Self-Interference

Identification ⇒ Avoidance ⇒ Mitigation

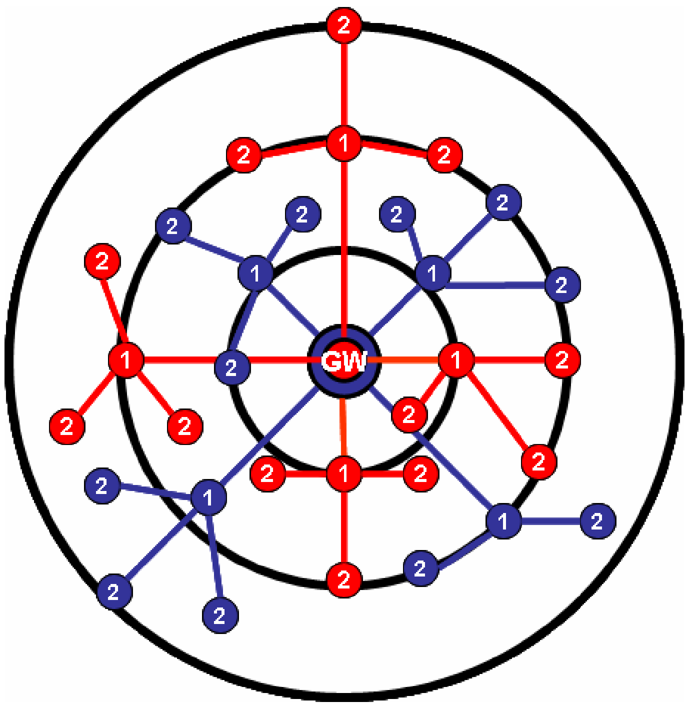

When deploying wireless networks, self-interference can become an issue if not given proper attention. This whitepaper discusses the concept of self-interference, how to identify it, how to avoid it by design and how to mitigate it in existing networks through proper parenting of nodes.

Introduction

Wireless links are impacted by the presence of various types of interference while receiving data. Wireless links are impacted by the presence of various types of interference while receiving data. The impact of interference can differ depending on the type and power level of the interference encountered.

This white paper is going to focus on one type of interference that can be generated within a wireless network, and that is self-interference, or interference generated by non-optimal placement of devices within the mesh. Other whitepapers will cover the topics of multipath (a function of Fresnel Zone) and external interference issues.

When you have completed this paper you will understand the following topics:

- What is self-interference?
- How is self-interference generated?
- What is spatial diversity and how does it help with self-interference?
- Strategic placement and location of the SkyGateway
- Strategic placement and location of the first-hop SkyExtenders
- Strategic placement and location of the second-hop SkyExtenders
- Considerations for hops beyond the second-hop
- Monitoring the network for self-interference

What Is Self-Interference and how is it Created?

Self-interference is a wireless network principle in which the transmissions from one device result in an unwanted signal being received on other devices. Specifically, when one device is transmitting to second device and these transmissions are “overheard” by a third device as it attempts to receive signal from a fourth device.

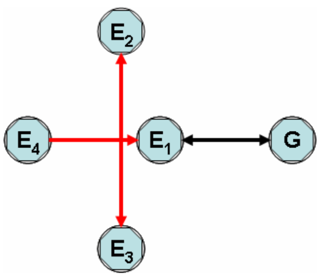

This is illustrated by the diagram, to the left. Self-interference is at its maximum, in this example, when E_4 transmits to E_1 at the same time that E_2 and E_3 are transmitting in either direction. The amount of self-interference generated at E_1 will be a function of the distance between E_2 and E_1 and E_3 and E_1. You will remember from basic RF theory, the power level decreases as a function of the square of the distance.

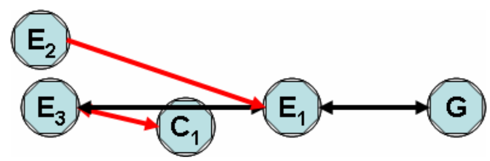

Another example of a scenario where self-interference is created is shown in the diagram to the right. In this example, C_1 is located between E_3 and E_1. Whenever E_3 transmits to C_1 at the same time that E_2 transmits to E_1 the transmission between E_2 and E_1 will be interfered with. As we proceed through this white paper rules of placement, that help reduce incidences of the scenarios shown above, will be covered in some detail. Another example of a mesh that generates a large amount of self-interference is shown below.

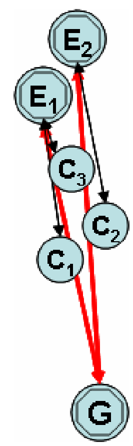

<u>**Device placement plays a critical role in the conditions that lead to self-interference**</u>. Self-interference results whenever a transmitting node interferers with the transmissions of another node in the network. The ratio between the transmitted signal power (carrier or “C”), and the interference power (designated as “I”), will determine the impact to link performance. If the C/I ratio is large enough there will be no impact to link performance. As the difference between carrier power and interference power (C/I ratio) decreases the impact to link performance will become more significant.

<u>The Carrier to Interference (C/I) ratio determines the level of impact that interference will have on the link</u>. The greater the power level of the received signal is, with respect to the interfering signal, the less impact the interference will have on the received signal.

<u>The chance of interference occurring is a probabilistic function</u>. If the interfering node, and the node(s) being interfered with, are inactive most of the time (low duty-cycle) then self-interference will rarely occur, though the potential for it might be high, based on device location and power levels. Conversely, if the interfering node, and the node(s) being interfered with, are very active nodes (high duty cycle) then self-interference will occur regularly and the impact to link quality, as evidenced by modulation rate, may be significant, depending on the power level, or RSSI, of the interfering signal.

The Near / Far Challenge

Another element that contributes to self-interference is one that is, in many wireless deployments, the one thought about the least. The author of this paper has been guilty, many times, of missing this issue, too. This element of deployment is called the near/far effect and it largely inhabits fixed wireless broadband installations, but can also impact dense mesh deployments if not carefully watched for.

The near/far effect arises whenever a strong signal, from a transmitting node that's very close to the receiver, interferes with a much weaker signal from the actual transmitter that's further away. The near signal (interferer) is stronger than the far signal making it impossible for the receiver to actually detect/decode the signal from the further transmitter. The near/far effect manifests itself whenever child nodes are not equidistant from their parent and both the RSSI and modulation rates vary as a function of the dissimilar distances between nodes.

The near/far effect can also occur whenever a group of child nodes are in very close proximity to one another (i.e. multiple SkyExtenders on a single rooftop to "reach" various locations within the network) and their child nodes (either SkyConnectors or SkyExtenders) are located at a distance significantly greater than the distance between the co-located SkyExtenders.

NOTE: This only applies when all SkyExtenders are configured to the same frequency and are on the same network.

The near/far effect is presented in the diagram to the left. The lack of equidistance manifests itself in poorly balanced power ratios. For example, the receive signal strength (RSSI), at the SkyGateway, for SkyExtender E_3 will be significantly lower in value than the interference generated, at the SkyGateway when E_4 transmits to C_4, C_2 transmits to E_2 or C_1 transmits to E_1.

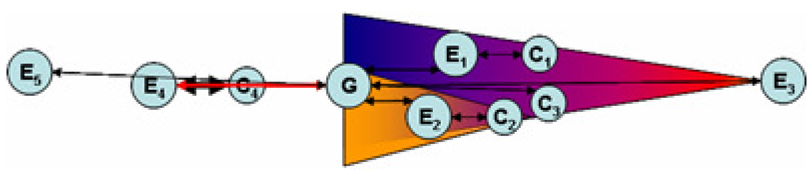

The same near/far conditions exist whenever E_3 transmits to the SkyGateway at the same time C_1 transmits to E_1, or C_2 transmits to E_2. The near/far effect may even manifest itself when E_4 transmits to C_4 but because of antenna rear-side rejection the impact of that interference may be somewhat less.

How Do I Tell If Self-Interference Exists?

The by-product of self-interference is a compromise in RF link integrity which results in modulation rates for the links compromised by self-interference dropping whenever simultaneous transmissions occur. These drops in modulation rates will most likely occur during the busiest times on the network – periods when modulation rates need to be the highest possible in order to accommodate the increased volume of traffic.

The System Warning Signs

There are several warning signs that can be viewed the link statistics for a given device. We will review these in some detail, here. There are two variables that to be concerned with when it comes to self-interference. The first is the Received Signal Strength Indicator (RSSI) and the second is modulation rate. Let's look at both of these and their relationship to self-interference within the mesh.

A warning sign that self-interference is impacting the network occurs during busy traffic periods on the network. During these busy periods, the link modulation rates may drop significantly - often from 36Mbps to 6Mbps or 12Mbps. If the self-interference is bad enough the links to some devices may drop altogether. This assumes that adequate RF evaluation of the various SkyExtender and SkyGateway sites has been completed and that no additional interference from non-SkyPilot equipment is present.

RSSI

The RSSI value directly translates into how much a given link will be impacted by interference – whether it is self-interference or interference from another source. Low RSSI values translate into low received signal strength., which translates into greater susceptibility to interference. RSSI values should be targeted at no less than 25. The higher the RSSI value the less susceptible to any type of interference the link will be. Subtract 95 from the RSSI value and you will obtain the approximate signal strength in dBm.

RSSI values do not drop in the presence of interference (self or otherwise) unless that interference is a direct result of the receiving antenna being in a predominant number of canceling Fresnel Zones, or is positioned well below the first Fresnel Zone. This condition is a result of multipath and will not be covered in this whitepaper as it's a function of microwave path design not self-interference.

Modulation Rate

The other value of importance is modulation rate. While RSSI is a display of receive signal power and susceptibility to interference **the modulation rate is an indicator of the integrity of the received signal**; the greater the amount of interference the lower the integrity of the received signal.

Modulation rates do not always follow RSSI values especially when self-interference is present. In many cases where self-interference is predominant the modulation rates may be low (6Mbps, 9Mbps or 12Mbps) but the RSSI values can be high (25 to 30).

In the situation where receive signal strength is high the modulation rates can be low, due to sources of interference with high values of co-channel (same channel) interference. The greater the amounts of channel overlap and interference power, the lower the resultant modulation rate.

Indicators

There are several indicators of self-interference that can be derived directly from tools at your disposal. The first and most obvious tool is the CLI. This can produce quick and effective data about the current operational characteristics of the network. Another tool that can be utilized to determine the need for and the impact of changes is the nodeInfo script (downloadable from www.skypilot.com). The nodeInfo script has the ability to generate self-interference output which identifies all potential interferers for a given node.

CLI Review

When interference is present on a link modulation rates may drop, or they may not. The reduction in the modulation rate is a function of two independent events: 1) The amount of activity, or duty cycle, on the affected link and 2) The amount of activity, or duty cycle, on the interfering link.

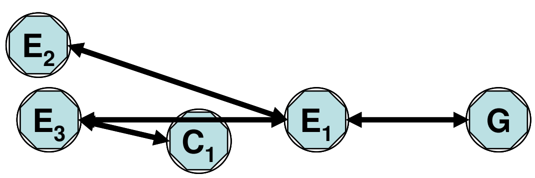

Let's look at this concept in a bit more detail. In the image, on the left, let's assume that E_1 and E_2 are very active and that E_3 and C_1 are not active at all.

The potential for self-interference to exist in this example is very high, as E_3 transmits to C_1 using the same antenna that it uses to transmit to E1. If E_3 and C_1 are not active the link between E_2 and E_1 will not be impacted. There is nothing to impact in that scenario.

Self-interference becomes an issue when E_3 transmits to C_1 at the same time that E_2 transmits to E_1. If the user connected to C_1 has a high bandwidth link and they are using that bandwidth, self-interference cause the modulation rates to drop and a possible loss of link if it becomes to prevalent.

When self-interference occurs we should be able to **log into E_1** and expect the following events to occur:

- On the link between E_2 and E_1 the **LRSSI and RRSSI will remain the same**.
- On the link between E_2 and E_1 the **LTxMod will remain the same**.
- On the link between E_2 and E_1 the **RTxMod will drop as a function of transmission activity** between E_3 and C_1 and the amount of transmission activity E_2 and E_1 (directionality important).
- Depending on the amount of interference generated by E_2 at C_1 the transmit modulation from E_3 to C_1 may also drop.

It's important, as the steps above are reviewed, to note the importance of transmission direction. Transmissions from E_1 to E_3 will not impact C_1. However, they may impact transmissions being received on E_2. The impact of the interference will be a result of the timing and power ratios between the interference and the received signal.

nodeInfo Output

nodeInfo is a utility provided by SkyPilot networks that presents both textual and visual data about self-interference within the network. We're not going to cover this in a lot of detail, in this section, but greater detail follows. The nodeInfo output provides the "receiving node" and the transmitting node (interferer) data as well as KML output to be viewed in Google Earth.

The SkyGateway – The Key to Mesh Performance

In the introductory comments about self-interference it was noted that "the chance of interference occurring is a probabilistic function." It went on to explain that the more active a potential interferer and receiver is the greater the probability that self-interference will occur.

If the amount of activity increases the potential for self-interference to occur it becomes increasingly evident that the SkyGateway is the most active node within the topology and has the greatest potential to interfere with the greatest number of active nodes on the network. If we are to mitigate the impact of self-interference within the mesh it becomes critical to locate the SkyGateway in such a manner as to substantially reduce the amount of potential self-interference it can generate. In all mesh deployments this should be a primary objective. This section will review placement rules, starting with the SkyGateway and the location(s) of its first-hop SkyExtenders.

The SkyGateway – Mesh's Most Important Node

The SkyGateway is the most important node in the mesh network. It is the source of all traffic destined for the mesh and it's the termination point for all traffic coming from the mesh. It's location within the mesh is crucial for the mesh to function holistically. Any elements of design compromised around the SkyGateway will compromise the performance of the entire mesh.

Locating the SkyGateway

This section will cover the location of the SkyGateway with respect to all other nodes within the mesh. At the risk of sounding like a broken record we're going to repeat that the node that requires the most careful attention to placement within the SkyPilot network is the SkyGateway – we can't say this enough. Compromising the location of the SkyGateway will lead to increased self-interference, multipath, reduction in system capacity or other issues.

The goal, with the installation of any wireless node is to strategically locate it in such a manner as to reduce the probability of it interfering with any other node on the network. To accomplish this objective **the most strategic location for the SkyGateway is at the center of the target service area**. Why the center?

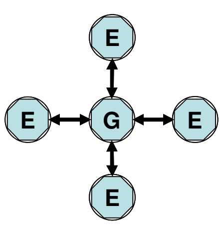

Locating the SkyGateway at the center of the coverage area allows it to efficiently use all of its sectors to communicate with links surrounding it. This reduces the potential for the SkyGateway to interfere with any given section of the mesh continually.

In addition to the SkyGateway being placed at the center of the coverage area, the first-hop nodes (typically SkyExtenders) should be placed at 90-degree angles to each other. The placement of devices around the SkyGateway is commonly referred to as **spatial diversity**.

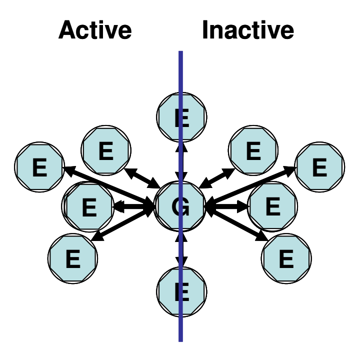

When all devices are located on one or two adjacent sectors of the SkyGateway very little physical separation between devices is realized and the probability of self-interference rises as a result of the SkyGateway continually transmitting in a single direction.

Based on projected and actual traffic rates the design "center" of the network may not always be the physical center of the network. This is an important concept to grasp and is further covered in the section on spatial diversity. The design end-goal is to **surround the SkyGateway with first-hop nodes that are active in all directions**.

What is Spatial Diversity and why is it Important?

Spatial diversity is a topic that is commonly seen in discussions centered on the elimination of self-interference within a wireless network. Self-interference is characterized by the fact that when a node transmits it has potential to interfere with the reception of other wireless nodes within that network. If the interference with other links is strong enough it will cause data loss on those links at a minimum, lowered modulation rates and in the worst case a complete loss of link. Spatial diversity can only be realized with wireless nodes that have directional antennas. Without directional transmission there can be no spatial diversity.

Some Concepts Required to Understand Spatial Diversity

In order to understand spatial diversity better a basic understanding of concepts surrounding SkyPilot protocol and node behavior within that protocol is required. More detail about these topics may be found in the Support section at www.skypilot.com. The information presented here will be a cursory overview for the purposes of explaining why spatial diversity is important and how it can be achieved.

There are two types of nodes within the SkyPilot mesh – a parent node and a child node. The SkyGateway can only function as a parent node and the SkyConnector can only function as a child node. The SkyExtender, however, may function either as a parent or a child node – depending on its current status with its parent node.

A parent node communicates with child nodes and either transmits data to the child, or allows the child to transmit data to the parent. The parent node uses a scheduling agent to determine which child nodes receive bandwidth, how much bandwidth, when the bandwidth will be allocated and in which direction that bandwidth will travel (either to, or from, the child node).

A child node must first listen for its parent. If the parent requires the attention of the child it will communicate that intention during the listening period and the child will respond accordingly – either receiving or transmitting data. If the parent node to the SkyExtender does not require the time/attention of the SkyExtender, the SkyExtender can turn itself into a parent node and communicate directly with its child nodes. Each SkyExtender in the mesh is a self-managing entity; managing bandwidth allocation between itself and all child nodes with active links to it. It performs this management without the aid/support of the parent node. This is an important concept to remember when dealing with self-interference issues – **every SkyExtender has the potential to be a parent and transmit independently of the SkyGateway, or its parent node**.

Spatial Diversity Defined

Spatial diversity is a topology design choice that allows for maximum separation of child nodes linked to any given parent node. This separation reduces the potential for self-interference as the parent will be required to transmit in different directions during each transmission opportunity.

The diagram to the right shows what would be considered perfect spatial diversity as it applies to the SkyGateway. Spatial diversity refers to the location of **active** **links** with respect to the parent node. Theoretically perfect spatial diversity has **active first-hop nodes** located at 90-degree angles to each other, surrounding the parent node. This provides maximum separation of antenna coverage areas. The SkyPilot hardware has embedded directional antennas and this arrangement promotes the most efficient use of those antennas.

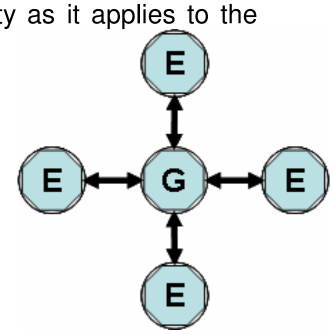

If each node around the parent node has approximately the data transfer rate the parent will spend equal time in all directions reducing the potential for self-interference in all areas of the network.

The following two illustrations show what spatial diversity is not and are both examples of poor spatial diversity. The diagram on the left shows all first-hop SkyExtenders located on one side of the SkyGateway, forcing the SkyGateway to transmit continuously in a single direction. This can lead to several issues including increased incidence of self-interference (a function of how active the SkyGateway is), increased incidence of multipath; especially standing waves and other issues within the mesh.

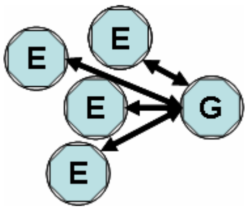

A similar scenario can be seen in the diagram to the right. This diagram indicates a seemingly “better” spatial diversity in that the first-hop nodes are located in a roughly 90-degree pattern around the SkyGateway. In this example the lack of spatial diversity is not due to node location it’s a function of node activity level. The nodes on the right are inactive as compared to the nodes on the left.

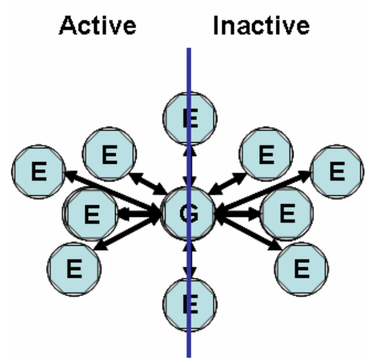

As in the previous example this leaves the SkyGateway actively transmitting on a single side leading to the same issues noted above but with slightly decreased frequency. The decrease in self-interference will be a function of the amount of active traffic to the nodes located to the SkyGateways right side.

Spatial Diversity, Multipath and Self-Interference

Lack of spatial diversity increases the chances for the development of multipath standing waves and self-interference. The more times a parent node transmits from a given antenna the greater the probability that those transmissions will interfere with other nodes in that region of the network. The more child nodes with active links to a single antenna the more time the parent node will spend transmitting to that sector further increasing the potential for self-interference. This constant transmission in a given direction also increases the potential for multipath and standing waves in that sector.

By locating **active** first-hop nodes in a spatially diverse manner around the parent node we decrease the number of transmissions from the parent node in a given direction. Spatial diversity forces the parent node to schedule transmissions to child nodes located around it.

The more a parent node transmits around itself the less chance that multipath will "build up" in a single location. This build up is typically referred to as standing waves. Standing waves develop as a function of the amount of time spent transmitting to a single region of the network; when transmission is more or less constant in a given area. Multipath and standing waves are two additional issues that support the need for spatial diversity.

Moving from the topic of multipath to self-interference, the illustration to the left depicts what occurs in the situation where the SkyGateway is located on one side of the mesh network. This condition not only applies to the location of the SkyGateway with respect to its active child nodes but it also directly applies to the location of the SkyExtender with respect to its active child nodes.

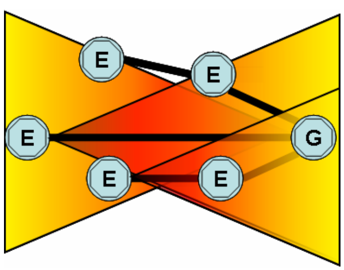

The areas with the darker red color are locations where the signals are the hottest in terms of transmission power. The signal power degrades toward the yellow areas as a function of distance between transmitting node and receiving node.

When the SkyGateway transmits to the SkyExtender in the center of the network that transmission will interfere with the transmissions between the other first-hop SkyExtenders and their child nodes. This is self-interference.

The same conditions exist when the second-hop SkyExtenders in this network transmit to their first-hop parents. They will cause self-interference with transmissions between the other first-hop SkyExtenders and the SkyGateway. Because they're using the same antenna, or one adjacent, the rejection of the interfering RF signal will be low.

The following illustration shows the same network with the SkyGateway now located at the center of the network with good spatial diversity. The original SkyGateway location has been replaced with a SkyExtender and the SkyExtenders furthest out have been parented to the first-hop in what would be, if this network were to continue development, a good example of spatial diversity.

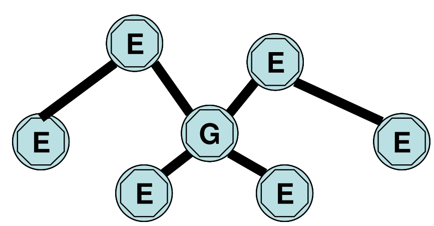

Location of First-hop SkyExtenders

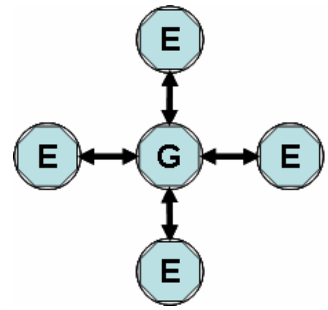

The location of the first-hop SkyExtenders with respect to their child nodes and to the SkyGateway is as crucial to network performance as the location of the SkyGateway itself. This is because the first-hop SkyExtenders are aggregation points for the entire mesh transmitting to and receiving from all nodes in their related quadrant (see spatial diversity above).

The **active** first-hop SkyExtenders should be located in a spatially diverse manner about the SkyGateway. In theory, there should be four **active** first-hop SkyExtenders, per SkyGateway deployed, and these four SkyExtenders should be located at 90-degree angles from each other, as shown in the diagram. In order to avoid issues with multipath and self-interference the first-hop SkyExtenders should not be located on one side of the SkyGateway. The spatially diverse arrangement provides the most strategic deployment method for the first-hop SkyExtenders.

Why? With this arrangement the SkyGateway will not spend any significant amount of time in any given direction. As a result the potential for self-interference in an area of the network will be reduced as a function of the amount of time that it spends transmitting in all other areas of the network. Given the same number of active subscribers the difference between the arrangement shown above and an arrangement where all first-hop SkyExtenders are linked to a single antenna on the SkyGateway would reduce self-interference in the mesh by 75%.

A Quick Note about Nodes Extremely Close to the SkyGateway

Because of the density of some mesh deployments there may be a "first-hop" ring around the SkyGateways that are closest to the SkyGateway mount location. It is HIGHLY recommended that these first-hop SkyExtenders be parented to a single SkyGateway and that these SkyExtenders have NO child nodes linked to them.

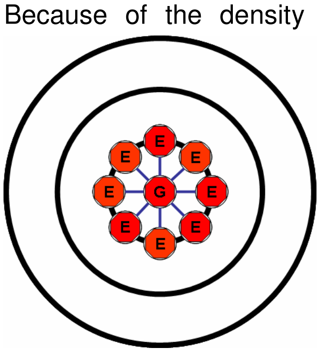

This will avoid any unnecessary self-interference occurring in the mesh. Once we're beyond this initial "first hop" ring, standard deployment practices will allow the addition of child nodes to the first-hop nodes. This only applies to SkyExtenders located extremely close to the SkyGateways. The expansion of the concept of the close-in hop and the first-hop nodes attached to other SkyGateways is illustrated in the diagram to the right.

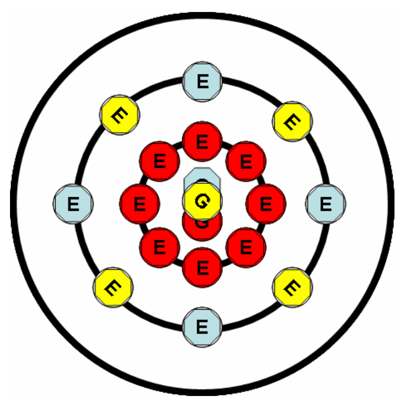

Notice the spatial diversity employed with the first-hop SkyExtenders in the blue and the yellow networks. The first-hop SkyExtenders are located at 90-degree angles to each other and their spacing in the coverage area is equidistant from the SkyGateway. This is the type of installation/deployment to shoot for when determining the placement of the first-hop nodes.

If further density is necessary additional SkyGateways can be added at the center of the mesh and additional first-hop nodes can be linked to the new SkyGateways. The new first-hop SkyExtenders will be interleaved between the already existing units.

Self-Interference and First-hop SkyExtenders

The first-hop SkyExtenders can contribute to and be impacted by self-interference from the mesh. In order to sufficiently mitigate the effects of self-interference on these links **good spatial diversity must be maintained around the SkyGateway**. In the absence of spatial diversity the SkyGateway will transmit more toward all network devices increasing the chance that self-interference will result.

First-hop SkyExtenders should be approximately equidistant from the SkyGateway. This will ensure that the link modulation rates and the link RSSI values are equal and reduce the susceptibility to interference that may be generated by the mesh network. Target the RSSI and the modulation rates as high as possible for these first-hop links as they determine system capacity.

Locating the Second-Hop SkyExtender

As mentioned in the opening paragraph, the location of the second-hop SkyExtenders is also key to the efficient operation and deployment of the mesh. This section will explore rules-of-thumb surrounding the placement of the second-hop SkyExtenders with respect to their parent node, each other and their child nodes.

In most networks there will be multiple SkyExtenders functioning as second-hop SkyExtenders. As with the SkyGateway the goal is to provide as much spatial diversity as possible between the second-hop SkyExtender and the child nodes linked to it.

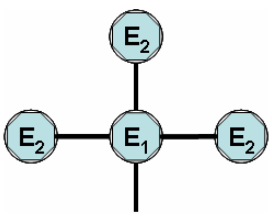

Perfect spatial diversity at this hop would look like the diagram to the left. A more real-world spatial diversity would look like the diagram to the right. This combination of parent node, and attached child nodes, is what is referred to as a topology cluster.

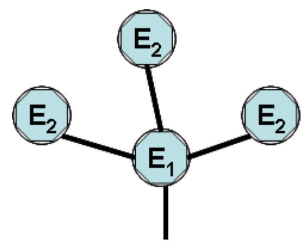

Links to Parent Nodes and Child Nodes

The discussion of spatial diversity, in the second-hop nodes is an excellent segue into the discussion of some rules surrounding the placement of all second-hop nodes. These rules, when followed, will reduce the amount of self-interference, multipath build-up and other potential issues within the mesh network. While it's realized that, in most real-world networks, you're not going to be able to achieve perfect node placement, working to get as closely as possible to these deployment guidelines will yield a highly successful municipal mesh deployment.

The first and most critical rule to follow is: **avoid forming child links on the antenna used to link to a parent node**. Forming a child link, on the antenna used to link to a parent node, will create conditions for self-interference that cannot be avoided, within the mesh.

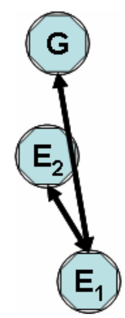

The example to the right is an illustration of this point. Second-hop SkyExtender E_2 is located between the first-hop SkyExtender E_1 and the SkyGateway. Whenever E_1 transmits to E_2, it will interfere with any transmissions being received by the SkyGateway on the antenna sector(s) facing E_1 and E_2. The closer the proximity of E_1 and E_2 to the SkyGateway the more destructive this interference will be. If there are child nodes located between E_2 and E_1 additional self-interference will result.

There are ways to rectify this situation. Let's take a look at a few of them. Assuming that RSSI levels and modulation rates between the second-hop SkyExtender and the SkyGateway were similar to the RSSI level and modulation rates of the first-hop SkyExtender to the SkyGateway the scenario could be resolved simply by parenting the current second-hop SkyExtender to a different SkyGateway, as is shown to the left.

It is important to maintain that "balance" between the RSSI and the modulation rate, on every link. If there are significant differences the best solution would be to add another SkyExtender, as shown on the right. The addition of the new SkyExtender will allow for better modulation rates and RSSI values between the first-hop SkyExtenders and the SkyGateway offering the highest resilience to interference. This also will provide better coverage density as a side-benefit.

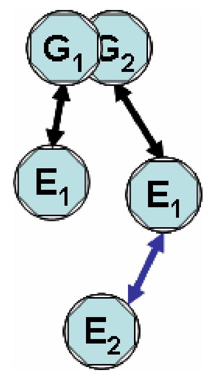

When determining positions of child nodes in the second-hop the rule-of-thumb for placement of the child nodes is **child nodes should be angled back, and away, from the parent node**. This allows the best spatial diversity to be achieved with the second-hop nodes.

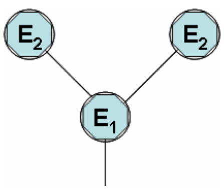

What does this look like, in practice? As you'll see from the diagram on the left the SkyExtenders in the second-hop (labeled E_2) are angled back, and away, from the parent node (E_1). In practice denser network topologies would attempt to keep the topology cluster (E_1 – the parent node, and E_2 the child nodes, limited to three nodes to reduce the amount of self-interference that could be generated within the cluster, or with other clusters on the same frequency, but located elsewhere in the mesh.

An example of how a "complete" second-hop would look is illustrated in the diagram, to the right. This arrangement provides 360-degree coverage of the area to be serviced. It also meets the requirements for spatial diversity in that each of the child nodes is angled back, and away, from the parent node.

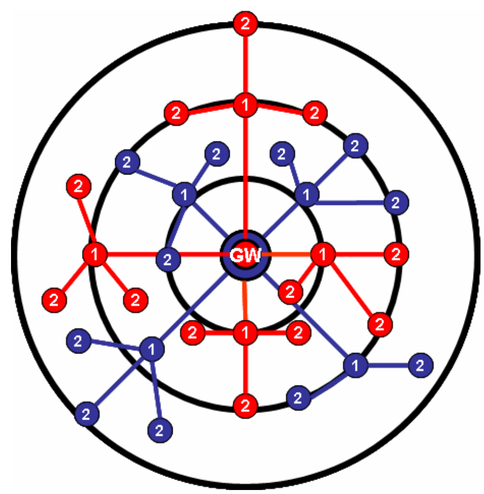

If there is a need to increase the coverage density, it can be accomplished by adding another SkyGateway and SkyExtenders which will link to the new SkyGateway. This interleaves the coverage areas of the SkyGateways and provides significantly denser coverage than would otherwise be possible with a mesh consisting of a single SkyGateway.

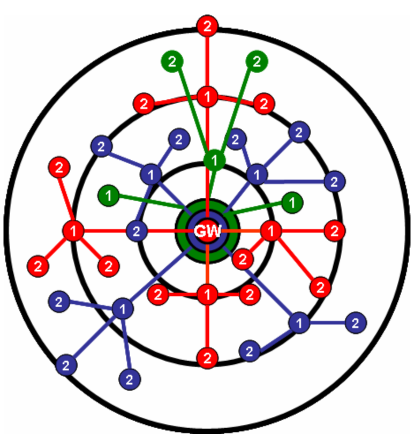

An example of interleaving another mesh network is shown in the example to the left. In this case there are three interleaved mesh networks used to increase coverage density and reduce self-interference. You can see that all the rules of spatial diversity have been followed in this example. In high density municipal mesh networks, this interleaving scenario is carried out with 4-5 SkyGateways at a single location.

Should coverage be required in the immediate vicinity of the SkyGateways, an additional SkyGateway should be added and links formed to first-hop ONLY nodes in order to serialize transmission to those nodes, as per the previous discussion.

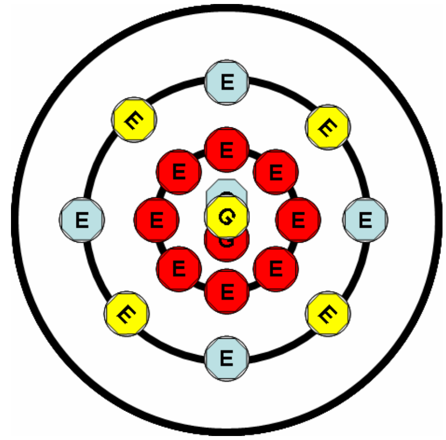

It's worthwhile to note that this strict attention to detail is most applicable to the first and second-hop nodes. These are the most active nodes in the mesh and they represent the greatest potential for self-interference if care and attention isn't paid. Beyond the second-hop, the focus changes from exact placement with respect to other nodes to a focus on ensuring that placement doesn't cause self-interference with first and second-hop nodes.

The Third-Hop and Beyond

We're finally to the "final" hop in terms of mesh network design. It is called the final hop because the rules that follow apply to all nodes beyond the third-hop. At this hop count we're not so much concerned about the perfect relationship between nodes as what we are the amount of self-interference generated within the first and second-hop nodes. The activity level in the third-hop and beyond becomes significantly lower as these nodes are aggregating less user traffic. If self-interference occurs among the third and greater hop nodes it's impacting a small amount of user traffic, if at all. If self-interference occurs with a first-hop node, it's impacting the aggregated traffic from the entire mesh behind it.

A node will receive data, traffic from multiple child nodes. The data is concatenated and sent to the parent node. This concatenation of data from multiple child nodes increases link transmission efficiency. Nodes that are multiple hop-counts from the SkyGateway are tasked primarily with handling small bursts of traffic from individual users. Nodes that are close to the SkyGateway, in terms of hop-count, are managing aggregated traffic from their child nodes.. Losing a radio frame at these nodes impacts a large number of users, in the mesh. This is why so much time was spent emphasizing careful node placement at the start of this section.

Locating the Third-hop SkyExtenders

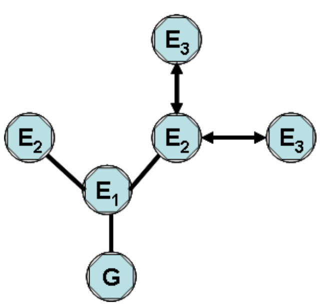

The critical factor in locating SkyExtenders in the third-hop, and beyond, is the elimination of potential self-interference with nodes operating in the first and second-hops of the mesh. There are a few rules of thumb that need to be considered when placing the third-hop SkyExtenders.

As a child node, the third-hop SkyExtenders should be located back, and away, from the antenna that their parent's parent is linked to. Ensure that placement of these nodes is such that when they transmit to their parent node, those transmissions will not interfere with nodes in the first and second-hop locations.

Reducing Self-Interference.
In the diagram above the SkyExtenders in the third-hop, labeled (E_3) are located "back" and "away" from the link between E_2 and E_1 providing not only spatial diversity, but also the avoidance of any self-interference with the closer hops in the mesh. The amount of spatial diversity, to be required here, is not as great as the links closer to the SkyGateway.

"Avoid forming child links on the antenna used to link to the parent node. In this situation, putting a third-hop node anywhere between E_1 and E_2 would be detrimental to the performance of the first-hop links on the network.

The SkyConnectors

In some implementations the use of a SkyConnector may be required to reach customers that may want a much higher Quality of Service, or may not otherwise be able to obtain wireless service, due to path obstructions, or issues, that cannot be overcome by the lower powered 2.4GHz side of the SkyExtender. This section will cover the basic rules surrounding the placement of the SkyConnector within the mesh.

Best Place to Locate SkyConnectors

SkyConnectors, like their SkyExtender counterparts, have some rules that surround their placement to ensure the best performance within the mesh. The vertical beam width and horizontal beam width on the SkyConnector has slightly less reach than a SkyExtender or SkyGateway, so the impact to the mesh isn't quite as bad if the placement isn't perfect.

When a SkyExtender transmits to the SkyConnector it transmits at the same power level that it would utilize when transmitting to another SkyExtender. It is important to locate SkyConnectors in such a manner as to avoid interference with other nodes, when the SkyExtender is transmitting to the SkyConnector.

Reduction of Self-Interference Generated by a SkyConnector

The rules about locating a SkyConnector are the same as those for locating a SkyExtender; the SkyConnector should be located back, and away, from the antenna used to link to the parent node. It should be located in an area such that when the SkyExtender transmits to it, the transmission will not interfere with other SkyExtenders. While this may be hard to avoid in some deployments this is the intended end-goal.

What If I Cannot Achieve Spatial Diversity In Some Areas?

In some areas of the network it’s will be nearly impossible to provide good spatial diversity. This is mainly due to a lack of resources upon which a SkyGateway or SkyExtender can be mounted. Not all areas, within a given municipality provide for perfect installations.

This following diagram displays a section of a deployed network where the location of the SkyGateway did not allow for spatial diversity. As can be seen from this diagram, all of the SkyExtenders are located on a single side of the SkyGateway. With the close proximity of the SkyGateway to the SkyExtenders in this mesh self-interference will be strong in all areas of the network.

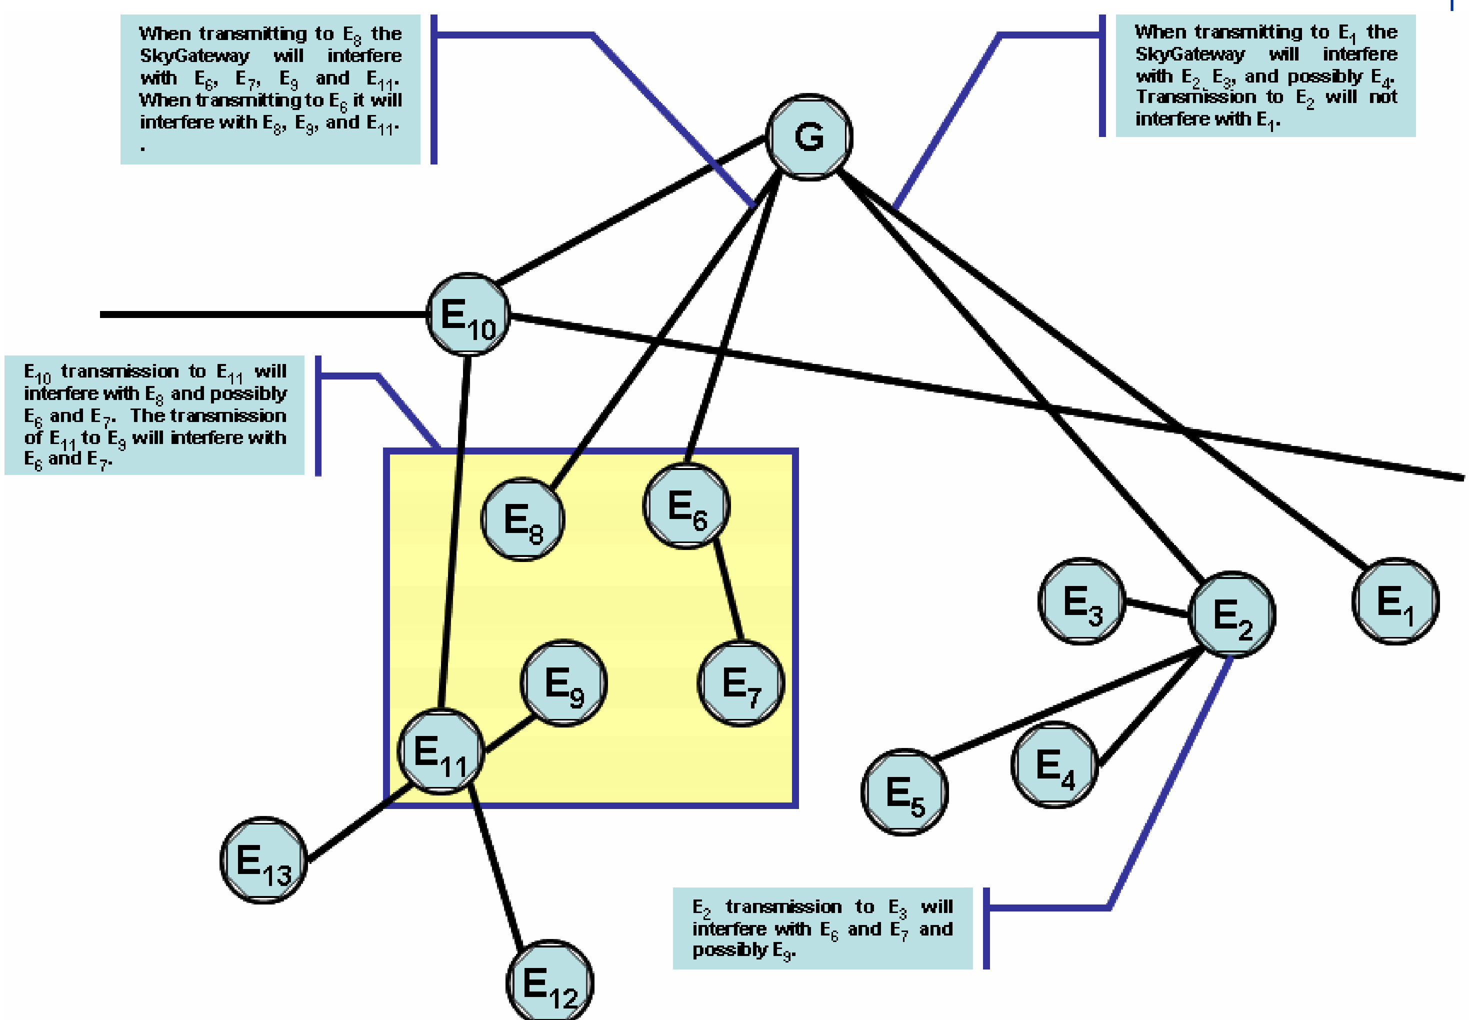

Notice that many of the second and third-hop SkyExtenders are located in the vicinity of other first-hop nodes, further increasing the chances of self-interference. The child nodes in this example are not pointing back, and away, from the parent node. The potential areas of self-interference are highlighted, on the diagram.

To correct this situation it is better to "serialize" all SkyExtenders to the SkyGateway; in other words, setup a preferred parent for all SkyExtenders to the SkyGateway and eliminate the second-hop nodes. If this presents too much of a load to the SkyGateway (i.e. you cannot serve up enough bandwidth to individual subscribers in this region) then the only choice to keep self-interference at an absolute minimum would be to add another SkyGateway (which is what this operator did) to the location and divide the SkyExtenders between the two SkyGateways.

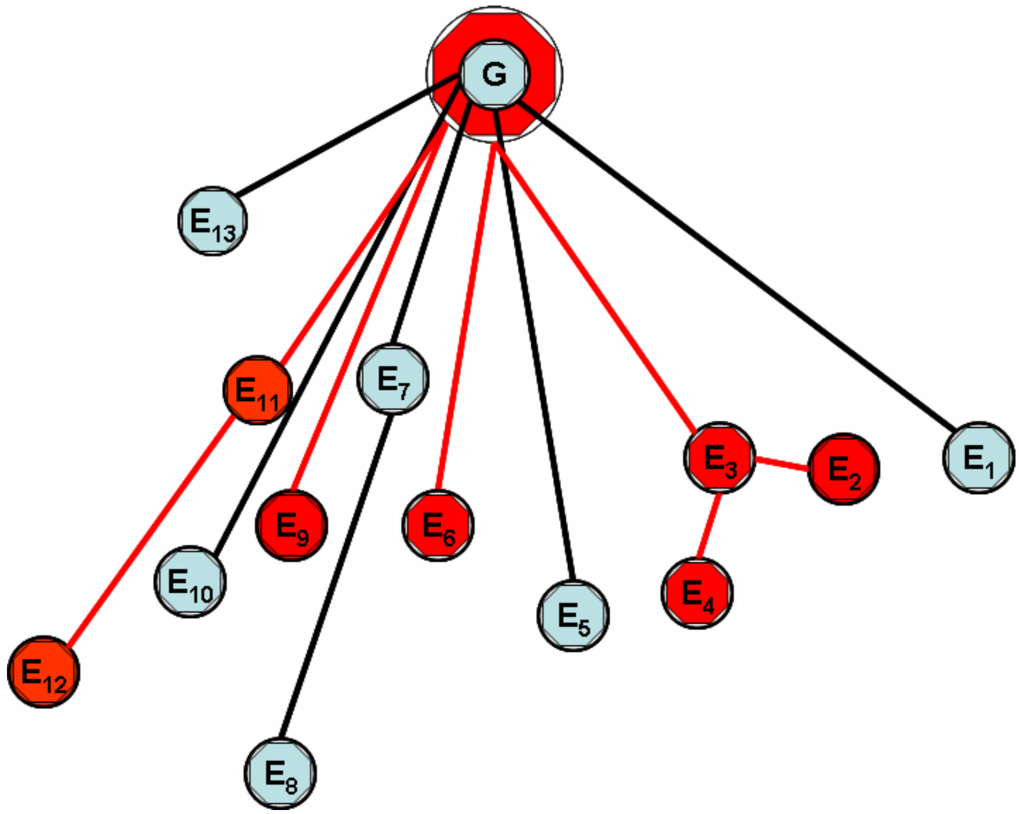

As can be seen from the example above two frequencies are used to segment the network. These frequencies are interleaved to produce maximum separation and reduce interference even further. In the instance where there are two child nodes (E_2 and E_4) transmission to either node will not cause self-interference with any other node.

There still is in this network some potential for self-interference when the SkyGateway transmits to E6. It is possible that this transmission may interfere with data transmission between E4 or E2 and E3. The node performance would have to be watched over time. If it's significantly degraded then E6 could be assigned to the other SkyGateway to prevent further self-interference.

Resolving Self-Interference Issues

This section will look at ways of mitigating self-interference within a network. We'll look at a few examples of how this can be done simply and effectively, as we did with the section on the inability to achieve good spatial diversity.

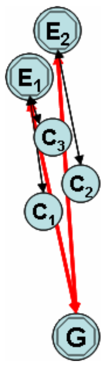

The example on the left shows a network segment where a large amount of self-interference potential exists. The diagram on the right shows this same network segment with preferred parents set to reduce the potential for self-interference. For sake of clarity, the SkyGateway is link to the SkyExtender E_2 via the red connecting line. This passes directly under SkyConnector C_3 but does not connect C_3 to the SkyGateway. The SkyConnector C_3 has a direct link to its parent SkyExtender E_1, in this example.

A cursory glance at this example reveals two issues:

1) Self-interference
2) The near/far effect (E_2 and E_1 are in very close proximity to each other but not their children, or their parent).

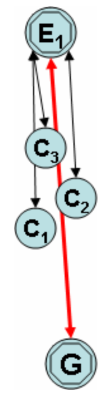

A first step to reduce the incidence of self-interference would be to make an attempt to remove one of the SkyExtenders – the elimination process involves removing the device with the worst links (lowest RSSI first, and modulation rate, second). If the network has periods of time when there is virtually no user data traffic you might want to reverse the order of qualification (i.e. modulation rate first and RSSI second). Always remember that RSSI directly translates into susceptibility to interference. Higher RSSI's translate into less susceptibility.

Once the SkyExtender is removed the SkyConnectors can be pointed toward the remaining SkyExtender, as shown in the diagram to the right. In this scenario, the SkyExtender is either communicating with the SkyGateway or one of the SkyConnectors. There can be no self-interference.

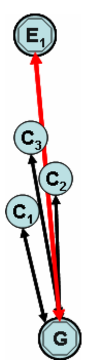

If the mesh extends to all sides of the SkyGateway and for hops beyond the E_1 SkyExtender it would be best to parent all of the SkyConnectors and the SkyExtender to the SkyGateway. While this arrangement may result in transmission inefficiencies, those inefficiencies will be much more desirable than the resultant self-interference if they are not parented to the SkyGateway.

In the event that both SkyExtenders are needed due to line of sight issues with the SkyConnectors we can try to reduce the amount of self-interference by parenting the SkyConnectors directly to the SkyGateway. While this is not always the recommended course of action, it is far better to parent a few SkyConnectors to the SkyGateway to reduce/eliminate self-interference with the SkyGateway or its first-hop SkyExtenders. Interference, in this area of the mesh, translates into data loss for all parts of the mesh.

Preferred Parenting

The SkyPilot system allows for Preferred Parenting. Through this process a new or existing node can be assigned a specific node as a preferred parent. When this parenting process is implemented the child node will form a primary link with the preferred parent if it's at all possible. Preferred parenting is used to tune an already existing network to further eliminate self-interference. Preferred parents can be assigned to SkyExtenders and SkyConnectors.

When establishing a preferred parent it is important to follow the location and placement rules described in the sections, above. This will ensure that the least amount of self-interference exists on the network.

Resolving the Near / Far Effect

The near/far effect arises whenever a strong signal, from a transmitting node that's very close to the receiver, interferes with a much weaker signal from the actual transmitter that's further away. The near signal (interferer) is stronger than the far signal making it impossible for the receiver to detect/decode the signal from the distant transmitter.

How can this issue be resolved? The diagram below is the starting point for the near/far effect. As can be seen from this diagram SkyExtender E_5 is much further away from the SkyGateway than SkyExtender E_4. To further complicate matters SkyExtender E_4 has a child node located between it and the SkyGateway. Whenever SkyExtender E_4 transmits to its child node, at the same time that SkyExtender E_5 is transmitting to the SkyGateway the near/far effect will be encountered.

A similar situation will occur on the opposite side of the SkyGateway. Transmissions from SkyConnectors (C_1, C_2) will be much stronger, at the SkyGateway than transmissions from SkyExtender E_3 to the SkyGateway, when these transmissions occur simultaneously. This is another example of the near/far effect.

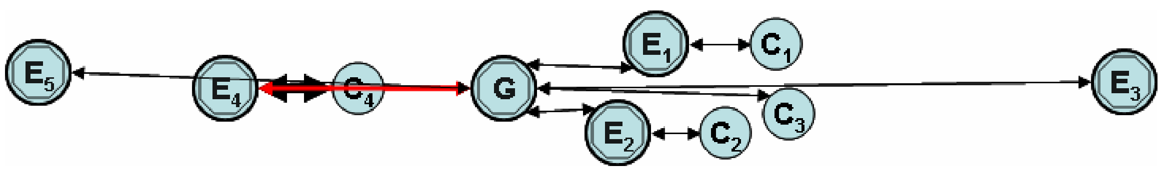

Preferred parenting

We've seen how to identify the near/far effect in the network. The next step is to resolve it. The first choice in resolving issues with power imbalance, or near/far effect, is to attempt to set preferred parents on all offending devices on the network. This resolution is the quickest and easiest to implement on a network. One change to the network above has been made in the illustration below. This change clearly shows the use of preferred parenting – C_4 was parented directly to the SkyGateway, in this example.

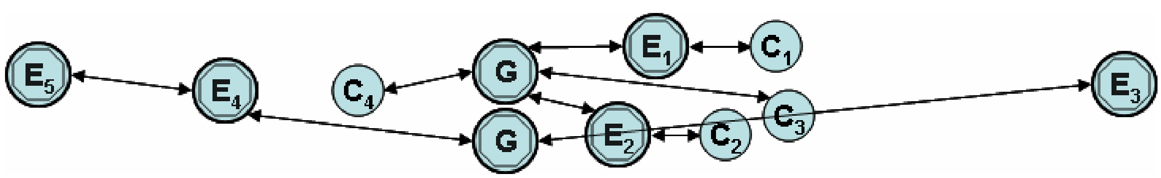

Additional SkyGateways – The Best Resolution

The strategic approach to resolving this issue is to configure a SkyGateway to handle all of the near devices, and a SkyGateway to handle all of the far devices. Both SkyGateways would be co-located and configured to different frequencies in order to avoid interference. This is illustrated in the above example.

By utilizing multiple SkyGateways, each handing devices at a different distance, the power imbalance issues that would normally occur can be avoided and the links can operate at their maximal efficiency as bounded by the network design.

Summary

Self-interference compromises link integrity between a node and its children as a result of two nodes transmitting toward the same receiver, simultaneously. This results in a decrease in modulation rates for all receivers impacted by the self-interference. Self-interference can be a result of SkyExtenders in close proximity to each other and other SkyConnectors, or near/far effect interference - which occurs as a result of low strength signals being over powered by high strength signals.

When confronted with the possibility of self-interference you should be able to go back to the design rules (i.e. child nodes should be pointed back and away from the parent node, and avoid forming child links on the antenna used to link to a parent node) and determine if the interferer meets this criteria, or not. If it does not, re-evaluate the link of the interferer, or the node impacted by the interference, to determine if there is a better link to be chosen or formed that will resolve the self-interference. Repeat the process for each link experiencing significant self-interference. Reroute for better antenna diversity or change frequencies for appropriate nodes.

Once links have been reworked and the mesh has formed run the nodeInfo script again and view the self-interference output to determine if the implemented changes were positive. Repeat the process until the network has as much spatial diversity and as little self-interference as can be accomplished. The more that individual networks can be interleaved and the greater the spacing between SkyExtenders of the same frequency, the less chance there will be for self-interference.

To reduce/avoid the potential for near/far effect it is suggested that two SkyGateways be configured. One SkyGateway will manage nodes in close proximity. The other SkyGateway will manage the more distant nodes. Both SkyGateways should be co-located on different frequencies so as not to interfere with each other.

Sometimes, strategically relocating a SkyExtender or SkyConnector may significantly reduce the amount of self-interference generated between it and another node.

Using nodeInfo for Additional Help and Support

SkyPilot has created a script called nodeInfo (downloadable from www.skypilot.com). The nodeInfo script queries SkyPilot nodes for data about the quality of their links, GPS coordinates, IP and MAC address, the RSSI and modulation rates and other information. It can also compute the potential for self-interference to exist on any link or antenna for any node.

The nodeInfo script is commonly used to generate an output files with information specific to self-interference. The two files of interest are the *.kml (Google Earth compatible) and a self-interference (csv) file. Both files perform useful functions and but have slightly different applications.

The KML file (Google Earth) is generated by querying nodes for their GPS coordinates, which then allows them to be plotted directly onto the Google Earth map. An individual icon is shown for each SkyGateway, SkyExtender, SkyExtender Dualband, SkyExtender Triband, SkyConnector and SkyAccess. Links are also shown on the output.

A user can click on any node icon to display information about the node. The informational text is displayed in a popup window. The data includes all current active and standby links, link modulation rates, link RSSI and interferers for the link. This provides both a textual and visual view of the links and the potential interferers.

The self-interference csv file contains the same information presented by the KML file only in a more text-based format for individual links. As CSV files can be imported directly into MS Excel and can be easily sorted and filtered, the information can be manipulated in various ways to view the impact of self-interference on the network. This data can then be compared to the Google Earth output and appropriate changes can be made to the network.

nodeInfo Output Example

skypilot_extDB

Telnet - Web Interface

```
IP                  : 10.0.0.31
MAC Address         : 01:32:ff
Active Image Name   : SkyExt.1.3p6.bin
Active Image State  : Accepted
Backup Image Name   : SkyExt.1.4p3_beta1.bin
Backup Image State  : Accepted
Active Image Bank   : A
Node type           : SkyExtender DualBand
Switchboard Version: B
BootROM             : 001.06.000.S
Preferred Parent    : 00:0a:db:01:19:a3
Uptime              : 9 days 04:13:44
Reboot Reason       : CLI Reboot

MAC Address Hostname                 IP Address      NType State     LRSSI RRSSI LTxMod RTxMod LAnt RAnt Range
----------- ------------------------ --------------- ----- --------- ----- ----- ------ ------ ---- ---- -----
03:0d:ac    RummyRVArea              10.0.0.143      cpe-o act mgmt     28    22     36     36    3    0  1350
02:17:e6    Communications           10.0.0.240      cpe-o standby-o    18    18     12     18    1    0  3000
03:14:64    mounch                   10.0.0.192      cpe-o act mgmt     17    16     12     12    2    0  1800
     Tx Interferer Hostname         IP Address      NType   Rx Interferer Hostname         IP Address      NType SNR
     ------------- ---------------- --------------- -----   ------------- ---------------- --------------- ----- ---
     03:0d:1f      gruy             10.0.0.158      cpe-o   01:0a:4a      skypilot_ext2    10.0.0.22       ext     7
     02:17:e6      TomCoxExt4       10.0.0.240      cpe-o   01:19:97      skypilot_ext4    10.0.0.24       ext     8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     01:15:4e      Chris_Paine      10.0.0.82       cpe-o   8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     03:09:46      CUSTOMER1_CPE    10.0.0.117      cpe-o   8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     02:08:8e      Jacksons6        10.0.0.214      cpe-o   8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     03:0b:3f      The_Builders     10.0.0.140      cpe-o   8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     01:0e:38      Sunday           10.0.0.72       cpe-o   8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     02:09:c0      ScottStill       10.0.0.205      cpe-o   8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     03:16:e3      LottExt6         10.0.0.186      cpe-o   8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     02:17:93      RussoRes         10.0.0.221      cpe-o   8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     03:16:1e      Huckfin          10.0.0.189      cpe-o   8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     03:09:d8      Laney            10.0.0.118      cpe-o   8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     03:0b:1e      Messersmit       10.0.0.131      cpe-o   8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     03:07:af      Wesgoggen        10.0.0.174      cpe-o   8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     01:32:0a      CUSTOMER1_CPE    10.0.0.139      cpe-o   8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     02:09:ae      AreaBuilders     10.0.0.209      cpe-o   8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     03:09:f9      Airport          10.0.0.130      cpe-o   8
     01:1f:99      skypilot_ext6    10.0.0.26       ext     02:17:81      WynneWin         10.0.0.220      cpe-o   8
```

The above example is a snapshot of the data presented when you click on an icon from within the Google Earth KML file. The top lines display the hostname of the selected node, along with other attributes (IP Address, MAC Address, firmware version, reboot reason, uptime, etc). The information that follows this text presents data about links to the node (black text) and the potential interferers to those links (red text).

In this example there are three displayed links to the SkyExtender. There are two active links as shown by the “act mgmt” in the “State” field. There is also one standby link, shown by the “standby-o” in the “State” field. Not all links for this SkyExtender are displayed in this example. The link to the parent, at a minimum, is missing. If it were present it would be displayed with a “State” value of “act path”.

One of the three displayed nodes has the potential for self-interference. This node has a MAC address ending in 03:14:64. The node hostname is “mounch” and the IP Address is 10.0.0.192. This CPE has an active path to the SkyExtender.

How is this data to be interpreted? Whenever the SkyExtender (skypilot_extDB) transmits to the CPE (03:14:64) there is potential for that transmission to be interfered with <u>at the CPE</u>. When the CPE transmits to the SkyExtender there may be potential for interference, but that potential is not displayed in this output. Clicking the icon for the CPE and looking at the link data for skypilot_extDB will display interference potential when the CPE transmits to the SkyExtender.

We’ve established that the transmissions from skypilot_extDB to the SkyConnector 03:14:64 can be interfered with. What further information can be extracted from the data presented?

The column **Tx Interferer** displays the MAC address of a node that, when it transmits, may interfere with the reception of data transmitted between skypilot_extDB and the SkyConnector 03:14:64. These transmissions ***must occur simultaneously*** in order for interference to exist. The following three columns provide additional information about this transmitting node including the hostname, IP address and node type. If multiple interferers exist, for transmissions between skypilot_extDB and SkyConnector 03:14:64, they will be listed in order of their MAC address. This is the case for the information presented above.

The next four columns display information about the **Rx Interferer**. The Rx Interferer is the node receiving transmitted data from the Tx Interferer when the self-interference potential exists. In this example, whenever skypilot_extDB transmits to SkyConnector 03:14:64 at the same time that SkyConnector 03:0d:1f transmits to skypilot_ext2, interference will occur and data transmitted to SkyConnector 03:14:64 may be lost. Following the Rx Interferer are the same columns as those that follow the Tx Interferer. They provide the hostname, IP address and node type of the node receiving the interfering transmission.

The last column, entitled **SNR** provides a picture of how bad the interference might get. The SNR value is calculated by subtracting the signal strength of the unwanted transmission from the signal strength of the wanted transmission. The more positive the number is the lower the chances of data corruption and data loss when the interfering transmission occurs. The more negative the number is the greater the chances of data corruption and data loss when the interfering transmission occurs. A value of 0 indicates that the wanted signal and the unwanted signal are equal. Values less than 0 indicate that the unwanted signal is stronger than the wanted one and values above 0 indicate that the wanted signal is stronger than the unwanted one.

Modulation Rate	Required SNR
6Mbps	3dB
9Mbps	4dB
12Mbps	5dB
18Mbps	6dB
24Mbps	10dB
36Mbps	13dB
48Mbps	19dB
54Mbps	22dB

The SNR values in the above table are required values. The listed SNR must be maintained in the presence of ALL environmental co-channel interference PLUS self-interference. If the SNR is not maintained at the required level then the modulation rate will drop to the appropriate rate given the channel SNR.

The potential for self-interference to impact a link is dependant upon the activity level, or duty cycle, of the transmitting interferer. Even with high QoS rates, SkyConnectors typically do not exhibit high activity levels for extended periods of time. If the Tx Interferer were the SkyGateway and the link being interfered with were a SkyExtender (especially in the first or second hop) the issue of self-interference becomes something that requires immediate attention. The SkyGateway has a high activity level as do the SkyExtenders in the first few hops. The probability of self-interference occurring, under those conditions, is very high. It is important to weigh the probability of the self-interference occurring, to effectively evaluate this data.

Another indicator of self-interference can be found in the TxMod columns. If self-interference is an issue the RTxMod value, for the SkyConnector 03:14:64, would drop whenever the mesh surrounding it becomes active. If activity in the mesh increases and modulation rates are decreasing it indicates self-interference. Focus on the high-impact links, first (SkyGateway and first-hop SkyExtenders) and then work your way through the mesh parenting devices according to the rules presented above.

Leading the Mesh Revolution

1100 Island Drive
Redwood City, CA 94065
408.764.8000
US Toll Free 866 SKYPILOT
sales@skypilot.com