possible, simple and achievable. It merely requires the expanding of one regulation in the Telecommunications Act 1997 , as discussed in chapter 2.

Index

- 3G Mobile, 6
- ADSL, 6
- Anarchy solution, 5, 7
- Anonymity, 24
- Antennas
 - Directional, 17
 - Omnidirectional, 17, 25
- Cable, 6
- Carrier license, 5, 6, 9
- Charities, 26
- Commercial investment, 6, 25, 27
- Commercial investments, 6, 7, 14
- Congestion, 18
 - Inter-cell, 18, 20
 - Intra-cell, 18
- Criteria, 7
- Directional antenna, 17
- Dynamic routing, 22
- Economic model, 24
- Example of Mesh development, 12
- Inter-cell congestion, 18, 20
- Interference, 6
- Intra-cell congestion, 18
- Last-mile problem, 20, 22, 27
- Mesh
 - Carriers, 9
 - Charging, 12
 - Congestion, 18
 - Consumers, 10
 - Definition, 8
 - Design, 8
 - Example, 12
 - Non-carriers, 9
 - Protecting itself, 23
 - Reliability, 22
 - Rules, 8
 - Scalability, 22
- Non-excludable consumption, 25
- Non-rival consumption, 25
- Not-for-profit organisations, 5
- Omnidirectional antenna, 17, 25
- Prioritisation, 23
- Private networks, 5
- Radiocommunications Act, 11
- Regional ISPs, 5
- Rural access, 5–7, 25, 27
- Status quo solution, 6, 7
- Sydney Wireless, 6, 18
- Telecommunications Act, 4, 8, 11, 27
- Terms and conditions, 23