

Spread Spectrum Power Control

Darryl Smith, BE, VK2TDS
Chief Engineer, Radioactive Networks.

Introduction

The FCC when they were drafting the Part 97 rules for Spread Spectrum felt that power control was something that should be included in operations of higher power links¹. One presumes that they felt that at lower powers the added complexity of power control was not needed given the low probability of interference.

When they were writing these rules they could never have imagined the number of spread spectrum devices that would potentially be sharing the bandwidth just a few years later.

The purpose of this paper is to encourage the use of automatic power control in spread spectrum operations for all users, even those operating under Part 15.

Structure

This paper will first look at point to point links, analyzing a possible power control scheme for these links before expanding these ideas to Point to Multi-Point links and MESH networks.

In each case we go back to basic principles, modeling the point to point link, and then extending the complexity of the model.

Spread Spectrum Rules under FCC Part 97

(d) The transmitter power must not exceed 100 W under any circumstances.

If more than 1 W is used, automatic transmitter control shall limit output power to that which is required for the communication. This shall be determined by the use of the ratio, measured at the receiver, of the received energy per user data bit (E_b) to the sum of the received power spectral densities of noise (N_0) and co-channel interference (I_0).

Average transmitter power over 1 W shall be automatically adjusted to maintain an $E_b/(N_0+I_0)$ ratio of no more than 23 dB at the intended receiver.

At this point I should make a few comments on these rules. The first is that I support the 'Automatic Power Control' requirement for all Spread Spectrum operations operating at powers above 1 W. I believe that the 100 W limit is artificial, and should not be in place, but I fully support the imposition of the 1 W limit, and would oppose raising it.

¹ This was actually suggested by the ARRL Future Services Committee. Later, many of those same members, including Phil Karn and TAPR, asked that the automatic power control provisions be removed

It should also be pointed out what the definition of 'Spread Spectrum' is according to the FCC Part 97 rules. Under the Part 97 rules the definition is any mode that requires significantly more bandwidth than the data rate being transferred. The common measure is a bandwidth of at least ten times the bit rate.

For 802.11b running at high speed, the bandwidth is 22 MHz, and the bit rate is 11 mBit/Sec. The only reason that 802.11b must operate under the Spread Spectrum rules is that when operating at low speeds the same 22 MHz is used for a 1 or 2 mBit/sec data rate. In addition for compatibility, when operating at the higher speeds, RTS and CTS packets are transmitted at the lowest possible speed, meaning that 802.11b MUST be operated under the Spread Spectrum rules.

Newer OFDM modes such as 802.11g operating natively would not be classed as Spread Spectrum since it is just a wide band mode. What must be taken into account is that many 802.11g units also send data using Spread Spectrum for compatibility.

OFDM is the equivalent of taking a lot of 1200 bps modems and bolting them together to create a workable system. Thanks to software this is not too difficult.

OFDM does not achieve a GAIN from the method of modulation and demodulation. It achieves it's gain thanks to channel coding, allowing any errors caused by narrow band interference to be limited to a specific modem or set of modems and therefore have the data recovered.

Packet Headers have better BER under 802.11b

One of the subtle additions to the 802.11b standard was that although the payload of a packet could be transmitted at any of a number of speeds, the header must be transmitted at a fixed low speed.

On the surface of it, it would appear that the reason for this was simply to get away from a Qualcomm-type situation found in their IS-95 CDMA phones. In these phones, all packets are decoded at every possible speed, with the system using the packet with the lowest number of parity errors.

But there are other advantages. Assuming that the power is fixed, data sent at a lower speed will be more likely to arrive correctly than data sent at the higher speed. That means that the header is effectively being transmitted at 11 dB higher power than payload of a packet, when operating at 11 MBPS.

This allows the power on a link to be tightly controlled to what is needed for a stable 11 Mbps signal, whilst maintaining a link margin for important protocol information, such as channel access requests

Power Control through RTS/CTS

Right now the question that the reader should be asking is why I have brought up the subject of the packet headers so early in this paper. The reason is that many of the ideas assume that there is some way of determining the power required for a link, whilst maintaining viable communications.

The RTS/CTS exchange operates at a power level of about 11 dB higher than other packets when operating at maximum speed. This enables a higher reliability of the control signals. This 11 dB can be used as a link margin. The idea would be to use the previous power transmitted power level as a seed for setting the transmission level for CTS/RTS packets.

Once these packets are received by their destinations, an analysis of the link margin can be performed to analyze if the power level was correct, or needs to be adjusted when the packet payloads are actually sent.

Power Requirements

There are two types of requirements for power, and enough energy must be produced from the transmitter for the particular case. They are

- Transmission power required for error free reception
- Transmission power required for carrier sensing

A viable implementation of 802.11b does not rely on the use of carrier detect thanks to the use of the CTS/RTS scheme to control transmission of packets.

Where a carrier detect is used, the amount of power required for error free reception is significantly greater than that required for carrier sensing. In many cases in Point to Multi-Point and Mesh networks, the transmission power for most of the stations within range is that which is needed for carrier sensing.

Put another way, when receiving a signal only to detect if the channel is available, then the error rate on the received signal can be much higher.

Power Control on a Point to Point link

The simplest case on power control is a point to point link. On a point to point link there are two cases to be analyzed

- Symmetric forward and reverse path link budget
- Asymmetric forward and reverse path link budget

Symmetric Point to Point Link

A Symmetric Point to Point Link is the simplest case of power control. Any receive signal strength measurements have a direct translation to the transmit power required for the reverse path.

Put simply, if the received signal strength is 3 dB higher than is needed for a given Bit Error Rate, the transmit power on the receiving station can be reduced by that by 3 dB.




Figure 1 - Single Ended Power Control

Such a scheme does cause the link to become asymmetric, with the end with the power control having a lower power output. The biggest advantage to this would be in the situation where one end of a link is on the top of a hill, and it makes sense to reduce the coverage of the station.

In general, it is more useful to distribute the power reductions between both ends of a link.




Figure 2 - Double ended Power Control - Open Loop

Simply duplicating the first scenario for both ends of the link as shown above will not work. For one reason or another, there will always be some type of imbalance, and the link will become totally asymmetric, with one end being fully attenuated and the other end fully powered, with the likelihood that the power distribution will oscillate over time.




Figure 3 - Double Ended Power Control - Closed Loop

One way to correct this requires some coordination between both ends of the 802.11 link to exchange information on the received signal strengths. Each end of the link would receive the signal strengths and reduce the power output by the excess receive signal. In order to maintain a stable link some type of filtering should be used to average the received power over a number of packets.

What happens is that a feedback loop is created generating a system of closed loop power control. Some type of dead zone is probably useful between the received power and the attenuation. How much dead zone is needed is dependant on the link.

Asymmetric Point to Point Link

An Asymmetric Point to Point link is a symmetric link where transmit or receive amplifiers are used. A scheme of adjusting transmit power based on the receive power will not work in this situation, since it will be unstable.

There are two possible solutions to an asymmetric link

- Make the link into a Symmetric link in software with an offset in the power calculations [Figure 1 – Single Ended Power Control]
- Use the symmetric link feedback scheme. [Figure 3 – Double Ended Power Control – Closed Loop]

Either of these solutions could be used.

Point to Multipoint link

A point to multi-point link can be thought of as being a specific case of a large number of point to point links, at least superficially. There are actually a number of components to a Point to Multipoint link that need to be separately power controlled.

- Multiple Hub to point connections operating at different power levels
- Broadcasts from the hub notifying other clients that they may (or may not) transmit on the frequency
- Beacons indicating the presence of a base station

In an optimal system, each of these components needs to operate on different transmit power level schemes.

Multiple hub to point links

Where a multipoint link exists, no stations connecting to a hub directly communicate between themselves, with the possible exception for carrier sense for collision avoidance.

We can therefore consider each link between hub and each station independently for power control. That is for each station to hub transmission is made with the correct power for the path from that station to the hub. Each transmission from a hub directed at a station will be made with the correct power to communicate with that station.

The hub therefore needs to create a database of the power levels to use for each station when transmitting to it, and to adjust the power on a per transmission basis. Depending on the power control hardware this may not be as easy as it sounds.

Channel Available Broadcasts

In some Spread Spectrum systems, stations are informed of when they are permitted to transmit with the use of a short broadcast. Where the presence of the broadcast indicates that the station may transmit, this broadcast should be made at a power level appropriate to the station being communicated with.

Conversely, where the broadcast is used to inform stations that they may not transmit since a particular station is being communicated with, the broadcast must be made with a power output in line with the maximum of the transmission powers from the hub to stations.

Hub Available broadcasts

Broadcasts notifying the availability of a hub are a strange case since they need to transmit enough power to advertise the availability of the hub without transmitting too much power all the time.

Several schemes are possible. One is to assume that all the stations that need to hear the broadcast are within the area occupied by the current users of the hub. This would allow the power output of the maximum required transmission power station to be used.

But this is probably not the best scheme to use. A better solution may be to use a scheme where each broadcast is at a different power level – with 90% being at this maximum value, and the remaining 10% being at the maximum transmission power.


In order to reduce the chance of interference, this 10% should probably be sequenced with some type of PN code distribution.

Mesh

Mesh networks are a problem because of their inherent lack of structure. A properly formed mesh network is constantly adapting to a changing network topology whilst maintaining stable routing.

Up until now in this paper we have been ignoring the question of routing totally. In mesh networks, routing goes hand in hand with power control.

An example of a simple mesh network is shown here. This same mesh network will be used later as an example of how to route packets through the network.



Hops Vs. Power.

In a traditional network the path that a packet takes is governed by rules that are set by the system designer, or the protocols use to maintain a stable network. Rules for packet routing might be

- Highest speed link
- Cheapest link
- Lowest latency link
- Highest available bandwidth

Once we start talking about a mesh network with power control we need to start thinking about other considerations. Some considerations might include

- Shortest intermediate hops
- Shortest intermediate hops through high density areas
- Longer hops bypassing high density areas

If you think carefully about these choices you will realize that the routing in these cases are not normally what you would consider for the best routing, since traditional routing has only been concerned about getting our packets through

Mesh power limited routing is also concerned about getting other packets through the network, and reducing the interference to other services using the same band. Creating a stable routing system whilst reducing the interference is a tradeoff and requires a good understanding of the underlying network routing protocol.

Short Links Vs. Long Links

Traditional wisdom states that a small number of long links is better than a large number of short links, once the bit error rates and latencies are sufficiently low. However once power control comes into the equation different decisions need to be made.


Let us assume for the moment that we have the following link, where the path is 6KM long, and 36W is required to get the data through this link.



36W is a lot of power needed, so we decide to use a repeater station in the middle of the link.




By halving the length of the path, the required power has decreased to 9W per link. This adds up to 18W, or half the total power. But this is not the end of the example. We can shorten the link again, this time creating 2KM links.



With 2KM links, we only need 4W per link, making the total transmitter power for the entire link 12W. Ignoring the power consumption of the equipment itself, the power used by the link has decreased from 36W to 12W. More importantly the peak power has decreased from 36W to 4W significantly reducing interference to other users and services.

Analysis of Power Outputs and Density

The graph below shows the relative field strength as a function of position as we move from one end of a spread spectrum link to the other end. As you can see, the peak power required where no repeaters are needed is significantly higher than where one or two repeaters are used.




What this graph is now showing is where all that energy is being dissipated. The excess energy is being transferred into an area that is significantly larger than the area that is required for error-free communications using shorter links.

Coverage Areas

The coverage area of the transmissions is also quite interesting. Lets assume that we are using isotropic radiators – antennas that radiate in all directions. In our case we are ignoring any energy that is directed anywhere other than parallel to the ground to make the visualization easier.

Of course we are also dealing with bi-directional devices, which increases the total area of the RF signal. Since we are attempting to determine the total area where the signals are present, we need to look at three regions – the exclusive region at both ends of the link, and the area in the middle where both signals are present.




The coverage area of a single link in the above mesh would be

$$Area = \pi r^2 + \pi r^2 - \left(\frac{4\pi - 3\sqrt{3}}{6} \right) r^2$$

Where $r = 6$, Area = 181.97


Lets compare this to the case where we have a repeater. In this case we have three circles overlapping, each with a radius of 3.



$$Area = \pi r^2 + \pi r^2 + \pi r^2 - 2 \left(\frac{4\pi - 3\sqrt{3}}{6} \right) r^2$$

Where $r = 3$, Area = 45.49

Now we have two repeaters



$$Area = \pi r^2 + \pi r^2 + \pi r^2 + \pi r^2 - 3 \left(\frac{4\pi - 3\sqrt{3}}{6} \right) r^2$$

Where $r = 2$, Area = 20.22

Spanning Trees, Mesh Routing and Minimum Power Networks

Routing in a non-deterministic network such as a Mesh is commonly performed with some type of Spanning Tree or “Shortest Path” protocol to build a map of the network. Whilst this map does not always show the best path through the network it does show a path.

Normally in a mesh network, the routing protocol is concerned at making sure that the path does exist. But there are times that we want more than that out of a network. In order to improve the performance for all users of the frequency we need to minimize the footprint of the entire network.

There are several ways to do this

- use many shorter links rather than one longer link
- Where multiple paths exist between two points, route packets using all paths
- Use only links with active power control


All these need to be taken into consideration when it comes to determining the metrics for use in the routing protocol.

Logical Vs Physical Signal Density Mapping

If we are attempting to minimize the power spectral density (PSD) of a network, we need to look at the network as a whole, analyzing the losses inherent in every path. Later in this paper I will describe some of the metrics to be used in choosing between links.

It is possible to estimate the path lengths and therefore the physical network topology by examining each station, and analyzing the amount of power required to achieve a given signal to noise ratio at the receiving station. This data could then be analyzed using a Least Squares iterative process to generate a 2D map showing a probably layout of stations, and viable signal paths. This would be a logical diagram of the network, as close to the physical layout that we can generate. To generate a physical map we would need to have the positions of all stations

In most routing systems, a path is assigned for all packets since the bandwidth available is more than the available traffic, and this simplifies the routing. Routing for minimal PSD is almost exactly the opposite of this – where given the choice between three paths



Path Losses between Nodes							
	A	B	C	D	E	F	G
G			25	15	40		
F				40			
E	35						
D	30		10				
C	20	40					
B							
A							

C-D cause mutual interference when transmitting.

In the simplest line of site case, there is a direct correlation between path losses and distance. Once non-uniform and non-unity antenna gains are taken into account along with losses caused when links are not line of site, the actual mapping becomes more an approximation than an accurate map.

Whilst this paper has been in production a lot of work has been done with direction finding in 802.11 networks. It is now possible to even produce a fairly accurate physical map of the network links.

Example of why Minimal Power Spectral Density routing is essential and difficult.

Assume in our network that we need to get packets from A to G. There are obviously three simple paths that can be taken, with some more complex paths possible too. An important component of the C-D link is that when one of these stations transmits, the other station is unable to hear any other stations.

If this station is transmitting	Then the following communications are lost
D	A to C B to C G to C
C	A to D F to D G to D

In order to facilitate communications to stations C and D we need to ensure that packets are not routed into through those nodes unless it is necessary. In this example, the path from A to G via node E provides a path that will not cause interference to other stations. Of course there will be contention for the transmission to node G, but this can be dealt with through the normal distributed coordination functions.

The question becomes how we would know that C and D cause interference with each other, in an automated environment. The answer is quite simple – Stations pairs that have very low path losses WILL experience the Near-Far problem we have just described, since they are so close. When one station is transmitting the other cannot be receiving from another station.

How to choose between links

In order to choose between links we need to analyze what our desired result is. We need to minimize the power density at all points in space whilst maximizing data transfer.

Each node is responsible for routing to adjacent nodes, and determining the best adjacent node for a packet to transit through.

Routing Assumptions

It is assumed that where the links are dynamic that there is no way to determine the correct path from any logical characteristic of the nodes, such as IP address, node names and the like.

There may be cases where static links are used, but these will tend to be permanent backbone nodes. Wherever possible static paths SHOULD NOT be used since they affect the routing for low PSD.

Metric	Description
Minimal Adjacent Connected Links	The number of links directly connected to either end of a link should be minimized. The traffic that the links generate can be accounted for elsewhere, but the fact that there are attached links is important when routing for minimal PSD
Minimal Link Length	Shorter link lengths are generally better than longer links in the since they tend to cover a lower area.
Low Path Loss	A link with a lower path loss is better than a link with a high path loss since this allows lower powers to be used.
Highly Directional Antennas	Highly directional antennas allow for the signal footprint to be tightly controlled. Reducing the gain on the system improves the performance of the total system too.
Active Power Control	Active Power Control reduces the power levels to what is needed for the communications, reducing the interference to other users.
Lower Interference	Some links will experience higher interference levels from any number of sources. These links should not be used as often since higher power levels are required for reliable communications. In addition there is the likelihood of interfering with the source of the interference.
More effective Antenna Systems	Links with more effective antenna systems and radio equipment are preferable to links with poor antennas and antenna systems with significant attenuation in the antenna system

Node Types

There are two different logical node types – Single and Multi linked nodes. Each of these nodes has different routing characteristics and requirements for routing.

The Single Linked Node has very simple requirements for routing, due to the fact that it can only route packets to one other station. Therefore in this situation, no additional routing software is required since there is only one path.

The complication is where the single node is also connected to the internet, or to an external network of some type. The node then needs to make a decision on some basis on the routing of these packets, such as by IP address.

Just because a node has a single link does not mean that power control should not be used. There are significant advantages to the entire network when all links use power control, as we have seen earlier.

PN Randomization of routing between multiple paths.

Where multiple paths exist with identical metrics, packets should be routed between the paths with an equal distribution. There are a number of ways to do this, with 'Round Robin' being the most common, with a fixed number of packets being routed through each path in turn.

An alternate way to do this is to route with the choice between links on the basis of a PN code. In this situation, a PN code would generate a binary value for each packet. If the value is different from the last PN value, the next link is used. In this way, the links are generally kept free of traffic for longer periods.

The problem with Round Robin in a network is that it will be likely that two devices will attempt to send data to the same third device, failing, and the next time they try the odds are that they will also transmit at exactly the same time again. Whilst this can be compensated for with some type of randomized timing, using a PN should give better overall performance.

Conclusion

In this paper I have started by looking at why power control on spread spectrum links is a good thing. From this basis, I have moved to looking at how to create a routing system based on power controlled links and reducing interference to all users. This paper is just a start, and a lot more work needs to be done in this area.

Prior Art – DCC Papers

Physical Layer Considerations in building a High Speed Network; Glenn Elmore, N6GN

MACA – A New Channel Access Method for Packet Radio; Phil Karn, KA9Q