Cambium Networks™

ePMP ABAB Frequency Reuse Deployment

This document discusses ePMP GPS Synchronization, Frequency Planning, Interference, RF Planning and Deployment, and Best Practices for ABAB Channel Planning.

Table of Contents

Figures

Tables

Introduction

The ePMP system is a wireless system that provides throughputs up to 114 Mbps in a 20 MHz channel and up to 230 Mbps in a 40 MHz channel and covers ranges up to 13 miles. The performance of any wireless system is affected not only by the noise and characteristics of the propagation channel, but also by the interference levels generated by neighbouring systems. The neighbouring transmitters can be part of the same network (self-interference) or part of another network (external interference). Self-interference occurs when signals come from a network under the control of one single operator, and it can come from the same tower or from several miles away. When a network becomes very large or very dense, it will be more exposed to self-interference and the performance of the system starts degrading.

A careful frequency planning can strongly decrease a network's self-interference. If self-interference is not considered during frequency planning, it can limit the overall network performance in the uplink or the downlink direction. This reduced performance will be experienced by the end users as not being able to achieve the expected throughputs and in some cases higher packet loss.

Several techniques can be used to minimize the effect of self-interference.
For example, the ePMP system uses highly directional antennas on the STA's site to focus the transmitted energy in narrow beamwidths with high gain. This minimizes the leakage of the signal from the sides and the back of the antenna, reducing the interference levels to other transmitting STAs. Also, the ePMP air interface protocol uses a centralized scheduler at the AP site which controls which STAs can transmit during the uplink. This reduces data corruption and retransmission, maximizing the throughput.

One of the most effective ways to limit self-interference is GPS synchronization. This paper outlines the advantages of GPS synchronized networks and describes self-interference scenarios in systems deployed with ABAB 4-sector 2-channel reuse configurations (which are not unique to ePMP networks) and suggests methods for decreasing the self-interference level and minimize the impacts to performance.

GPS Synchronization

GPS synchronization is one of the most effective ways to limit self-interference. GPS synchronization allows all APs to share an accurate timing reference so that all can transmit and receive at the same time. With GPS synchronization the network can scale to significantly larger sizes without any AP's transmission interfering with the reception of the others and a more efficient frequency reuse across sectors and towers can be employed. By deploying all sites in a network using the same framing parameters (frame size, duty cycle, max range), the downlink and uplink times of all sectors never overlap.

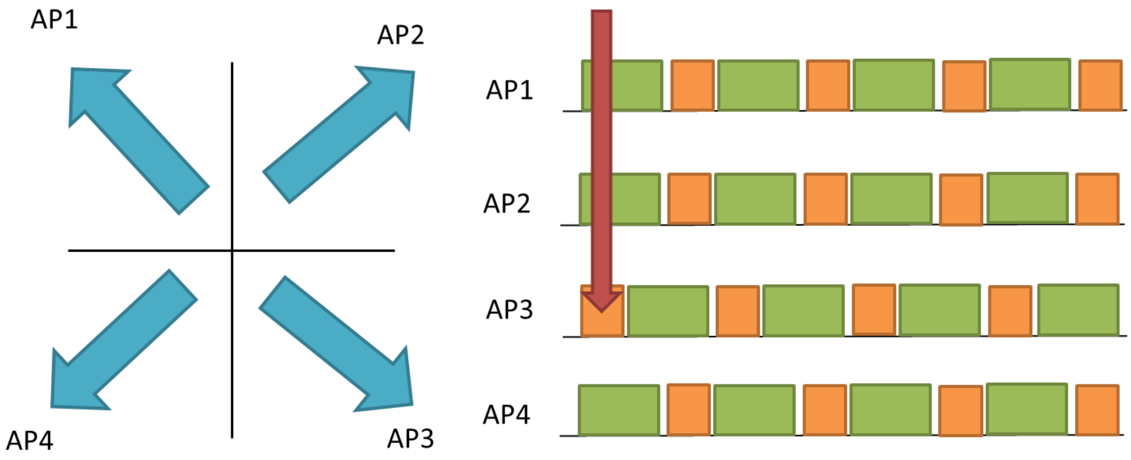

Figure 1 - Non-synchronized sector

Figure 1 shows a cluster of four APs in the four sector of one tower. If the sectors are not synchronized (for example, in Figure 1 AP3 is not synchronized), the receive time of one AP (orange region) can overlap with the transmit time of another AP (green region), making it impossible to reuse the same channel frequency on multiple sectors and also creating a lot of interference in sectors using adjacent frequencies.

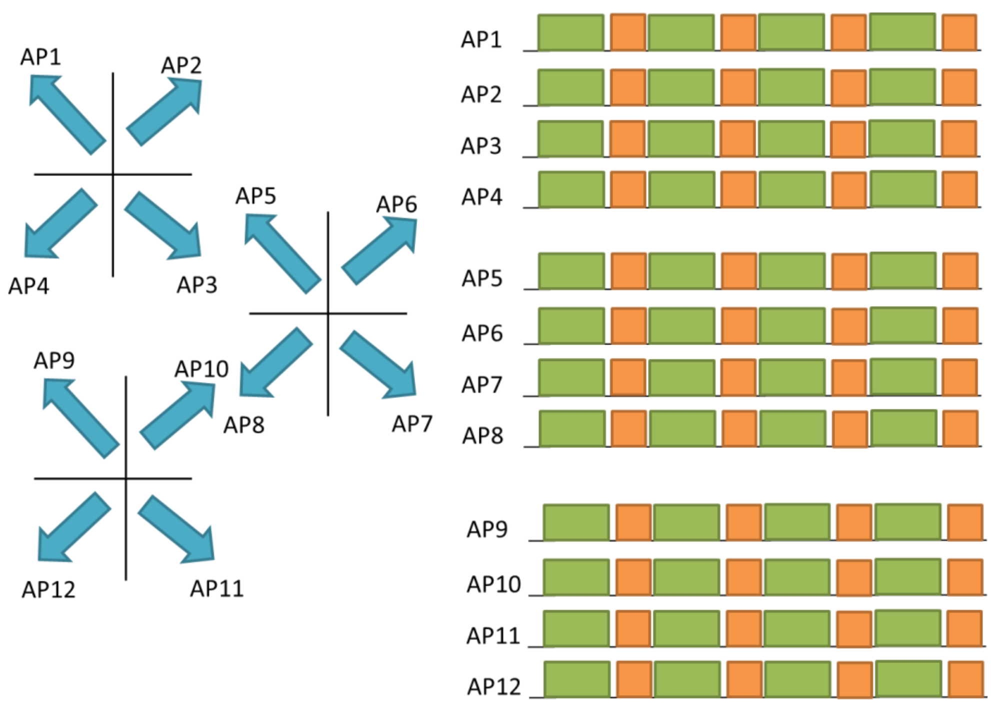

Figure 2 - Synchronized network

Figure 2 shows a synchronized network, in which not only all the APs in one cluster share the same accurate timing information, but also all the clusters across the network. Now the transmit and receive times of all APs are aligned and the network can grow to any size.

GPS synchronization eliminates any source of self-interference due to an AP receiving when an adjacent AP is transmitting and a STA receiving when another STA in the same sector is transmitting. The only self-interference left in a synchronized network comes from sector to sector or tower to tower interference where uplink noise comes from STAs on other sectors and where downlink noise comes from APs on other towers or sectors.

Of course, additional noise / interference could also come from competing systems within the same area on the same channels.

Frequency planning

As mentioned above, the ePMP system can deliver up to 114 Mbps (230 Mbps) to STAs connected to one sector using a 20 MHz (40 MHz) channel.

For an efficient network deployment with four-sector sites it is recommended to use four channels (A, B, C and D) for a total of 80 MHz (or 160 MHz) of spectrum.

The reason why a two-channel four-sector deployment is not efficient is shown in Figure 3.

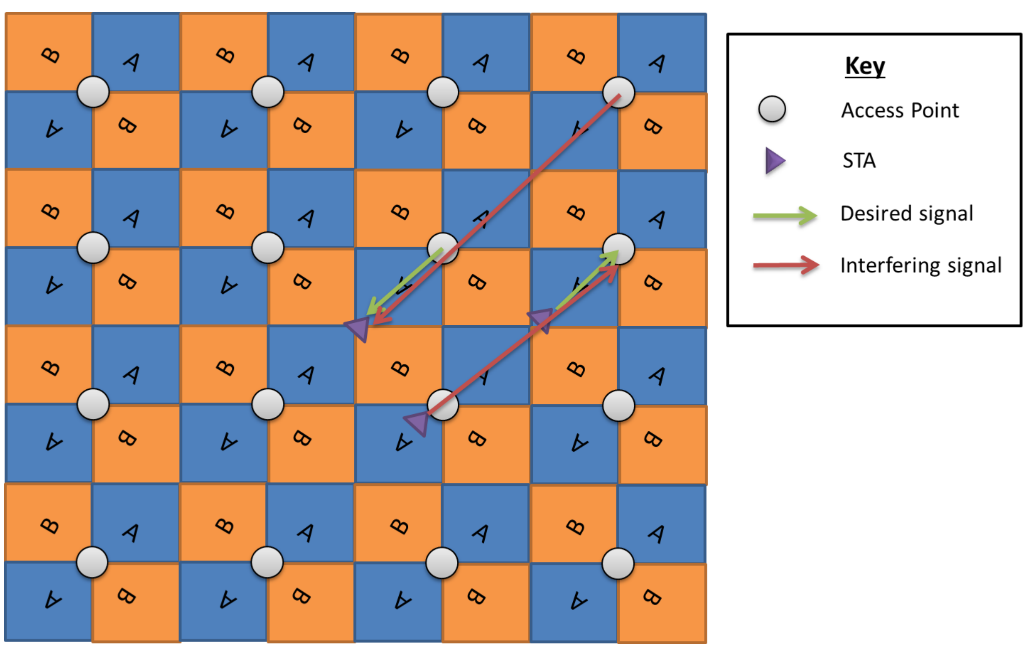

Figure 3 - Two-frequency four-sector deployment

If only two channels (A and B) are available, the only efficient solution is having channel A back to back on each tower, with channel B back to back on the two remaining sectors of the tower. The reason is that the high front-to-back ratio of the AP antenna provides more isolation when the same frequency is reused in the sector in the back compared to the case in which it is used in an adjacent sector, because the side isolation is not as good as the back one.

In this type of deployment the highest level of self-interference at the STA's receiver occurs when the STA is at the edge of a sector using channel A, as shown in Figure 3, and it additionally receives an interfering signal from another AP also on channel A. The attenuation of the interfering signal compared to the one of the desired signal is given by the additional distance between the STA and the interfering AP. With the distances between the STA and the two APs having a ratio of 3, as shown in the Figure, the interfering signal is attenuated by $20\log(3) = 9.5\ \mathrm{dB}$. This attenuation is enough only for operation with MCS1. Higher MCS values cannot be used because the attenuation is lower than the CCI requirements for all other MCS levels.

In the same type of deployment, the highest level of self-interference at the AP's receiver occurs when the AP on channel A receives a signal from a STA at the edge of the cell, and at the same time a STA in another sector also using channel A transmits uplink data. The interfering STA can be very close to the adjacent AP on the same channel, making the ratio of the distances of the AP from the two STAs only 2. In this case, the attenuation of the interfering signal is only $20\log(2) = 6\ \mathrm{dB}$, which makes it impossible to communicate even at the lowest MCS level. Note that this worst case scenario is somewhat attenuated by two reasons: first, the interfering STA is very close to the AP it is connected to, so its transmit power could be reduced by the Automatic Transmit Control algorithm; second, the interfering condition only occurs if the two STAs in the example are scheduled to transmit at the same time. With many STAs connected to each AP, the probability of this occurring is low. On the other hand the interference in the downlink is persistent, because all APs transmit for the duration of the downlink, which is synchronized across the whole network. In this deployment scenario the self-interference of the network is clearly limiting the performance of the system.

A more efficient deployment uses four channels (A, B, C and D). How can an operator select which frequency to allocate to which sector in order to minimize the self-interference of the whole network?

Two solutions are shown here. Figure 4 shows a deployment in which each tower uses all four channels; Figure 5 shows a deployment in which alternating towers use either the A and B channels or the C and D channels.

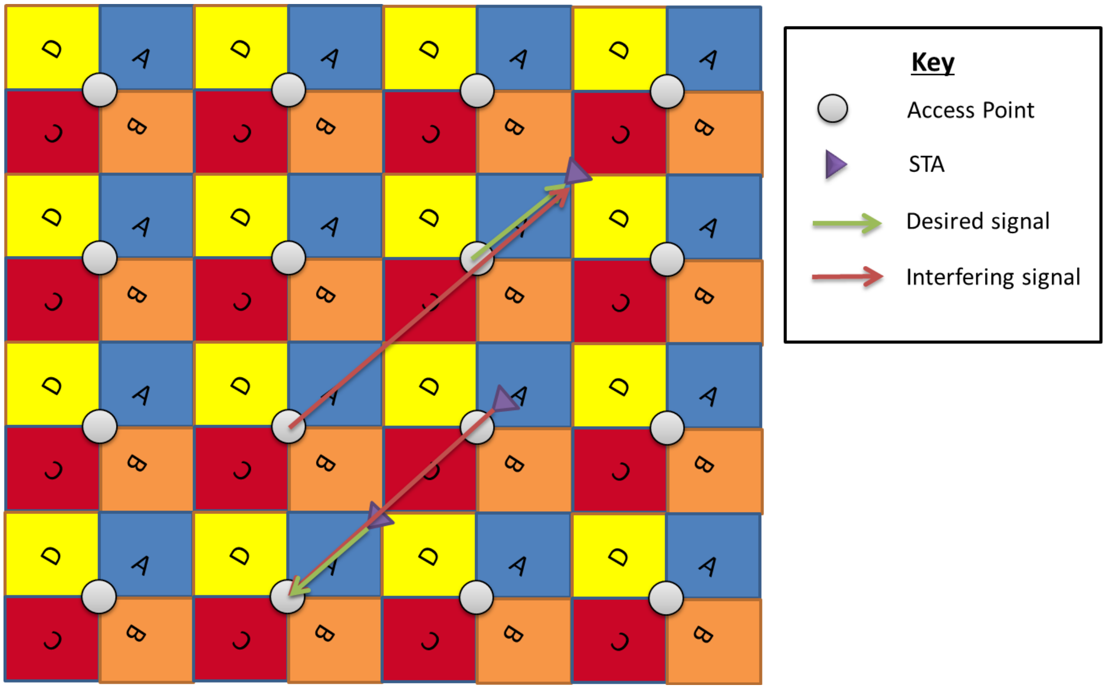

Figure 4 - Four-frequency four-sector deployment - option 1

In option 1, each of the four sectors in the same tower use different frequencies. Figure 4 shows the example in which sectors facing each direction reuse the same frequency. In this example, the same self-interference scenario described in the two-frequency four-sector deployment occurs.

Figure 4 shows a STA at the edge of the sector using channel A receiving a desired signal and an interfering signal from an AP also using channel A. The difference in received power between the desired and interfering signals is given by the additional distance between the STA and the interfering AP. The ratio between the two distances is 3, which means that the interfering signal is only 9.5 dB weaker than the desired one. This link can use only MCS1 for communication, regardless of the cell size.

Figure 4 also shows two STAs transmitting in the same uplink time to two APs using channel A facing the same direction. As seen in Figure 3, the ratio between the two distances can be as low as 2, giving only 6 dB of attenuation, which is not enough even for communication at MCS 1.

It is possible not to assign the same frequency to sectors facing the same direction, but rotate the four frequencies in each adjacent sector. With any rotation there are always two adjacent sectors sharing the same frequency and facing the same direction, which creates more self-interference than the deployment shown in Figure 4.

Figure 5 shows another option, in which one tower uses channels A and B only in opposite facing sectors, while the next tower uses channels C and D only in opposite facing sectors.

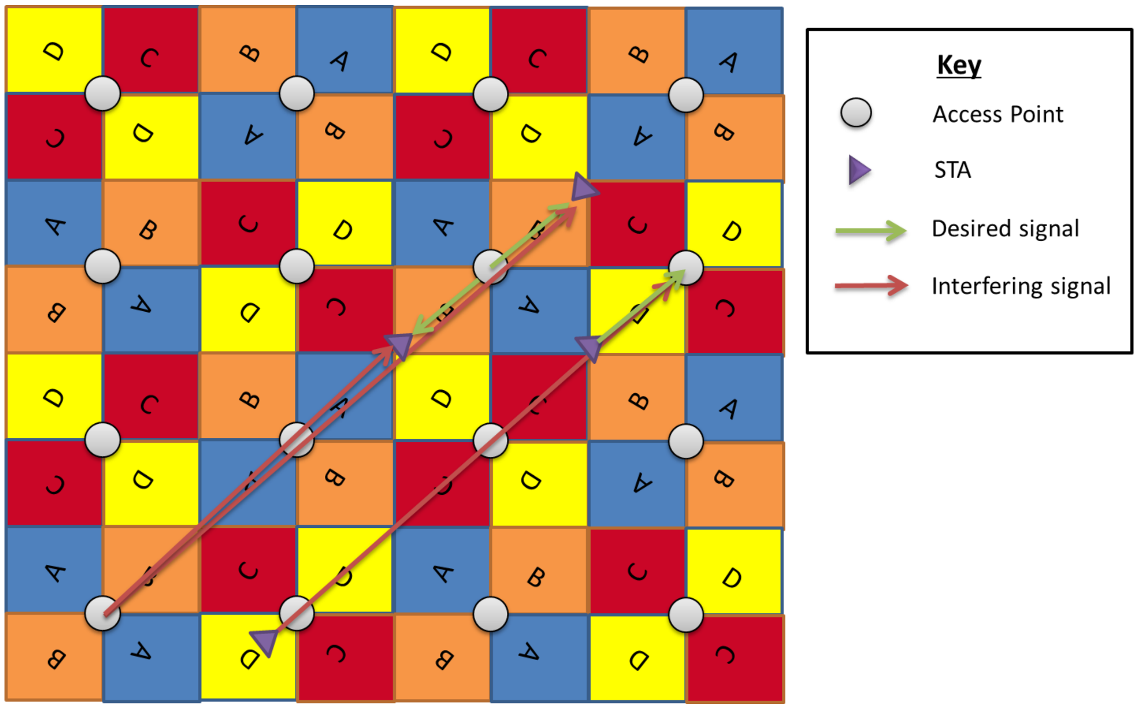

Figure 5 - Four-frequency four-sector deployment - option 2

With this type of deployment a STA at the edge of the sector using channel B receives an interfering signal from the closest sector also using channel B whose AP is five times farther than the AP transmitting the desired signal. This means that the interfering signal is attenuated by $20log(5)$=14 dB, which is enough for communication at MCS10.

Also, when two STAs transmit at the same time to two APs using the same channel and facing the same direction, the ratio between the two distances is 4, which means that the interfering signal is now attenuated by $20log(4) = 12\ dB$, which is enough attenuation for communication with MCS 9.

This is the reason why the deployment shown in Figure 4 is the recommended one with four-sector towers.

Additionally, in any type of deployment another typical concern is the interference a STA experiences from the back lobe of its antenna when receiving data from its AP and also another AP on the same channel located behind the STA. With the frequency planning shown in Figure 5, the signal strength of the interferer is attenuated by the front-to-back ratio of the STA's antenna plus the additional attenuation of the interfering signal due to the different path length. The ratio between the distances of the STA with the AP it is connected to and the STA with the interfering AP is 3, therefore the total attenuation is equal to the front-to-back ratio of the STA's antenna plus 9.5 dB, which is enough for communications at higher MSC.

Self-interference from co-located sectors

The previous section showed the recommended frequency planning for minimizing self-interference across towers of the same network. This section describes self-interference scenarios when the self-interference comes from an AP (downlink) or a STA (uplink) of another sector in the same tower.

Sources of Downlink Self-Interference

Downlink self-interference occurs when one or more APs in the same tower emit signals that degrade signals transmitted by other APs. It is critical that co-located APs are set to the same transmit power to minimize interference to and from co-located sectors.

Let us consider the case of one tower with four sectors and two frequencies, where frequency A is deployed back-to-back, and so is frequency B.

Example 1: AP on sector A1 and AP on sector A2 on the same channel

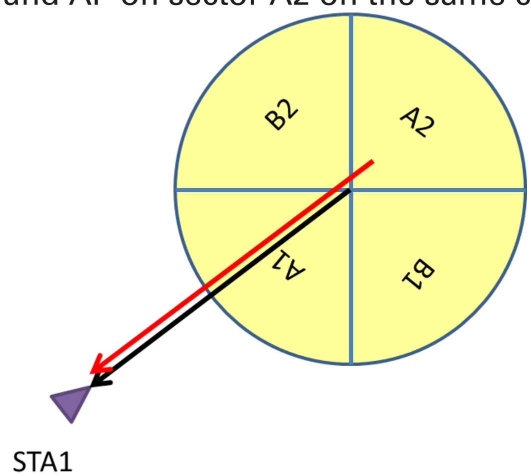

Figure 6 – Downlink interference from back-side on channel A

Figure 6 shows an AP using channel A (sector A1) communicating in the downlink with STA1. During the downlink time, the AP on sector A2 is also transmitting using channel A. The signals transmitted by the AP is sector A2 may leak from the back lobe of the antenna, degrading the signal transmitted by the AP in sector A1. The interfering signal is attenuated by the front-to-back ratio of the AP antenna. If this value is not large enough, STA1 will not be able to communicate at higher MCS levels.

This front-to-back isolation can be further enhanced by increasing the distance between the sectors on opposite sides of the tower or by placing additional obstructions between the two AP sector antennas.

Example 2: APs in sector A1 and sector B1 are on adjacent channels and STA1 is on the border between the two sectors

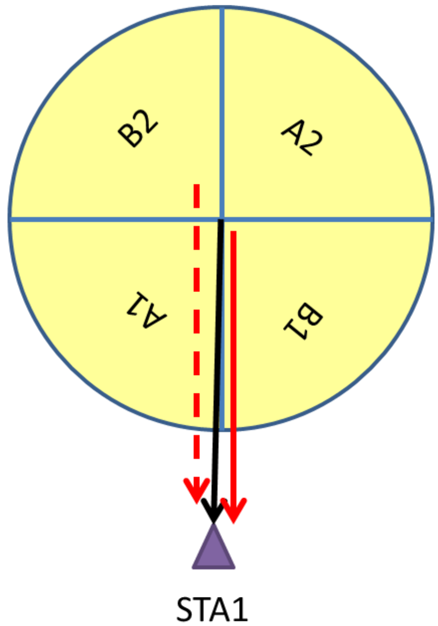

Figure 7 - Downlink interference from adjacent channel

Figure 7 shows an AP transmitting on channel A (sector A1) to a STA at the edge of the sector. The signal to STA1 could be degraded by the signal transmitted by the AP in sector B1, and possibly the signal transmitted by the AP in sector B2, which both use a channel B, adjacent to A.

This source of interference is prevented by the adjacent channel rejection of the radios and the antenna systems.
The roll-off of the off-channel transmit power for B1 signals arriving at STA1 could cause interference. The roll-off of the off-channel transmit power for B2 signals arriving at STA1 is further reduced by the AP in sector B2's antenna isolation.

This potential for interference is reduced by providing a channel separation between adjacent channels. In the ePMP system the recommended minimum channel separation is 5 MHz.

Sources of Uplink Self-Interference

Since all STAs utilize Auto Transmit Power Control (ATPC) to maintain a target receive level at the AP, uplink interference from co-located sectors using the ABAB frequency reuse can cause issues.

Example 3: Uplink interference from back side on-channel STA

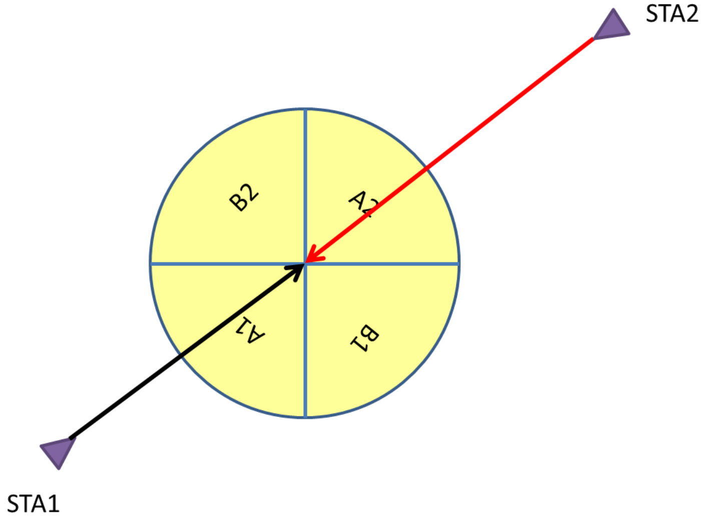

Figure 8 - Uplink interference from back-side on channel STA

Figure 8 shows STA1 communicating on channel A with the AP of sector A1. At the same time, STA2 communicates also on channel A with the AP of sector A2. The signal from STA2 can interfere with the signal from STA1 if the two STAs are scheduled to transmit at the same time. The interference is stronger if the distance between STA2 and its AP is much lower than the distance between STA1 and its AP.

This source of interference is minimized by the front-to-back ratio of the sector antennas. If the isolation provided by the front-to-back ratio of the antenna is not sufficient to remove the interference, STA1 would not be able to communicate at higher MCS levels. This front-to-back isolation can be further enhanced by increasing the distance between the sectors on opposite sides of the tower or by placing additional obstructions between the two AP sector antennas.

Example 4: Uplink interference from STAs in adjacent sectors using adjacent frequencies

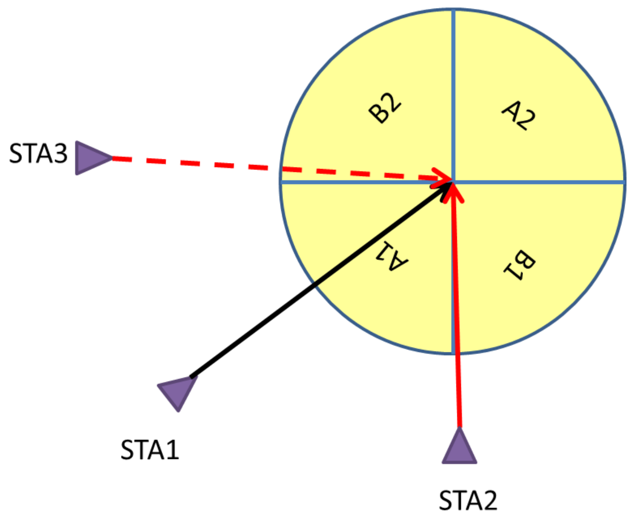

Figure 9 – Uplink interference from adjacent sector off channel STAs

Figure 9 shows STA1 communicating on channel A with the AP of sector A1, while STA2 communicates on channel B with the AP of sector B1 and STA3 communicates on channel B with the AP of sector B2. Additionally, STA2 and STA3 are at the edge of their sectors. The signals from STA2 and STA3 can degrade the signal from STA1, especially if the distances between STA2 and STA3 to AP2 and AP3 respectively are much smaller than the distance between STA1 and AP1. The antenna pattern of the sector antennas typically prevents this kind of interference by rejecting adjacent channel. By rejecting this off-channel power, the CINR or carrier to interference ratio is kept at a level to allow full performance of the in-channel STA.

This potential for interference is reduced by providing a channel separation between adjacent channels. In the ePMP system the recommended minimum channel separation is 5 MHz.

RF Planning and Deployment

Two key parameters must be considered for RF deployment planning:
1) Adjacent Channel rejection. In order to limit interference from the adjacent channel, the use of a guard band of 5 MHz between channels is recommended.
2) CINR/RSSI. In order to reach the desired throughput and service level for an individual STA and to maximize the overall system performance, the CINR/RSSI should be maximized. The table below shows an example of CINR/RSSI requirements for downlink and uplink MCS performance for a 20 MHz channel in the 5.8 GHz band.

Table 1 - CINR and RSSI requirements for 20 MHz channel in 5.8 GHz band

Downlink				Uplink		
MCS	Min CINR (dB)	Min RSSI (dBm)		MCS	Min CINR (dB)	Min RSSI (dBm)
MCS 15 64QAM 5/6	28	-70		MCS 15 64QAM 5/6	30	-68
MCS 14 64QAM 3/4	27	-71		MCS 14 64QAM 3/4	28	-70
MCS 13 64QAM 2/3	25	-73		MCS 13 64QAM 2/3	25	-73
MCS 12 16QAM 3/4	20	-78		MCS 12 16QAM 3/4	21	-77
MCS 11 16QAM 1/2	17	-81		MCS 11 16QAM 1/2	16	-82
MCS 10 QPSK 3/4	14	-84		MCS 10 QPSK 3/4	15	-83
MCS 9 QPSK 1/2	12	-86		MCS 9 QPSK 1/2	11	-87
MCS 1 QPSK 1/2 SS	9	-89		MCS 1 QPSK 1/2 SS	8	-90

In RF planning, it is necessary to perform both CINR and RSSI analysis on each sector. In simplistic terms, the RSSI is the amount of power being received by the radio and CINR is the margin between the useful signal and the noise floor. The AP/STA system selects the modulation level, and thereby the available throughput, primarily based on the CINR (Carrier to Interference and Noise Ratio).

Additional comments on RSSI and CINR:

- A given RSSI does not always equate to the same MCS state.
    - MCS1 uses a single stream transmission in which only one stream of data is sent by both transmit antennas. The received signals are then combined using the MRC (Maximal Ratio Combining) technique providing gain from environmental effects. This scheme improves CINR without a corresponding increase in RSSI, which means that the same RSSI on two sites can have a 3 dB difference in CINR and therefore provide different levels of throughput.
    - Presence of external noise can degrade performance
- A tower layout design based on RSSI coverage alone can yield a satisfactory result if there is no interference (external noise or self-interference).
- A tower layout design based on CINR analysis is needed in addition to the RSSI where known interference power is present within sectors, for example when using multiple towers and ABAB configurations
- Frequency reuse plan and the environment (propagation exponent) will have a significant impact on expected CINR. If the propagation environment has less clutter from buildings and trees blocking the path between sectors of adjacent towers, the signals will propagate from tower to tower and cause interference between towers.
- Reuse patterns: the recommended reuse pattern is described in section 4, with two sets of two channels used back-to-back in adjacent towers.
  When the reuse pattern ABAB is used in a tower, several techniques can be used to limit the effect of self-interference:

    - Down-tilt can be utilized to minimize same frequency power arriving from one tower sector to another
    - Sector alignment should aim to provide isolation between same frequency sectors by taking advantage of the AP antenna patterns and geography
    - Reducing AP Receive power target to lower STA transmit power
    - Reducing AP transmit power to reduce power arriving at adjacent tower STAs. While this may seem counter-intuitive, lowering the power of APs can actually increase the overall system performance especially on the shorter links that can achieve higher modulations already.

Best Practice Recommendations for ABAB Channel Planning

The following table describes several techniques to minimize the sources of self-interference discussed above.

Table 2 - Best practice recommendations for ABAB channel planning

Interference technique	Downlink Relief	Uplink Relief	Description	Trade-off
AP Down-tilt	Yes	Yes	Reduces the downlink power that arrives at sectors on remote towers	Reduces the maximum range
Lowering AP transmit power	Yes	No	This reduces the downlink power that arrives at remote sectors but must be done consistently on all sectors on all tower	Reduces the maximum range
Lowering AP Receive target	No	Yes	For a sector that is suspected of having STAs causing interference to a remote AP Uplink, lowering the offending STA Transmit power can relief	Not possible for STAs at the edge, as they would lose the connection. Other STAs with use lower Uplink MCS
STA up-tilt to tower on especially short links as required	Yes	Yes	A STA aimed properly at the tower will reduce that STA's transmit power and increase the desired STA receive power	
STAs registered to correct AP sector	Yes	Yes	A STA connected to the wrong sector (especially if it connects to the AP facing the opposite direction on the backside of the tower) will transmit more power and have lower	

			RSSI than if in the right sector	
Avoid high interference channels	Yes	Yes	When selecting the frequencies for RF planning, test available frequencies and choose the ones with lowest interference levels.	
Proper orientation of AP sectors on a tower	Yes	Yes	If sectors on a tower are not tilted or aligned at a exact 90 or 180 degrees from each other as planned, then the desired RSSI levels will be reduced and the required STA transmit powers end up higher.	