June 2008

Understanding Mesh Hardware: The Evolution of Wireless Mesh Networking

What is Wireless Mesh Networking?

Mesh Networking is a technique for routing network traffic that automatically creates efficient pathways and guides traffic around any hardware disruptions. Wireless Mesh Networking achieves this goal using radio transmitters and one of the wireless data standards, such as the many IEEE 802.11 standards. The route to any particular destination in the network, through the Mesh radio units or “nodes”, is continuously evaluated, and if there is a better (e.g., more reliable or efficient) route available, the connections between nodes are reconfigured without administrative intervention. There are multiple technologies that are capable of creating the Mesh, such as OLSR, AODV, OSPF, and many others. Each meets the basic goals of the Mesh, but emphasizes different details such as end-end throughput, route recalculation speed, link status detection, and so forth. The focus of this discussion is proving client access using wireless Mesh. The points made are just as applicable to more of a fixed wireless, backhaul only application. If this is your application, just ignore the first generation and client-related problems below.

Advantages of Mesh

In a traditional Wireless Network there are often single points of hardware failure, and the whole network is vulnerable to transitory interference at any one point. If you combine this with applications or users that are intolerant of slowdowns or outages, such as Voice over IP, IP video surveillance, or Internet, you have a painful situation. Wireless Mesh Networking can address these problems in two ways:

- Self-configuring

  In OLSR, the state of each link is evaluated at a configured interval, typically one to five seconds. If a particular link is degraded because of congestion, interference, or other reason, that link is flagged as bad and traffic is routed around, if possible. The Mesh architecture also allows for multiple WAN connections, which allows traffic to avoid a whole congested sector of the network and reach the Internet through an alternate gateway.

- Self-healing

  Building on the link evaluation described above, if a link is down entirely the network can be reconfigured to provide connectivity. The Mesh can even detect if the WAN gateway, from the ISP for example, is not available. In this case all Internet bound traffic is redirected to the backup WAN gateway, which could be an entirely different ISP for maximum safety.

What is not Wireless Mesh Networking?

“Mesh” is a network route configuration technique, but equipment manufacturers often associate it with other high level or proprietary Wireless Networking features. This has tended to bundle the very important Mesh technologies with other high cost features, and made Mesh in general not very affordable. You should evaluate these other features independent of the reliable and efficient Mesh architecture, as they may entail significant additional cost.

- “Zero” configuration

  These devices are just deployed straight out-of-the box, and when powered on try to discover an existing Mesh network and add themselves to it. This could be a very handy feature for networks that are set up, moved, and broken down frequently, or installed by inexperienced personnel. In a permanent Mesh infrastructure, these devices maintain numerous unnecessary connections and may make bad “automatic” decisions about radio channel, power, Mesh peers, or other settings. This should not be confused with the “self-configuring” routes fundamental to the Mesh, which choose among properly provisioned Mesh nodes.

- Full Duplex Wi-Fi

  Some manufacturers have made proprietary modifications to the 802.11 standard to separate the receive and transmit functions, which are ordinarily shared on a single radio, to two dedicated radios. There is no technology downside to this, assuming it all works correctly, but it does at least double the hardware costs, for twice as many radios, plus the cost of the proprietary technology. In considering the cost, you should determine if your network will have heavy usage both upstream and downstream, or if it will mostly consist of subscribers downloading large files and HTML pages, downloading and sending small emails, and so forth.

- High throughput Wi-Fi

  Similar to full-duplex Wi-Fi, some manufacturers allow two radios to be connected to the same endpoint, theoretically doubling the throughput. Again the only downside seems to be cost. In designing a network, you should try to determine if you will have more than 54Mbps (20-25Mbps actual), but less than 108Mbps (40-50Mbps actual) at any single node. For general network performance this may not be necessary, as some Mesh standards, like OLSR, take the current utilization of links into account, and so may route traffic away from crowded links. In the case of a single chokepoint, a hardwired Ethernet or dedicated point to point link might be worth considering. Many radios also have a “turbo mode” that uses two adjoining channels, so in theory can provide 108Mbps. General testing reports are that these turbo modes provide about 5Mbps additional throughput outside the lab (AKA the real world) due to additional interference.

The Evolution of Wireless Mesh Networking

Like many technologies, Wireless Mesh has evolved from a bleeding edge over-hyped application to a cutting edge and effective solution. We could go back as far as twenty years where Wireless Mesh starting being discussed and funded at DARPA, the source of the Internet itself, not to mention mind control and death rays. However, we are only really interested in the commercial Mesh solutions that have become available in the last six years or so.

Single Radio Access/Mesh (First Generation)

The commercial Wireless Mesh units that became available around 2002 used the simplest possible hardware architecture. A single radio maintains the backhaul links to all nearby Mesh nodes, and the same radio also services all access requests from Wi-Fi clients. These units were very expensive, like any bleeding edge technology, but with press adulation and promised like “ROI in 4 months” and “300% ROI in the first year” the first Mesh generation was hard to resist. At least until wide scale deployments were made. An example of the first generation single radio Mesh would be the original (and current single radio) Tropos HotZone.

Advantages of First Generation Mesh

As the first generation of Mesh products, the only comparison at the time was against traditional access point (AP) Wireless Distribution Systems (WDS) or networks of hard wired access points.

- Self-healing and self-configuring

  By definition the first generation Mesh is a true Mesh, so it has the advantages of the Mesh architecture.

- Ease of deployment

  Having a single box that does it all is an advantage in deployment, since just a single mount, data, and power connection must be provisioned. There is no real advantage here over a traditional AP/Bridge or Dual Radio AP that would be deployed the same way.

- Advanced features

  Some first generation products include advanced features like "zero configuration" not available in an AP. As discussed previously, the need for these technologies should be assessed separately from the need for a safe and reliable Mesh architecture.

- Relative Cost

  A Wireless Mesh Node costs more than a simple Access Point due to the increased computation required to maintain the Mesh routes in real time, not to mention the fact that is it just better. In general the first generation Mesh units are going to cost less than second or third generation units. However, this is may only be true within a particular product line. Certainly some vendors will be offering second or even third generation Mesh units for less than some available first generation units.

Disadvantages of First Generation Mesh

The most serious problems with the first generation Mesh architecture are inherent in the single radio design. Some of these are just inevitable limitations, while some are just bad design choices.

- Backhaul and client access share bandwidth

  The 54Mbps (20-25Mbps usable bandwidth) available on the 802.11g radio in a first generation wireless Mesh network must be shared between every wireless client seeking LAN or WAN access and every backhaul radio with data to send or receive. The nature of 802.11 is that radio resources are shared more or less evenly between stations requesting to send or receive data. So in the case of the first generation wireless Mesh, subscribers with laptops and Mesh data backhaul link are treated equally. Clients sharing access on a regular AP work fine. If you have 20 clients on an AP that can send/receive about 20Mbps, they are each getting about 1Mbps. But in the case of the wireless Mesh, a Mesh node with 9 clients and 1 active backhaul link is only going to provide about 2Mbps per link, including the critical backhaul link that everyone downstream is depending on. See Figure 1. Looking downstream, if there are 10 clients one hop away, they are dependent on that 2Mbps backhaul link, which is only 200Kbps each. This is marginal service, but you might continue this disaster another hop downstream where 10 links are sharing 200Kbps. At 20Kbps many HTTP pages will not even load before timing out. Let us not dwell on 3 hops . . . Of course in reality not everyone in the network is broadcasting all the time, so generally data squeaks through up to a hop even with a reasonable number of users. But from the example we can clearly see the danger that some clients may be shut out, and certainly scalability of the Mesh is basically non-existent.

  This problem was not new to wireless Mesh. The same problem exists if you create a non-Mesh wireless network using AP/Bridge access points. You just can not allow clients on the backhaul radio links. In the best case they are consuming 50% of their own backhaul, in worse cases they are completely drowning out downstream users. If the minor cost savings can not be resisted, this configuration, Mesh or otherwise, should only be used at edge or endpoint units that have no downstream clients.

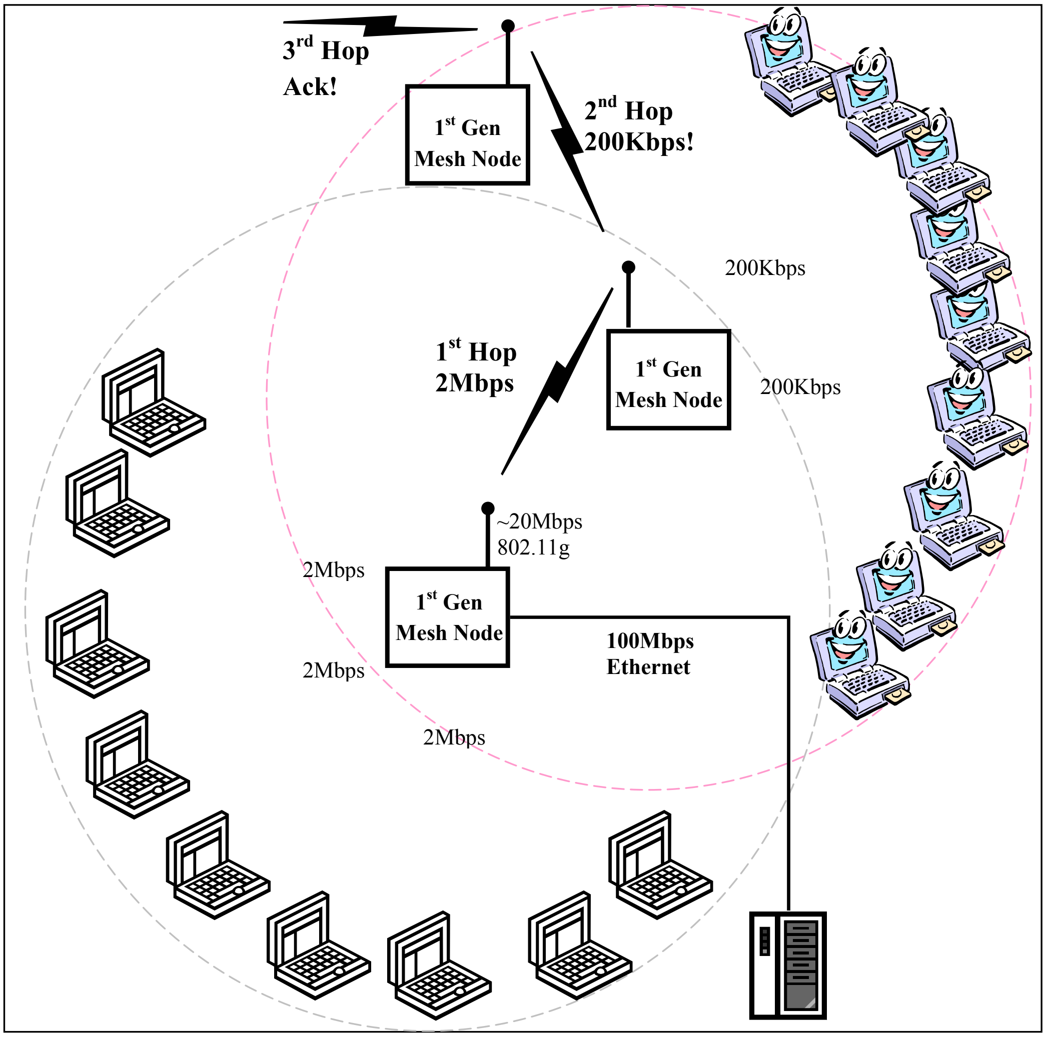

Figure 1: Sharing Client Services and Backhaul on One Radio

- Internal interference maximized

Every radio, client and Mesh, in a first generation network must be on the exact same channel. This includes the Mesh units themselves as well as any connected laptops, CPEs, or other Wi-Fi equipment. This causes tremendous contention for the use of that one channel. The way Wi-Fi is designed, only one radio uses the channel and the others wait. A single radio also only transmits or receives, so every transmit from every radio in range must wait for its turn. Really the radios use a random back off rather than organized turn taking, but the effect is the same under normal circumstances (see hidden node below for problems with this system). The shared bandwidth of everyone taking turns is still only 54Mbps

(20Mbps actual). In a very small network, or single AP, this is fine, as the switch off between radios and turn taking is very fast. But when the backhaul must share radio time this way, each additional laptop and Mesh unit in range reduces the total available throughput.

Looking at Figure 2 illustrates this. It is interesting to note that the actual connection speed point to point when everyone else is waiting is 54Mbps (~20Mbps useable). But this is only for a few, or maybe a single, packet. The speed estimates in Figure 1 take into account that fact that everyone is doing more waiting than sending/receiving.

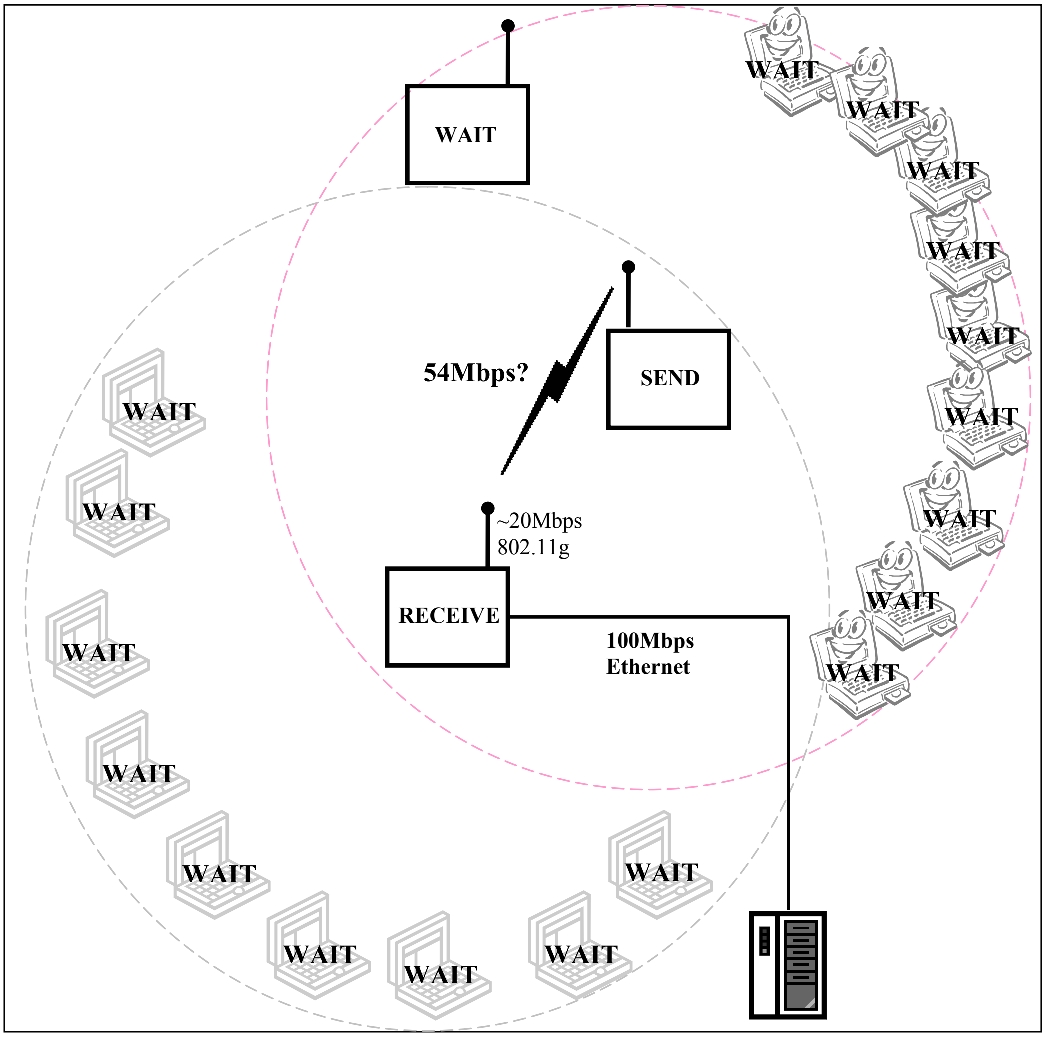

Figure 2: Radio must wait for every other radio in range

- Extremely limited multi-hop throughput

Due to every radio in the network being on the same channel, whenever there is transmitting radio every other radio in range must wait. In the worst case scenario, every backhaul link in the network would be waiting for every other radio, limiting the end to end throughput for N additional hops to (1/2)^N, e.g. 1=10Mbps, 4=5Mbps . . .10=20Kbps, etc. In the best case only pairs of radios would wait, limiting throughput to 1/(N+1), e.g. 1=10Mbps, 2=7Mbps, 10=2Mbps, etc. Performance in the real world will fall somewhere between the two.

- Backhaul and client links cannot be effectively served by a single antenna

  A single radio of course means a single antenna. However, data backhaul and client access have different requirements in terms of signal propagation. See Figure 3a and 3b for the two mounting options for a first generation Mesh. The backhaul links need to as high as possible, to avoid both free standing obstacles like building and trees, as well as obstruction of the Fresnel zone from the ground itself. Clients, by contrast, are located at ground level for the most part: in a home, a first floor business, boats in a marina, or in an open area. One kind of link will lose out wherever the antenna is located, and you cannot do without a solid link for either (choose the backhaul if you must. . .). In example 3b, even there was no tree, the client is still will not have a good connection because of the antenna deflection.

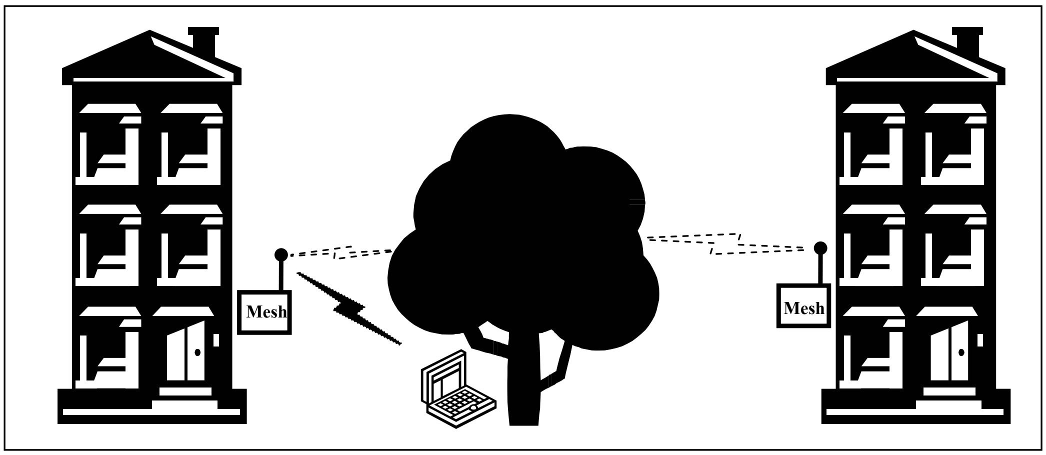

Figure 3a: Antenna location too low for backhaul

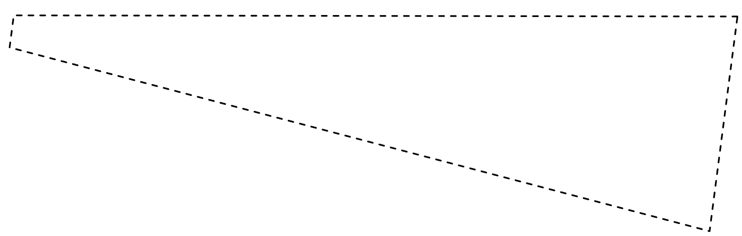

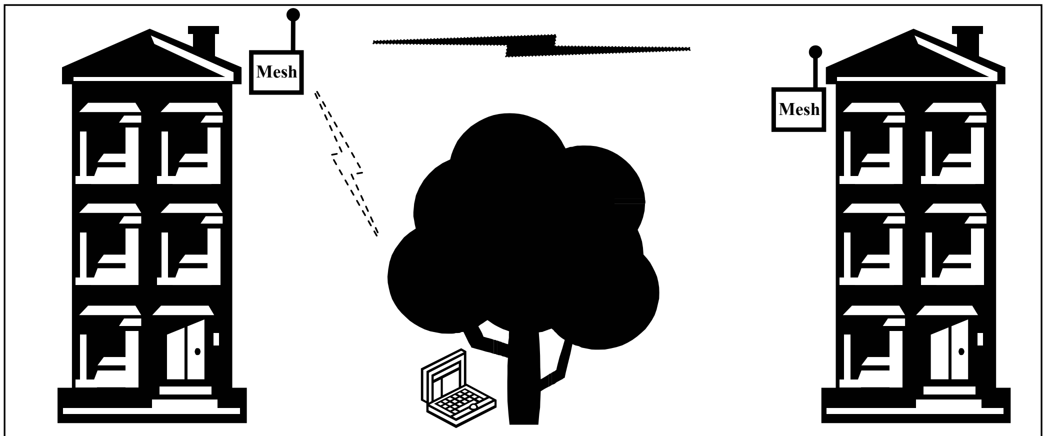

Figure 3b: Antenna location too high for clients

- Impossible to address external interference

As a Wireless Mesh network scales geographically, you will encounter strong interference from various sources. This could be another WISP or Wireless Mesh, commercial or HAM radio, airports, industry, high tension power lines, cell towers, or any of a hundred other things. If there are only one or two sources of such interference in your location you can switch to the third channel. If there are more, however, you are totally stuck. See Figure 4. The single radio in every first generation wireless Mesh must be on the same channel, so if you change channel to avoid interference in on spot, you make it worse in another spot. You could possibly avoid this doubling up Mesh units at a location with an Ethernet link, but then you are really just creating a DIY second generation wireless Mesh.

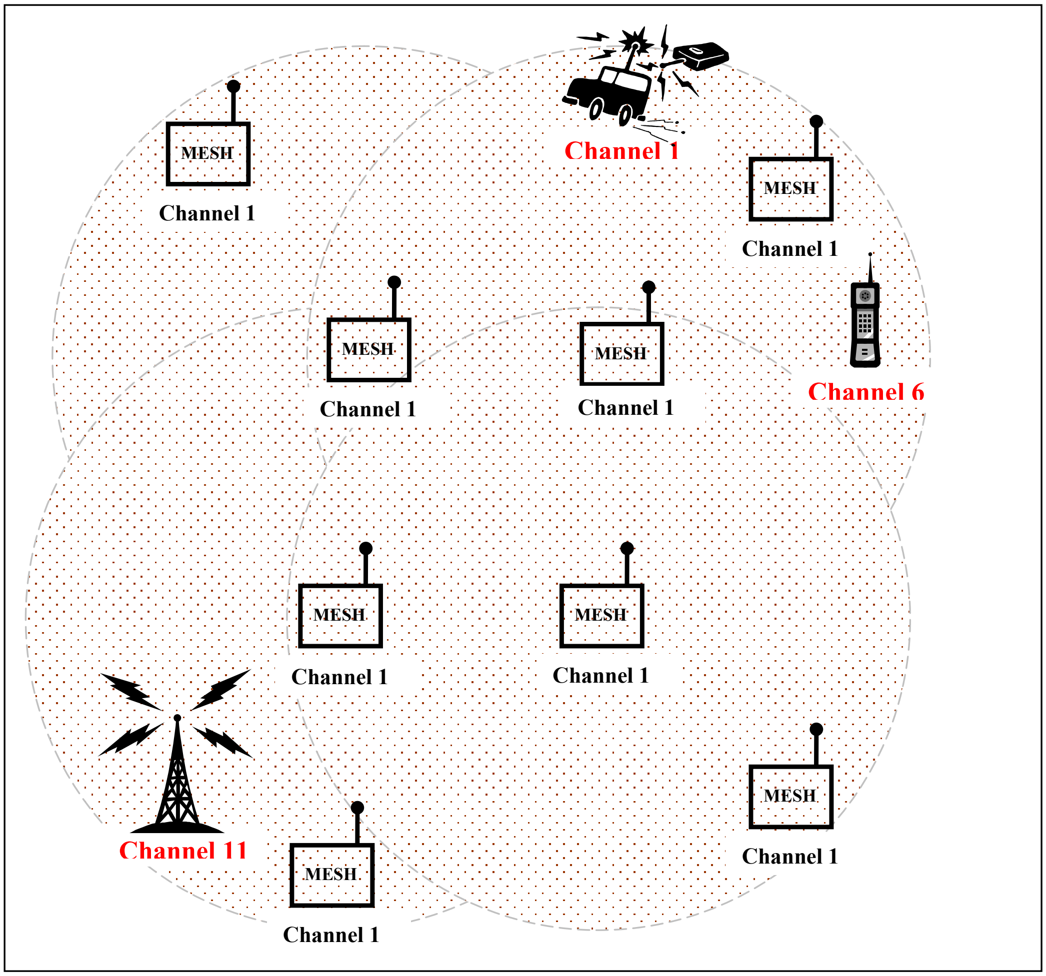

Figure 4: External Interference

- Inflexible spectrum choice

  Combining the client and backhaul services into a single radio means, realistically, that you are limited to only 802.11b/g. There are just not really any 802.11a, (or 900MHz, 4.9GHz, etc.) based laptops. 802.11b/g is the most crowded spectrum and limited to only 3 non-overlapping channels. 802.11b/g subjects the wireless Mesh backhaul to constant low-level interference from every wireless router and laptop in the world, not to mention garage door openers, radio controlled toys, pagers, and any of a thousand other things.

- "Hidden Node" problems

  Also due to every radio being on the same channel, if two Mesh nodes or clients have a peer in common, but are too far away to sense each other, transmit collisions will occur. Wi-Fi is dependent on the radios being able to hear each other to avoid conflict, but if two radios talk to a third radio in the middle, the third radio cannot tell them to take turns. See Figure 5. This is like two people whispering in your ears.

They cannot hear each other and you keep saying "What?". All the third (middle) radio can do it keep telling each side to repeat what they just sent. This results in much lower, or no, data throughput between the three radios. Wi-Fi has a failsafe traffic management technology called RTS/CTS that can be enabled, but in many cases the extreme overhead from RTS/CTS worse than the "hidden node" overhead.

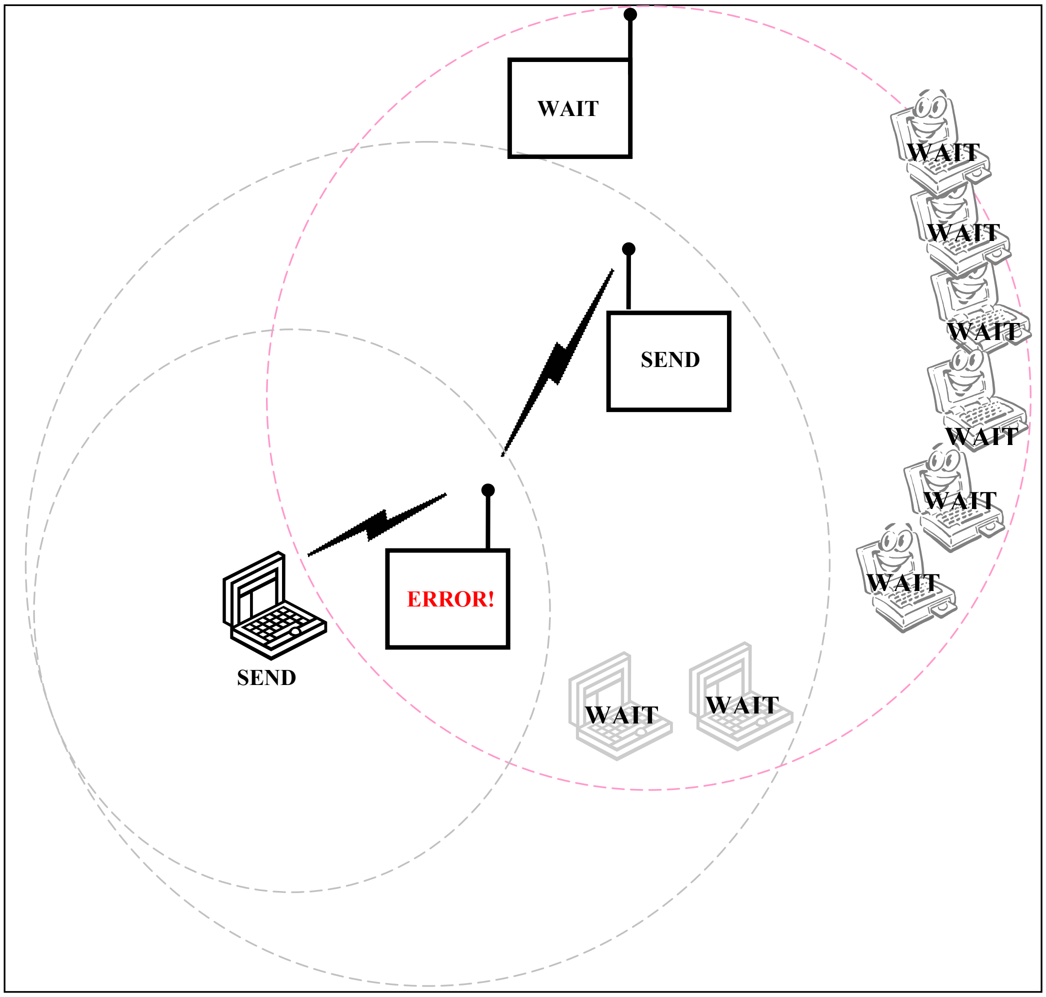

Figure 4: Hidden Node Problems

- Misbehaving clients are directly connected to the Mesh backhaul radio

  There are a tremendous number of Wi-Fi clients out there and a tremendous amount of different Wi-Fi hardware. While theoretically every piece of 802.11 hardware is compatible, some may work better, in combination or alone, than others. It is also impossible to predict what clients may have going on in their laptops. If they cause some radio problem maliciously or inadvertently, the shared backhaul is lost as well. The Mesh self-healing design mitigates the risk somewhat, but it is an unnecessary risk. If a bad client brings one unit down, they may very well be in range of additional units and cause problems there as well.

Conclusion

The first generation design is still here with us today, and seems to comprise much of the low end of the Wireless Mesh market, particularly the down-market products from some commercial Mesh vendors. Although these units are the lowest cost in most cases, careful consideration must be taken as to whether a network of first generation Mesh units is going to work AT ALL for your application.

The single radio backhaul/access architecture was a bad design in 2002, and worse today that there are better Mesh options. The idea of allowing clients to consume their own backhaul bandwidth for AP access was a bad AP/Bridge design before it was a bad Mesh design, so it really has nothing to do with Mesh.

That being said, for a very small temporary network of a few nodes, some of the plug-and-play Mesh units that you just power up and let them do their thing could be handy. Just keep in mind that if a first generation wireless Mesh network scales beyond a few dozen users or a couple of hops, anyone not connected directly to the base Mesh unit is going to get little or no service.

Another possibility for first generation Mesh hardware is to use it for fixed wireless, and configure the units to keep the clients off the backhaul radio. In this configuration it functions the same as the second generation Mesh units. Deploying a first generation Mesh as a hardwired CPE would be a good application, for example.

Dual Radio Access/Mesh (Second Generation)

The second generation of Mesh hardware added a dedicated radio for client connections.

The potential Mesh integrator may be daunted by the list of problems above. The developers of Mesh products were daunted as well, and in seeking to correct the first generation shortcomings it was wise to focus on the worst problems first. Clearly the worst problems are related to clients sharing the same radio as the backhaul. These are the issues that can keep the Mesh form working at all with more than a few users and a few hops. A second radio just for clients was the obvious and straightforward solution. However, problems related to the Mesh backhaul radio itself, independent of clients, still remained. In many case these problems were improved from the first generation, since the clients made them terribly severe, but not gone. The Second Generation Mesh hardware most often appears as a “Dual Radio”, where two radios are connected to a single board in a single enclosure. More recently, however, popular “1+1” systems have become available from a few manufacturers, where the Mesh Backhaul and Access Point are in separate enclosures, but easily install with a single power and data link. The ValuePoint SuperMesh 8000 is an example of a “1+1” Second Generation Mesh.

Advantages of Second Generation Mesh

- Dual independent radios resolve some First Generation problems:
    - Client and backhaul bandwidth no longer shared

      Backhaul and client bandwidth are no longer shared, so every client in the network shares the available bandwidth in the network, rather than consuming the backhaul traffic of downstream clients. Any problems of overloading a particular point will remain local to that point only, which is only fair. An example network is provided in Figure 5. Client throughput is evenly divided, and the backhaul is adequate for the first couple of hops.

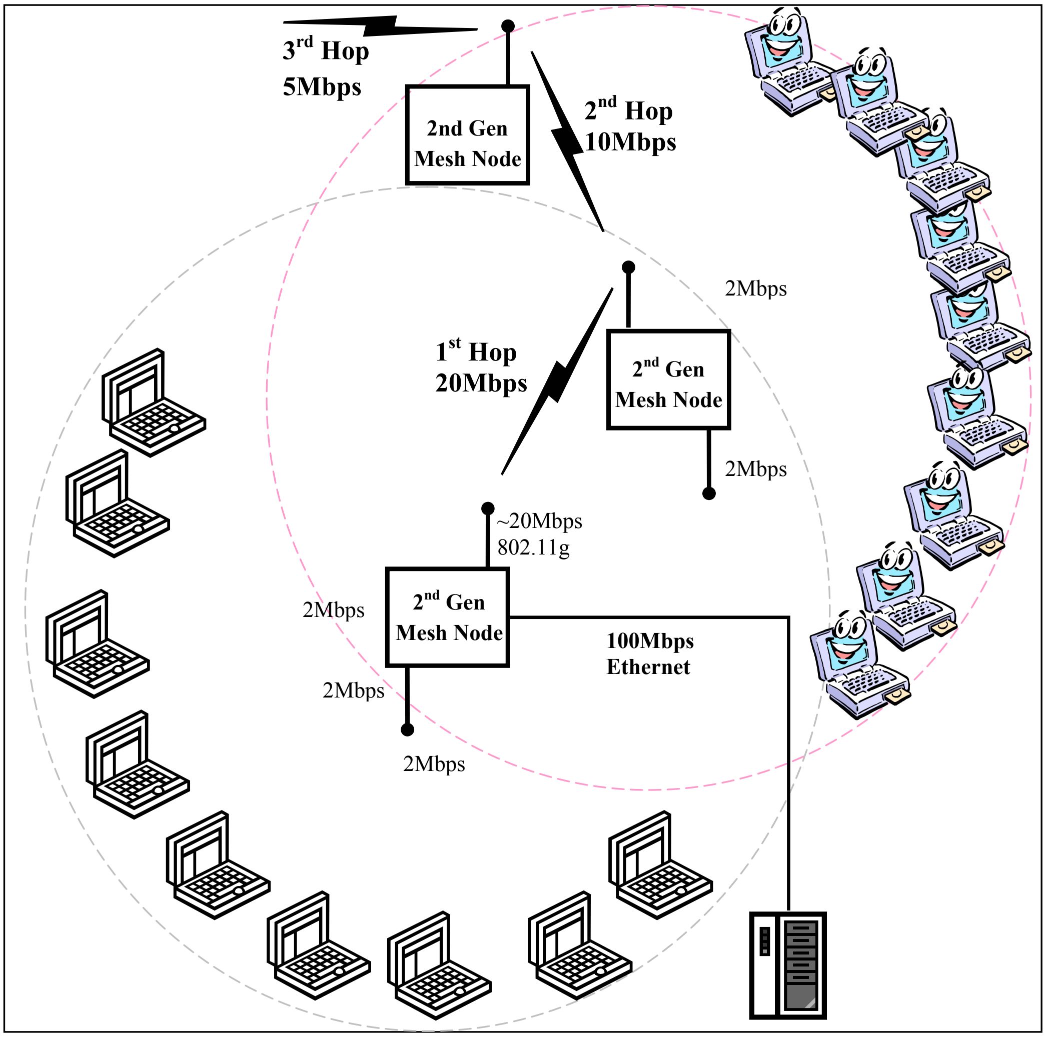

Figure 5: Second Generation Mesh Network

- More flexible spectrum Choices

  The Mesh backhaul radio no longer needs to connect to any clients, so really any spectrum can be used. The most typical is 802.11a on 5.8GHz, but proprietary radios for 900MHz, 4.9GHz (Public Safety Band) or other channels are available as well (at increased cost of course).

- Misbehaving clients are isolated from the Mesh backhaul radio

  The Dual Radio design keeps client mischief off of the backhaul radio. Even if that radio has a problem it will not affect the network as a whole. The 1+1 design takes it a step further and moves clients onto an entirely separate board, which is even better.

- Self-healing and self-configuring

- Ease of deployment

  For Dual Radio Mesh, we have a single box that does it all, similar to the first generation units. "1+1" systems typically are designed to be easy to install as well, with one unit providing power and data to the second unit. Both systems are still just a single mount, data, and power connection. However, in both cases there is a little more to install and provision than the first generation Mesh.

- Advanced features

  As with the first generation, some Second Generation Mesh products include advanced features like "zero configuration". As discussed previously, the need for these technologies should be assessed separately from the need for a safe and reliable Mesh architecture.

- Relative Cost

  Second generation units are going to cost more than first generation Mesh units, other things being equal, but less than Third Generation Muli-Radio Mesh. Still, this is may only be true within a particular product line from a particular vendor.

Disadvantages of Second Generation Mesh

The second generation dual radio and 1+1 Mesh units still suffer from many of the first generation problems, though not so excruciatingly as the first generation Mesh.

- Internal interference still a problem

  Now the clients are gone, but every Mesh unit must still wait for every other Mesh unit in range. All of the units are still on the same Channel, 802.11a channel 149 in Figure 6. When a pair of Mesh units are broadcasting, see Figure 6 A and B, every other unit in range must wait. In figure 6 it should be noted that in a real Mesh the units at C and D would have something to do, as they are not in range of any the transmitting radio.

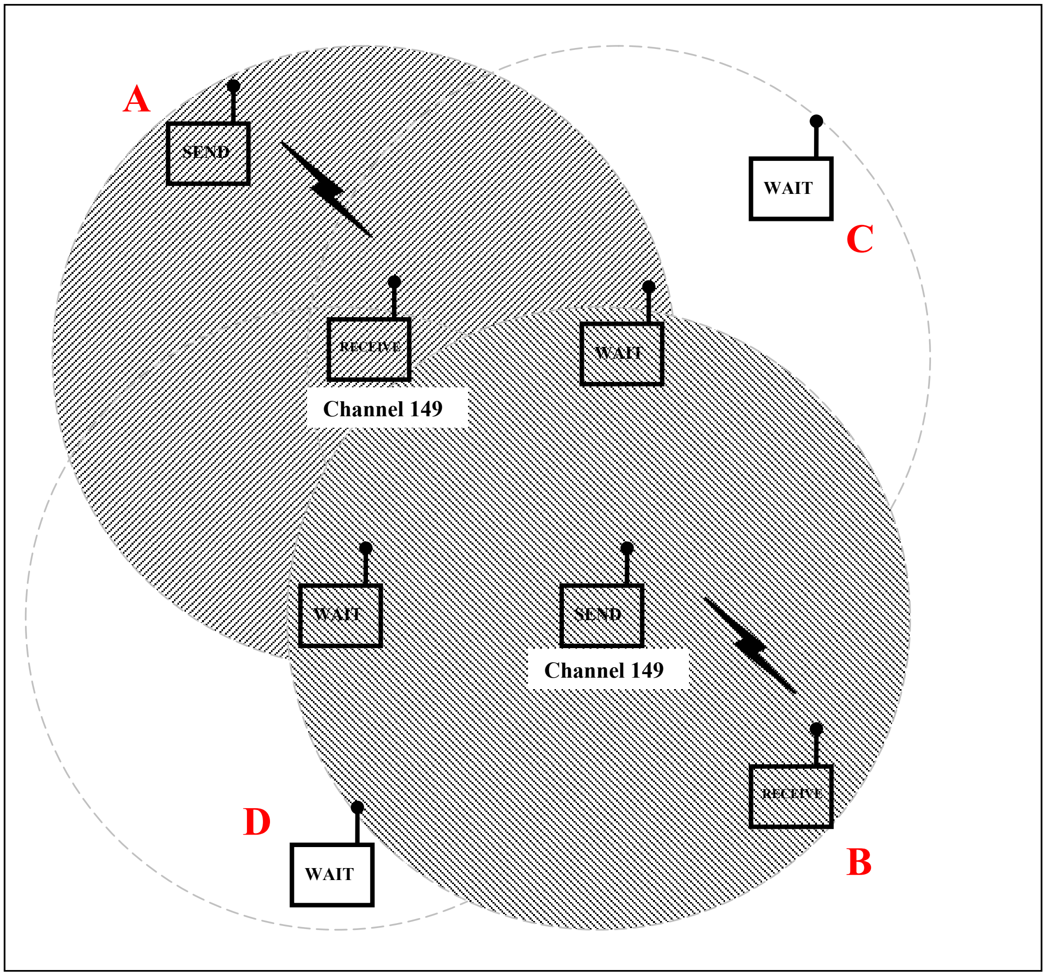

Figure 6: Backhaul Radio still must wait for every other radio in range

- Multi-hop throughput still a problem

As seen above, radios are still waiting to send in the second generation Mesh. The horrible crowding caused by clients has been resolved, but every Mesh link is still on the same channel. Still looking at Figure 6, we can see that in order to get traffic from A to B, the central pair must send the traffic while the other radios wait. The throughput from A to B is probably about 10Mbps, since the packet must wait every other cycle. However, if traffic needs to be simultaneously transmitted from C to D, which is extremely likely, then the packet must wait three cycles, reducing the throughput to about 5Mbps. Clearly in a case where many radios are in range of many other radios over multiple hops, the end-end throughput becomes very low.

Figure 6 also illustrates one of problems with the advanced auto-discovery feature on some first and second generation Mesh products. If every node in the Mesh is going out and automatically linking to as many other nodes as possible, this multi-hop problem is greatly exacerbated. In fact, the best practice is to tune the radios and Mesh links to minimize unnecessary interference while maintain maximum connection speeds.

- External interference still difficult to address

  The interference problem illustrated in Figure 4 is not fundamentally changed by removing the clients from the picture. The whole network must still use a single channel, which may suffer interference somewhere in the network. However, the ability to use 802.11a, which has more possible channels, or another frequency like 4.9GHz (for some applications only) makes it more likely that a single channel might work system wide. As the Wireless Mesh Network scales to cover a large geography, irresolvable interference is still likely to appear.

- Dual Radio antenna placement problems

  This problem is the main reason for the creation of the 1+1 Wireless Mesh products. The Dual Radio Second Generation Mesh products are potentially an improvement over the first generation. Most Dual Radio Mesh Units are deployed with just a second omni-direction antenna connected to the base unit, and so are no better than the first generation Mesh units. From a Dual Radio an antenna cable can be run to another, more appropriate location. However, low loss antenna cable is very expensive, and high loss cable would cut into the coverage area for clients. By contrast, the 1+1 Mesh units allow the client Access Point (AP) to be located anywhere within 100m using inexpensive Cat5 Ethernet cable.

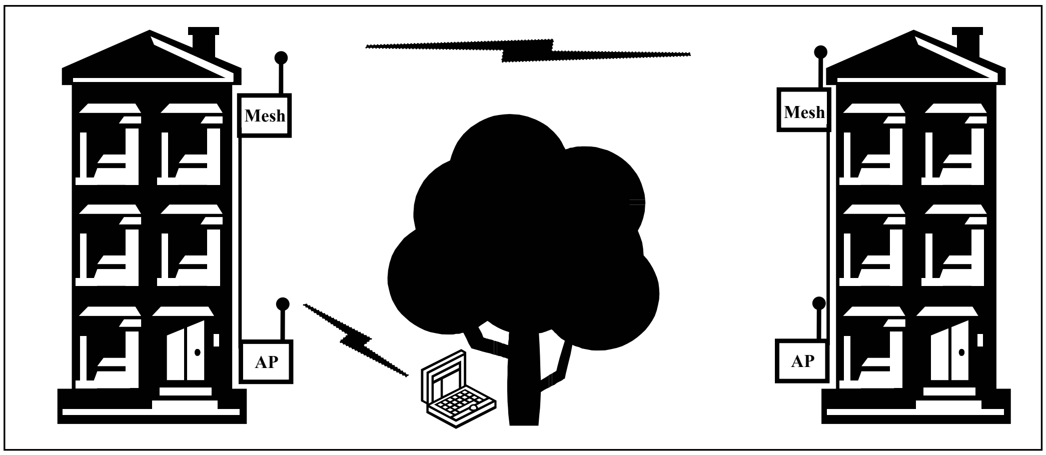

Figure 7: 1+1 Mesh solves antenna placement for clients and backhaul

- "Hidden Node" problems

  Because their broadcast range is so short and their numbers to great, the clients were by far the worst culprits in causing hidden node problems. The second generation fixes that client problem, but hidden node situations are still likely to occur. In Figure 8, the Mesh units at A and B can not hear the broadcast from the sending radios, so they may send their own overlapping radio traffic and cause an error. Unlike the case with dozens or hundreds of clients, in most cases this should be a manageable number of errors. As the network scales, and the more links each Mesh node maintains, the more traffic will be lost to this problem.

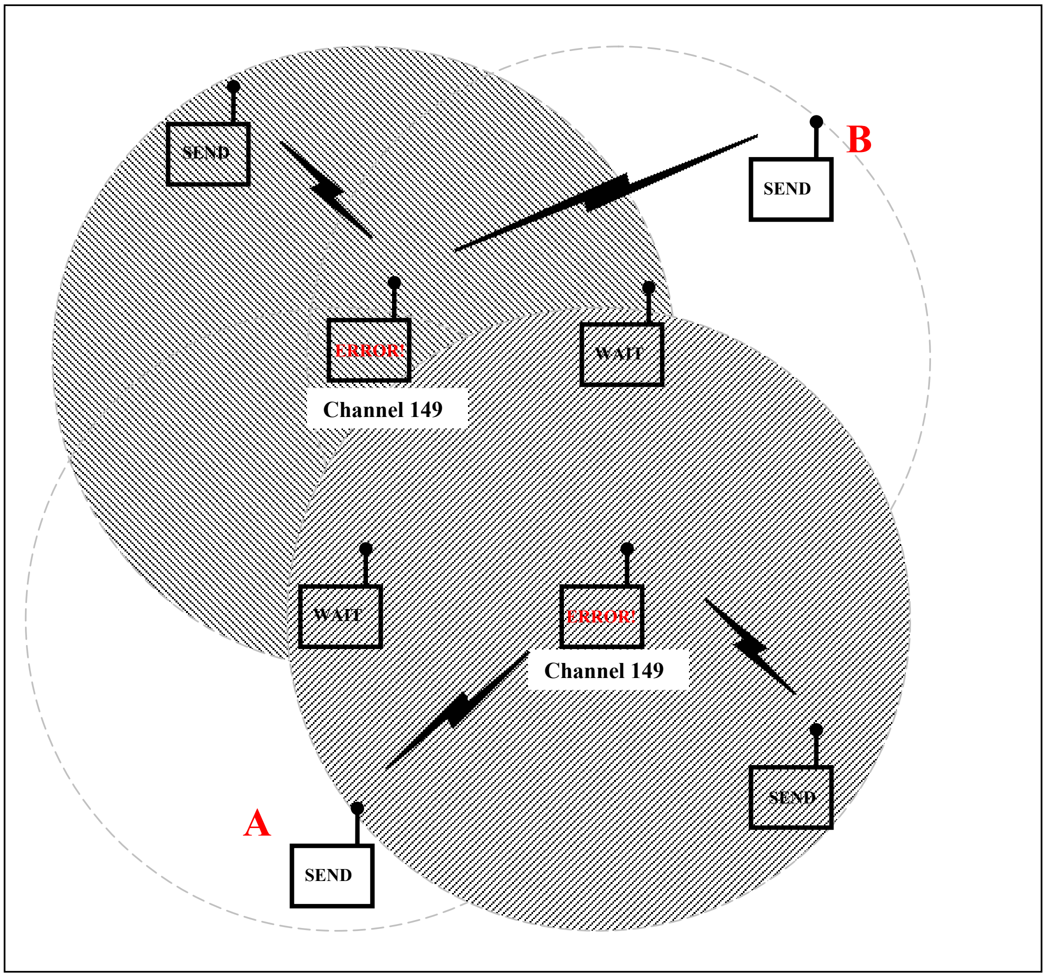

Figure 8: Hidden Node Problems in Second Generation Mesh

Conclusion

Despite the attempts to address the problems of the First Generation Wireless Mesh, the second generation Mesh still does not provide us with a scalable high performance wireless Mesh network. It is hard to recommend the first generation Mesh design for any application, but the second generation design seems perfectly well suited for a small Mesh network that does not need to scale beyond two hops. Beyond two hops throughput for the network is going to be degraded, particularly if all of the Mesh units are operating and maximum power and attempting to make connection to every other Mesh unit in range.

Multi-Radio Mesh (Third Generation)

The third generation of Wireless Mesh Networking adds two or more radios per node dedicated to the backhaul only.

Having beaten the client problems in the second generation, wireless Mesh manufacturers still came up against the fundamental problems of having every radio on the same channel. This led to the third generation of Mesh hardware. Third generation Mesh units provide two to four radios at each Mesh node. This allows each part of the network to run on a separate channel, even combinations of 802.11a, 802.11g, and 900MHz, or other. This resolved the last of the Mesh product specific problems in the second generation, so that only inherent Wi-Fi problems like hidden node still remain. Similar to the second generation mesh, these units come in a Tri or Quad Radio version, with the client radio in the same enclosure, or a “1+x” configuration with a separate Access Point.

Remaining Second Generation problems solved by Third Generation Mesh

- Internal interference is minimized

  In an ideal configuration, each link in the Mesh has a dedicated radio on a channel with no interference. In this configuration every link in the Mesh functions at maximum efficiency, and throughput and latency are ideal. See Figure 9 for an example of third generation Mesh that covers a wide area without any contention for channels in the backhaul. In some cases it may be necessary or desirable to still share links, perhaps creating two links per radio. Sharing links in this way reduces overall throughput, but not to the catastrophic levels seen in the first and second generation Meshes. Another great feature of the multi-radio Mesh units is that when you combine them with the second generation 1+1 units, the second generation problems go away because they can still be on different channels. The second generation units must be limited to the edges of the Mesh to keep this benefit, however.

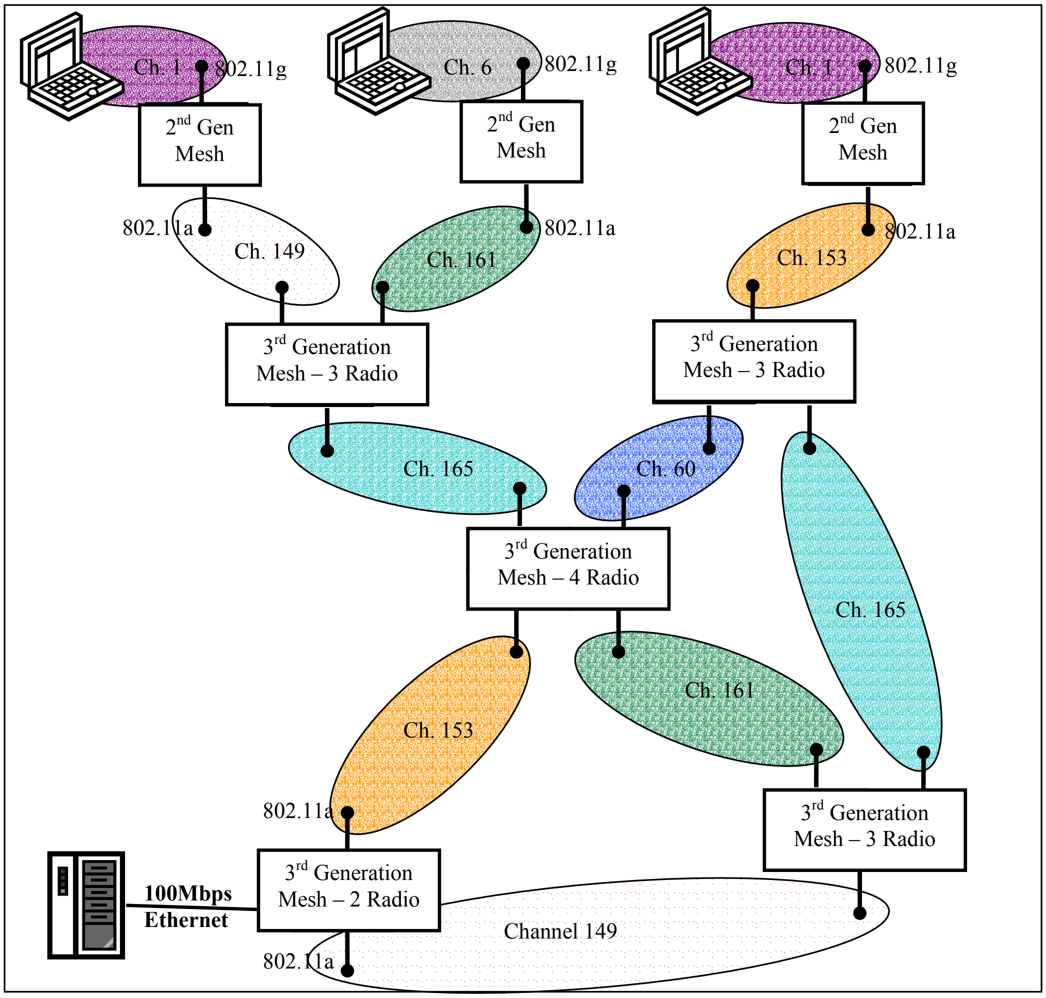

Figure 9: Third Generation Mesh Channel Use

- Multi-hop throughput is high

As seen in figure 10, it is possible to maintain full speed end-to-end in a third generation Mesh. Traffic from A to B is delivered at the full 54Mbps, almost as if they had a direct link. The throughput each hop is reduced by a small amount due to administrative overhead and radio cycling time between send and receive modes. Most importantly, no radio is waiting on this route. Up to seven hops in a third generation Wireless Mesh Network generally provide enough bandwidth for video and low enough latency for voice over Wi-Fi. You will notice at C a link that is not doing anything. This link will become active if the Mesh determines there is traffic that needs it, or if there was a failure in another Mesh node. D illustrates using a shared link, in this case replacing a three radio unit with a two radio unit, for some cost savings, while only impacting some edge units with client access.

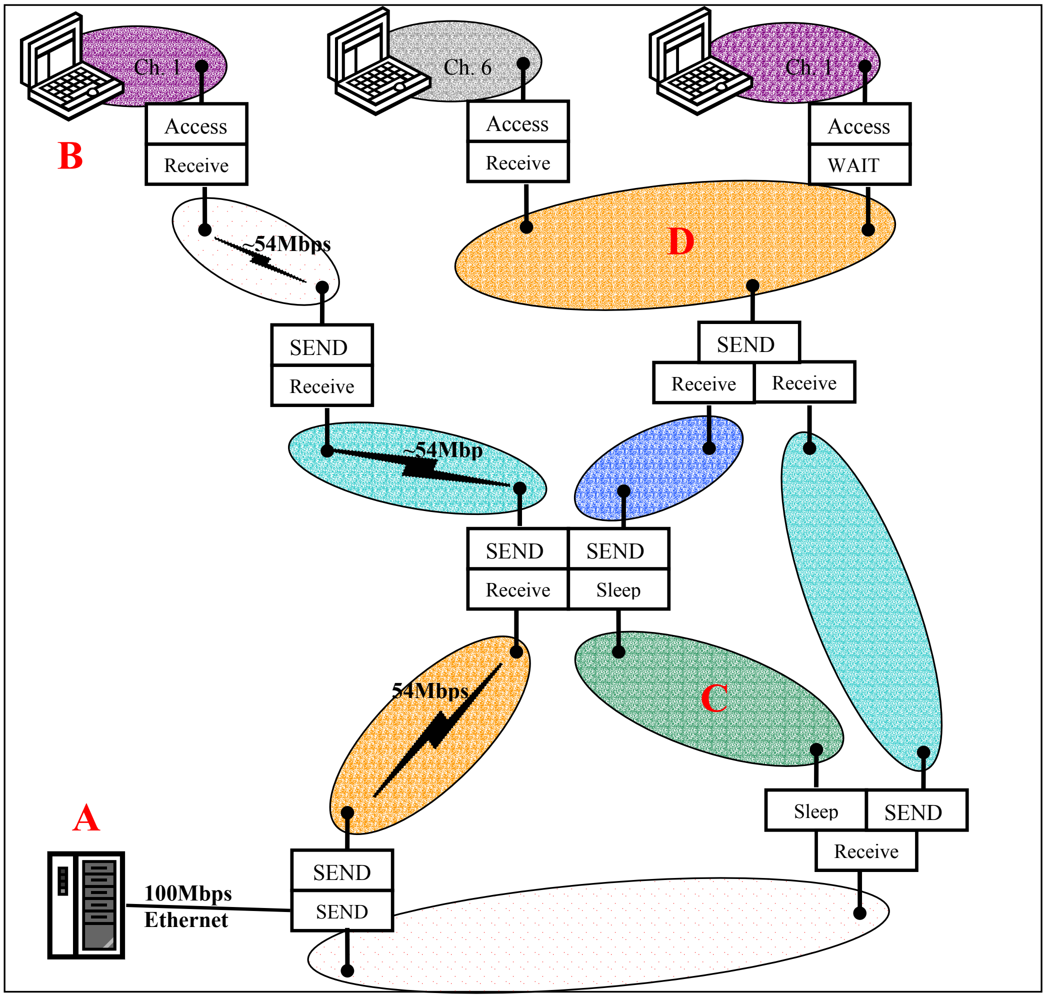

Figure 10: Multi-hop throughput in Third Generation Mesh

- External interference is manageable

The complete flexibility of radio and channel selection makes it easy to adapt to any localized interference. A separate challenge is still figuring out what the sources of interference are, and predicting transitory signals. 802.11g channels for client access will still generally be subject to interference, due to the limited channels, but this is just the nature of 802.11g.

Remaining Disadvantages of Third Generation Mesh

Multi-radio Mesh works great. Nothing is perfect, though, so here are some persistent issues.

- Dual Radio antenna placement problems

  Third generation Mesh units will come in a Tri/Quad Radio configuration or a "1+x" configuration depending on whether the client access is built onto the board or in a separate enclosure. The question of client access and antenna location is the same as the second generation Mesh. See figures 3a, 3b, and 7.

- Possible "Hidden Node" problems in shared links

  When the Wireless Mesh Network consists of entirely of independent two radio links there are no hidden node issues in the backhaul. Sharing backhaul links for cost efficiency, such as in Figure 10 D, may introduce this problem again. This does not have to be a strong reason against sharing links when appropriate in terms of cost and desired service levels, just make sure you consider this issue in your network.

- Cost

  In the third generation of Mesh, the hardware itself has become very complex. The cost of all that high-power radio hardware must be added to the cost of the powerful computation required for calculating Mesh topology for two to four separate links. Whatever Mesh manufacturer you choose, the multi-radio Mesh units will be among their most costly wireless hardware. Additional costs may be associated with high power radios to create longer backhaul links and Wireless Mesh Networks utilizing bands other than 2.4 or 5.8GHz.

Conclusions

The dream of Wireless Mesh Networking from DARPA was to create a network over radio waves that would be as good as good old Ethernet. After a bumpy start with the first generation, and persistent design limitations in the second generation, the third generation delivers. The third generation Mesh network is reliable and provides excellent throughput across long distances. Properly designed, there should be approximately 20Mbps available end to end. This falls short of Ethernet's 90Mbps, but the cost savings of not having to run hundreds of miles of cable are well worth it.

When you go to build your network make sure you consider the pro and cons of the second generation and third generation Mesh hardware. We really do not recommend the shared Client/Backhaul first generation Mesh. There are just not that many applications where first generation Mesh can succeed unless you kick off all the clients. You can further reduce overall cost by avoiding Mesh hardware that has proprietary features that you do not necessarily need, or that may in fact reduce the overall performance of you network, like "zero configuration".

Utilizing "1+x" hardware in the core of your wireless Mesh network and "1+1" units at the edges, and sharing links judiciously, you will be able to create reliable, high performance, and cost efficient network that meets the needs of your customers.