Measuring the Effects of WiFi Low Rate Transmissions on the High Rate WiFi Traffic in the Same Frequency

Oguzhan Eren (Original Document Date: December 2017, Updated: October 2018)

Introduction

Due to the nature of the WiFi standards and AP/Client implementations, roaming and automatic rate selection provides endless connection rate possibilities in an enterprise WiFi network with lots of APs and clients.

Backward compatibility is another strong aspect of WiFi networks which should be handled correctly. To provide correct backward compatibility and the correct "coverage" at the same time, it can be seen that many networks easily allow low-rate communications to be able to provide better coverage and backward compatibility.

For example, a regular 802.11ac network with 3x3 MIMO support can provide 1.3Gbps link rate to 3-stream clients and 867Mbps link rate to 2-stream clients.

However, these high data rates, such as 1.3Gbps or 867Mbps, is only available in a very short range, compared to the all RF range of the WiFi cell. So, most probably, many clients in a real world WiFi environment, won't be able to use max rates at all times, instead many low rates will be used.

For example, it is very likely that a communication from a client which has 867Mbps rate could easily be "intercepted" or "lowered" by another client which is transmitting 6Mbps on the very same channel that the high-rate client operates. So, in this case, the high rate client or the AP will be affected by this low rate client, but the question is how much?

This paper explains the effects of low rate transmissions on the high-rate ones, how and why low rate transmissions can affect high-rate Tx in the same channel and how this negative effect can be solved.

What happens in the real world?

In a real-world RF environment, it is possible that a WiFi AP can be installed in an enterprise environment with the hopes that it can provide "modern" speeds but it can be seen that it suffers from low speeds because of:

-a sticky client which is still attached to a distant AP (which belongs to the same organization or not) with a lower rate and trying to download data as much speed as possible.
-a neighbor AP which is installed with default settings with the low rates enabled.

Either the low-rate neighbor AP belongs an organization's network and tries to cover maximum area with low rates enabled, or being it is a neighbor AP which runs on default configuration; these low rate communications affects that organization's clients a lot, if they are running on the same channel, which is a very good possibility in today's crowded WiFi bands. We can also define this low-rate communications as CCI - Co Channel Interference - because they are sharing the same channel with the high-rate clients and affecting them negatively.

What is the Reason Behind This Effect?

802.11 channels are half duplex and there is a "contention" to reach to the medium to be able to transmit over it. So, when somebody is talking, the other clients in the same medium, meaning the same channel/frequency, cannot talk to prevent collision.

So, when a client tries to access to the medium, it first needs to understand if the medium is busy, via energy detect or via NAV (virtual carrier sense). When there are other clients are talking in the same channel, the client which wants to Tx, reads the duration value from the 802.11 MAC header from the transmitting client, waits for that duration without any transmission. This is virtual carrier sense mechanism and, if NAV timer is zero and "energy detect" says there is not RF energy detected on the medium, then the STA can start IFS, then "contention window" and then, if still no transmission on the medium, transmit.

However, it is well-expected that there can be other stations in the channel. This contention, or fight for access to the air, is not limited with the SSID or AP used. Regardless of the SSID or AP, if two stations are using the same channel, that's enough to start contention among them. So, definitely there will be other stations in the air especially when we think about how far WiFi signals can travel in reality.

Figure 1 shows that the modulation is possible with higher rates only within very short distances but the signal goes well far away with low modulation rates.

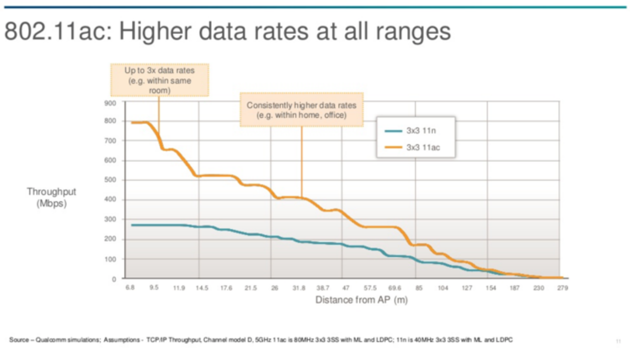

Figure 1: 11ac and 11n connection speeds dramatically drop with range. Source: Qualcomm

According to the Figure 1 values from Qualcomm, 802.11ac high data rates, such as gigabit or nearly gigabit connection rates are possible only within the 10-meter radius of the AP. However, on the other hand, the same signal can travel 300 or more meters with which, the AP easily accepts very low rate connections over those large distances. This means, in a crowded city environment with lots of neighbor access points transmitting, actually the signals are being propagated much more than needed. This is a side effect of high rates; those rates need strong SNR level to decode high QAMs such as 64 or 256bit.

For 11n rates; propagation distance and drop rate behaves the same. It starts with lower speeds but propagates nearly the same, as the blue line represents in Figure 1.

This proves that the low rates connections are possible from very large distances and this communication is being broadcast all over the RF domain with very large number of users are possibly contenting at the same time. Simply, this means that the RF domain is actually too large than expected which makes low rate connections possible and this affects high rate clients badly.

So, when it comes to make a transmit attempt when there is another client is talking and a client is hearing this, it should wait until the end of the communication and then transmit. The wait time is defined by the "duration" value and that value is too large for low rate clients because they are transmitting with a low rate and which means their transmission takes too much airtime to transfer over the air because of the slow rate they are using. So, for example, to transmit 1KB data (as the MPDU payload only, excluding other preambles) over the air takes 10 microseconds for a high rate client, the same amount of data (as the MPDU payload only again) will be taking 1330 microseconds for the low rate client. 10 microseconds versus 1330 microseconds. This is really a big difference. Of course, this calculation does not count the legacy preamble and VHT/HT preambles that every payload part of a frame should be started with.

These LP and V/HT preambles should be sent at the lowest rate, so the real-world difference between low and high rate client is not 133-fold. The real-world values will be explained in the "Test" section.

LAB Setup

The below lab setup is used to test this effect:

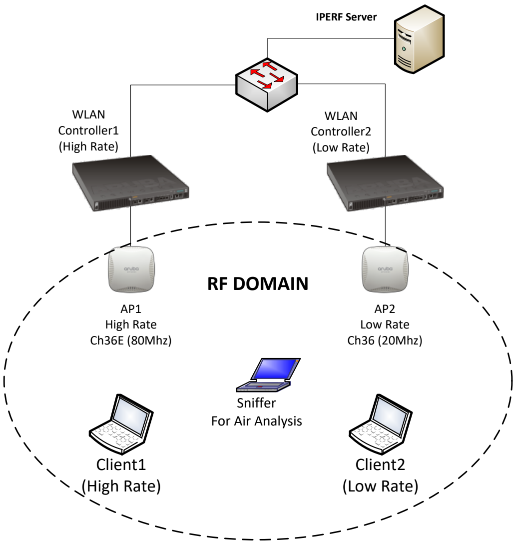

There are two clients for testing purposes. One is to simulate high-rate connection, which is connected to the 802.11ac AP with 867Mbps (MCS9-11ac) at all times.

The other client's connection rate is variable in each test case. Test cases are explained in detail with the respective score for the each one.

Each test is conducted as the high-rate client is downloading traffic with the max speed, the low-rate client is also trying its best to download as much as data it can, with its low-speed connection to another AP in the same channel, at the same time. Both high rate AP, high rate client, low rate AP and low rate client are in the same RF domain, hearing each other within the CCA threshold and NAV (802.11 header in the MPDU) as well.

Test Cases and Results

1- Single Client, Baseline.

In this test, only the high rate client downloaded data from the iPerf server.

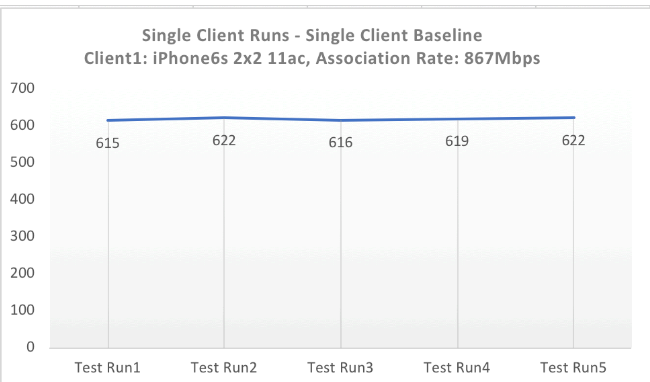

Figure 2: Single Client Baseline Test Runs. Y axis represents throughput, in Mbps.

This data represent the single client baseline, which is around 620Mbps throughput, over 867Mbps association rate.

2- Dual High-Rate Clients, Baseline.

In this test, dual high-rate clients are tested to define the two-client baseline. Two clients were connected as high rate, 867Mbps.

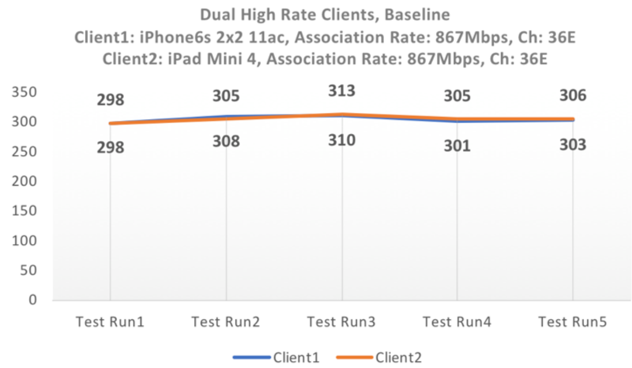

Figure 3: Dual Client Baseline Test Runs. Y axis represents throughput, in Mbps.

As a baseline, it can be said that when there are two high rate clients are downloading traffic at the same time, contending for accessing to the channel and both are connected with 867Mbps rate, each one gets around 300Mbps.

From now on, we will decrease Client2's association rate step by step (867Mbps baseline, 144Mbps as 20Mhz baseline, 115Mbps, 43Mbps, 26Mbps, 13Mbps and as the lowest step 6,5Mbps) and observe how this affects Client1's performance.

3- High-Rate and Low-Rate Concurrent - Low Rate is 144Mbps (11n, MCS15, 20Mhz, SGI)

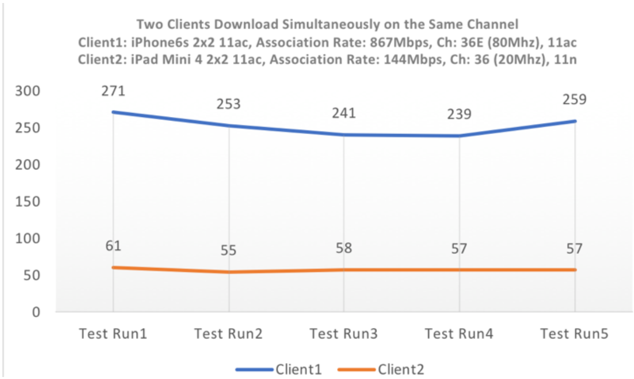

Figure 4: 867 vs 144Mbps Test Runs. Y axis represents throughput, in Mbps.

In this test, 80Mhz and 20Mhz coexistence was tested in general, with max connection rate values in each channel width, both clients were trying to download as much data as possible, at the same time. Client1 is connected as 867Mbps, however, the other client is on Ch36 (20Mhz) attached to an 11n AP with highest possible dual-stream rate which is 144Mbps. SGI was enabled.

As low-rate client (144Mbps) averaged as 57Mbps over the five test runs, the high rate client got the average download speed as 252Mbps, which is **16%** lower than the baseline score which is 300Mbps.

4- High-Rate and Low-Rate Concurrent - Low Rate is 115Mbps (11n, MCS13, 20Mhz, SGI)

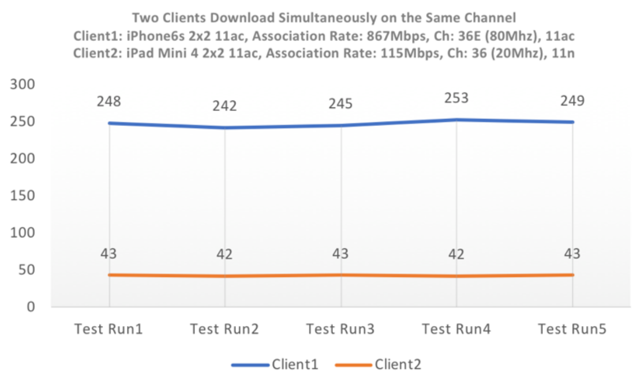

Figure 5: 867 vs 115Mbps Test Runs. Y axis represents throughput, in Mbps.

While the client 2 was occupying the channel with 115Mbps rate, download rate of the high-rate client averaged at 247Mbps.
In this test, the high rate client, despite its high connection rate (867Mbps), lost its **18%** performance compared to baseline score (300Mbps) because of the low rate client which was connected as 115Mbps.

5- High-Rate and Low-Rate Concurrent - Low Rate is 43Mbps (11n, MCS10, 20Mhz, SGI)

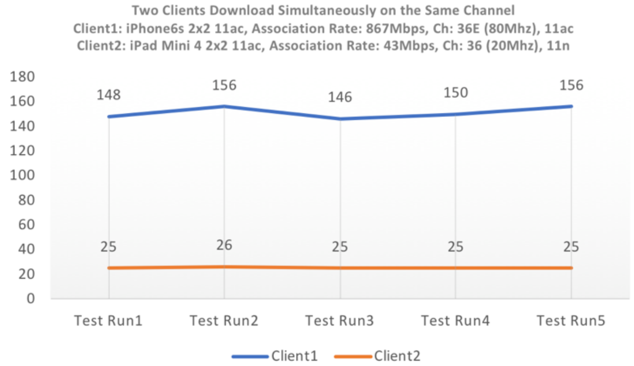

Figure 6: 867 vs 43Mbps Test Runs. Y axis represents throughput, in Mbps.

When the low-rate client was connected with 43Mbps rate, the performance drop on the high-rate client was **50%**, as it averaged as around 150Mbps over the dual client baseline of 300Mbps.

6- High-Rate and Low-Rate Concurrent - Low Rate is 26Mbps (11n, MCS3, 20Mhz, LGI)

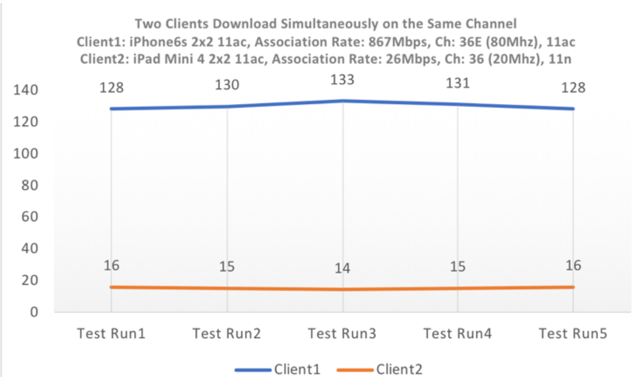

Figure 7: 867 vs 26Mbps Test Runs. Y axis represents throughput, in Mbps.

In this test, low-rate client's connection speed was 26Mbps and the performance drop on the high-rate client was **57%** as it is averaged as 130Mbps.

7- High-Rate and Low-Rate Concurrent - Low Rate is 13Mbps (11n, MCS1, 20Mhz, LGI)

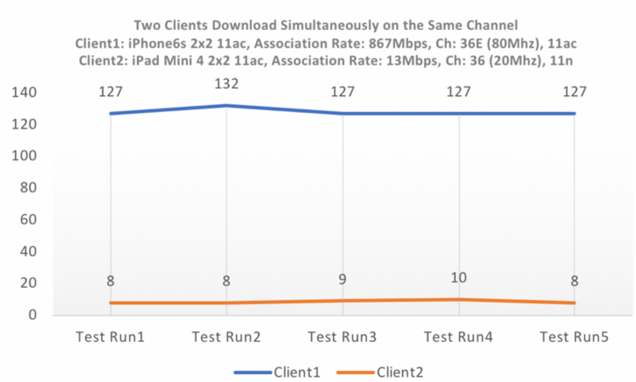

Figure 8: 867 vs 13Mbps Test Runs. Y axis represents throughput, in Mbps.

In this test, the low-rate client connected as 13Mbps and the performance drop on the high-rate client was **58%**, with the average score of 128Mbps over the five test runs.

8- High-Rate and Low-Rate Concurrent - Low Rate is 6.5Mbps (11n, MCS0, 20Mhz, LGI)

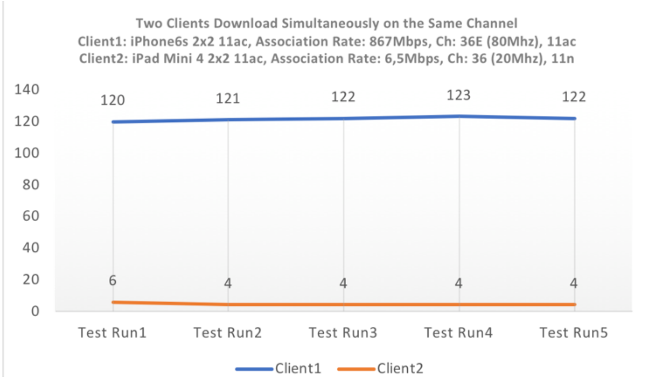

Figure 9: 867 vs 6,5Mbps Test Runs. Y axis represents throughput, in Mbps.

In this last test, the low-rate client connected with 6,5Mbps data rate and pulled data as much as possible while the higher-rate client was also downloading at the same time, within the same RF domain and they were on the same channel (high rate client was on 80Mhz channel, low-rate client was on 20Mhz channel, both channels' base 20Mhz subchannel was the same: 36). With this test, the performance drop, due to the low rate client, was measured as **60%** on the high-rate client.

Evaluating the Test Results

As the baseline set at ~300Mbps and ~300Mbps for each high-speed client during the "Dual High Rate Client Baseline Test", it is clearly seen that there were different performance drops when one of the clients are low-rate, instead of high-rate.

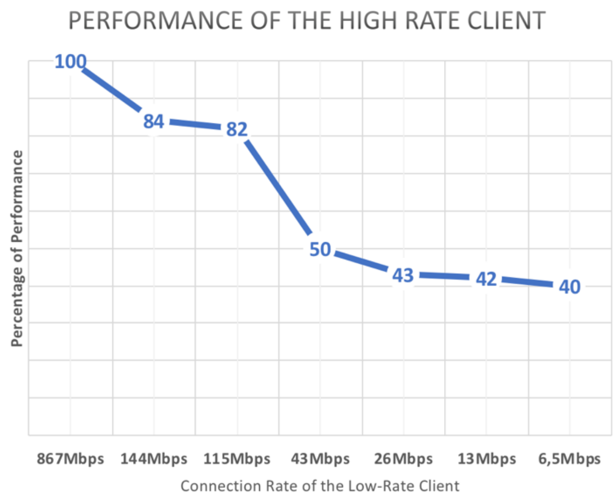

Figure 10: Performance effect of each low-rate Tx rate, on the high-rate client.

According to the tests conducted with different low-rate values, Figure 10 above summarizes the scores. The baseline means that the neighboring (low-rate or secondary) client has max data rate, so the first client's data rate is measured as around 300Mbps on average and this value is set as the baseline and represented as "100" in the Figure 10.

When the second "low-rate" client runs at 144Mbps, high rate is affected by 16% performance drop and this performance loss is represented as "84 performance over 100" in Figure 10.

In the lowest end, if the low-rate client runs at 6.5Mbps, the performance drop is measured as **60%. This is a very significant result that only <u>one</u> low-rate client can affect and decrease the total performance of a neighbor channel by 60 percent. Big hit.**

This type of situation can easily be seen in the real life, when there are many clients are "fighting" for the channel access in the same RF domain. However, it is highly possible to see that the term "RF Domain" can easily be underestimated.

In a city environment, a 5Ghz AP or a basic home modem which supports 5Ghz, can be installed near a window facing out to the city view, broadcasting signals to all over the city. How far signals can go?

Before answering how far signals can go, we need to understand what is the level of RF signal at the Rx side which makes it wait. This is defined as PD - Preamble Detect - threshold and it is -82dbm, according to the standard. However, it is the highest signal level that a receiver must wait if there is signal (if the legacy and/or v/ht preambles can be decoded) higher than -82 dbm. However, many "modern" enterprise APs can detect preambles and payload data even with lower signal levels. For example, Aruba's mid level AP315 access point can decode HT20 6Mbps signal with an only -90dbm which is actually very weak. However, this means that the AP will "wait" if there are signals in the air and their energy level at the AP Rx is merely -90dbm.

Now, the above question, which was "How far signals can go?" can be enhanced: "How far signals can go and make an AP wait?"

So, if a signal is sent from an AP which is open to the air and it is broadcast 5Ghz signal via 20dbm EIRP, the signal can reach 100 meters with -66dbm, 500 meters with -80.7dbm, 1000 meters with -86.7dbm and 1450 meters with -90dbm. Theoretically, this AP in this example, can make another AP wait which is 1450 meters away, under the completely open-air conditions. This means that an AP can easily make others wait (in the same channel) within a very large area.

The above "distance vs dbm" can easily be calculated with a "free space loss calculator".

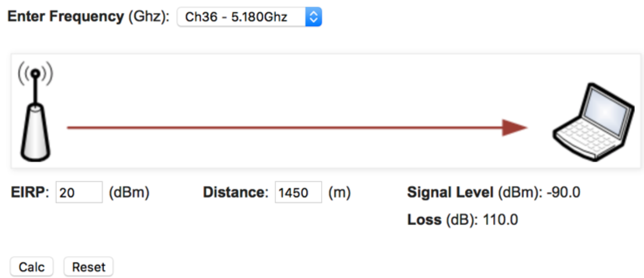

Figure11: Free Space Path Loss Calculator, Source: wifinigel.blogspot.com

If the same AP can transmit signals with a low rate (because of distance, obstacles, client interference and lowering rate etc.) to its clients, the same AP can easily affect negatively many high-rate clients in a very large area. Many high rate clients will "hear" this traffic, even hundred meters away and see significantly low throughput because of this low-rate APs and clients. The effect rates and amounts are explained above.

Potential Solutions to the Problem

AP Rx Fine Tuning Possibilities; CSR and Channel Quality Analysis

When an AP hears one of more low-rate transmissions within the same RF domain and the transmission is coming from far away APs, this means that the low-rate transmission is coming outside of the organization, hence organization's IT department has no control over it.

In this case, according to the channel access rules of WiFi and also the "unlicensed" nature of the overall WiFi operation, it is not possible to find the exact Tx point of low-rate transmission and ask for a fine tuning or complete shutdown. Rather, the affected organization should accept this interference and deal with it.

There are some ways in enterprise vendors to deal with it. The most classical solution to this problem is to have an automatic channel selection algorithm in the WLAN infrastructure which automatically calculates the quality of all channels and select the best channel accordingly. This way, it may be possible to select another -better- channel to leave the channel which is occupied with low-rate transmissions and overall quality is decreased a lot.

Another infrastructure-based solution to these distant and low-rate APs could be "Cell Size Reduction" techniques possible with some vendors.

Cell Size Reduction (CSR) means that AP can artificially lower its PD sensitivity level, to not to hear weak signals. So, that means AP won't assert CCA medium busy because it will ignore the preambles coming from distant clients.

For example, based on the Aruba's AP315 example above, it is possible for AP315 to decode 6Mbps preambles and also payloads which are decoded with -90dbm signal level. However, with CSR setting, this sensitivity level (-90dbm for instance) can be set as -80dbm. Which means that AP won't hear (actually it will hear but ignore) all transmissions which are weaker than -80dbm. Even if there are transmissions hitting to the AP with, for example, -81dbm; AP will ignore them and will not go into "waiting" mode. Rather, the AP will go ahead and start Tx sequence as if CCA is idle.

In the AP315 example, it is stated above that a distant AP which is 1450 meters away from an AP315 (completely open environment, full LoS) can affect AP315 and can make it "wait". What is the current "affecting range" now, after the sensitivity is manually adjusted to -80dbm, instead of -90dbm?

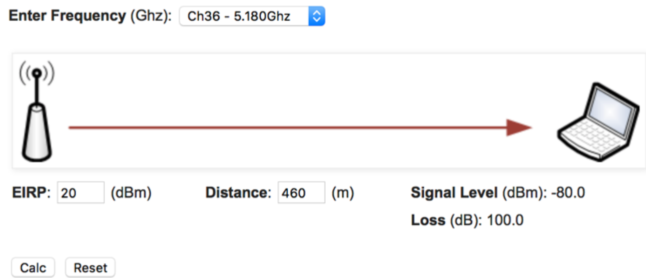

Figure12: Free space path loss after 460meters, 20dbm EIRP, 5180Mhz.

The same free space path loss calculator shows that the -80dbm radius is now 460 meters. So, with the help of CSR, the affecting range is decreased nearly 1000 meters, in an open space. This means a huge "shrink" in the cell size and now AP is more protected against distant low-rate APs. However, this setting only affects APs, not the clients. If clients can hear distant APs, they can also go into medium-busy state, which is actually unlikely for clients to hear better than APs.

Infrastructure Based Tunings If the Low-rate Clients Are From "Inside"

Other than "CSR" and channel quality-based channel selection algorithms explained above, mainly if the low-rate AP or APs are inside the organization and belongs to the same WLAN infrastructure with the high-rate clients; then it is possible to use sticky client mitigation algorithms offered by some vendors.

It is possible for a WLAN controller to take an action when there is a low-rate client connected to a distant AP (belongs to the same organization) while there are other potentially good-rate APs nearby. In this case, WLAN controller can either use 802.11v based BSS transition for clients which support this, or make deauth to "steer" those clients to "better" APs to which clients will connect with high-rates and do not lower the performance of the RF domain which actually consists of many APs and many clients.

In addition to the above, some vendors also offer airtime fairness algorithms (limiting the airtime further for the distant client) for this type of problems. The efficiency of these algorithms should be measured in another controlled study.

Trimming Low Rates

It is widely accepted that low rates should be removed to get better WiFi performance. Does this solve the low-rate client problem?

Before explaining if this solves the problem or not, it should be explained that which rates are trimmed mostly. Figure 13 is a screenshot of a default SSID configuration with all rates enabled:

```
Parameter                               Value
---------                               -----
SSID enable                             Enabled
ESSID                                   aruba-ap
Encryption                              opensystem
Enable Management Frame Protection      Disabled
Require Management Frame Protection     Disabled
DTIM Interval                           1 beacon periods
802.11a Basic Rates                     6 12 24
802.11a Transmit Rates                  6 9 12 18 24 36 48 54
802.11g Basic Rates                     1 2
802.11g Transmit Rates                  1 2 5 6 9 11 12 18 24 36 48 54
Station Ageout Time                     1000 sec
Max Transmit Attempts                   8
RTS Threshold                           2333 bytes
```

Figure13: Aruba's default parameters for an SSID profile.

Figure 14 shows the recommendations from Aruba's VRD called "RF and Roaming Optimization for Aruba 802.11ac Networks":

Feature	Default Value	Recommended Value
Data rates (Mbps)	**802.11a**: Basic rates: 6,12,24 Transmit Rates: 6,9,12,18,24,36,48,56	802.11 a & g: Basic rates: 12,24 802.11 a & g transmit rates: 12,24,36,48,56

Figure14: SSID low rates removal recommendations according to Aruba.

According to these recommendations and removing of low rates, the below item is expected after the change. This is the positive point about trimming low rates.

-If there is a client connects to an AP which belongs to the organization or can be managed somehow, this AP won't accept 6.5Mbps clients and the "867Mbps vs 6.5Mbps" test will be invalid. Therefore, it will be possible to avoid 60% performance drop.

However, the below items, on the other hand, are the negative points about trimming low rates:

- Even after the low rates are removed from the organization's APs correctly, only 6 and 9Mbps are removed. So, it will still be possible to see 13Mbps clients which leads to **58% performance drop on the high-rate client(s).**

-Some very high-density recommendation guides advise removing more rates from the bottom end, and start with 24Mbps as the lowest rate. However, this also makes 26Mbps client connections possible which also affects high-rate clients badly, measured as **57% performance drop for the high-rate client(s).**

When it comes to measuring airtime consumption of two users, high-rate and low-rate (and, given that control traffic rates, MPDU sizes and MPDU per A-MPDU numbers are configured as equal and only payload rates are different as 867Mbps versus 26Mbps) airtime utilization of two scenarios are calculated as seen in Figure 15 below:

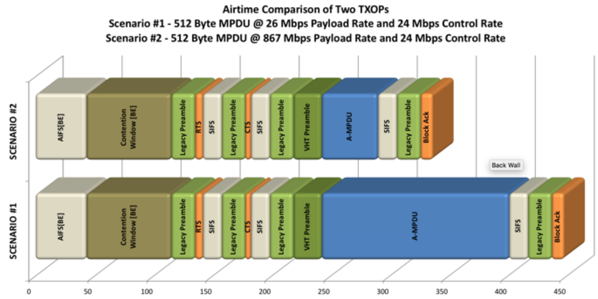

Figure15: Comparison of the time duration, needed to transmit the same payload with different data rates.

According to the Figure 15, created by Aruba's WiFi Frame Time Calculator, shows that 867Mbps traffic would need to use air for 47,6 microseconds but 26Mbps traffic would need 158.8 microseconds for the same amount of data to be sent over the air. This shows the efficiency at using air medium with different payload rates. So, trimming all rates below 24Mbps could still lead to the situation in Figure 15, meaning inefficient usage of the medium.

On the other hand; if the low-rate AP cannot be reached for management (for example, it doesn't belong to the corporate, but belongs to anybody else in the neighbourhood and "bleeding" into corporate's RF environment) it still won't be possible to trim the low rates, at all.

Another point about removing low rates is that actually this configuration does not "shrink" the size of the cell and does not protects the cell from interference, because even if the APs are configured with, for example 24Mbps minimum rate, they are still using BPSK - 6Mbps rates for all preambles (legacy and/or v/ht) and this is a requirement by the standard. So, even if all APs in an environment set as 24Mbps minimum rate, their effective PD interference range still equals to 6Mbps range because of the preambles they send with 6Mbps. So, trimming low rates does not affect cell size in terms of interference.

It can be said that trimming low rates effectively provides better roaming decisions for clients and fairly lower airtime utilization (due to the fact that the new "lowest" rate for some frames will be higher, so they will be sent faster); however, it does not protect APs from performance drops due to low-rate clients.

Summary and Conclusion

With this test, it is possible to validate the effect of a low-rate client on a high-rate client(s) in the same WiFi channel, can be seen as high as 60%. This is due to the fact that WiFi propagates more than expected (because of improved Rx sensitivity in the modern AP platforms) and a neighbor AP can easily make many APs and Clients wait, which are on the same channel. So, CCI (PD-CCA) radius is very large.

In addition to the large CCI radius, the low-rate clients consume more air to transfer the same amount of data, because of their low-rate, which means there will be less airtime for the high-rate clients in the "same air".

Even though low rates are trimmed aggressively (no rates allowed below 24Mbps), this again means that 57% performance drop for high-rate clients.

This problem can be solved partially with cell size reduction measures available in some vendors, if the transmitting low-rate STAs' signal at the receiving AP is weak, compared to valid high-rate clients. So, APs can choose "not to hear" those weak signals.

It is also possible to solve the problem with the help of active sticky client re-connecting techniques if the low-rate clients are also connected to the same organization's APs and can be manageable via 802.11v BSS transition or deauth moves.

References

- RF and Roaming Optimization for Aruba 802.11ac Networks
- CWAP Certified Wireless Analysis Professional Official Study Guide: Exam PW0-270, 1st Edition.
- Validated Reference Design, Very High-Density 802.11ac Networks Theory Guide Version 1.0
- Free Space Loss Calculator http://wifinigel.blogspot.com
- The WiFi Evolution: An Integral Part of the Wireless Landscape, Qualcomm.
- Aruba WiFi Frame Time Calculator - Very High-Density 802.11ac Networks Planning Guide Version 1.0