

# Analyzing Frame Utilization and Spatial Utilization Statistics in the PMP 450m AP

The PMP 450m access point (AP) supports MU-MIMO operation, which changes the way data is scheduled and utilization is reported in the AP GUI.

## Scheduler in the PMP 450m AP

The PMP 450m AP schedules user data in MU-MIMO or SU-MIMO mode dynamically in every frame, adjusting to the changing traffic conditions.

If there is groupable traffic in the data queues, the scheduler tries to schedule as many subscriber modules (SM) as possible, up to the platform limit, which is seven for the 5 GHz version and four for the 3 GHz version.

The group uses as many symbols as necessary to transfer the data of the SMs in the group. If the group does not fill the frame, additional data is sent in SU-MIMO mode. The ratio of MU-MIMO traffic and SU-MIMO traffic changes frame by frame and it depends on the traffic conditions. Figure 1 shows examples of data traffic in two frames.

If a sector has very little traffic, data does not queue up long enough for the scheduler to form large groups. Most data is sent in SU-MIMO mode as soon as it arrives at the AP.

As traffic increases, the AP cannot schedule data as soon as it arrives in the queues, and the queues fill up. With more data in the queues, the AP can schedule more data in MU-MIMO mode, which means both more SMs per group and more symbols for the group.

As explained below, it is possible that even when there is data waiting in the queues, the scheduler cannot form groups if the SMs do not have enough spatial separation between them. In this case, the groups will have fewer SMs, or data is transmitted in SU-MIMO mode.

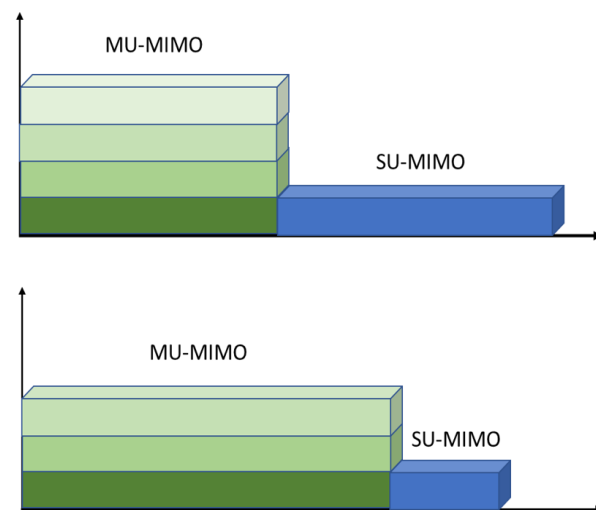


Figure 1 - Examples of frame scheduling with 450m

In order to understand how the frame is utilized, the user can refer to the Frame Utilization statistics and the Spatial Utilization statistics in the PMP 450m AP GUI.

## Frame Utilization Statistics

The Frame Utilization tab under Statistics shows how the frame is utilized both in time and in the grouping dimension.

Let us consider the example in Figure 2. This example is taken from a 5 GHz PMP 450m, where up to seven SMs can be grouped. The 3 GHz PMP 450m AP can group up to four SMs, and its Frame Utilization statistics are modified accordingly.

MU-MIMO Utilization		
Slot Grouping		
Group Size	% Downlink Distribution	% Uplink Distribution
1 (ungrouped)	34.9	11.5
2	0.3	0.0
3	0.3	88.5
4	0.5	0.0
5	3.2	0.0
6	32.3	0.0
7	28.7	0.0
Additional Statistics		
	Downlink	Uplink
Average MU-MIMO Group Size - Data	4.5	2.8
Multiplexing Gain	2.9	2.2
Sector Utilization		
	Downlink	Uplink
SU-MIMO	31%	8%
MU-MIMO	57%	58%
ACK	3%	29%
MU-MIMO ACK		5%
Broadcast & Multicast	0%	
Total Utilization	91%	100%

Figure 2 - Example of frame utilization statistics

## Slot Grouping Table

In this example, when groups are formed, they have a distribution of the group size as shown in the Slot Grouping table. In the downlink direction, 34.9% of the slots are transmitted in SU-MIMO mode (equivalent to groups of 1), 0.3% of the slots are transmitted in MU-MIMO mode with 2 SMs per group, 0.3% of the slots are transmitted in MU-MIMO mode with 3 SMs per group, and so on.

Note that 100% here corresponds to the total time used to send user data. It does not include, for example, time used to send acknowledgments or broadcast and multicast traffic, or unused symbols.

## Additional Statistics Table

The average group size of the MU-MIMO transmissions in the downlink direction is:

$$34.9\% \times 1 + 0.3\% \times 2 + 0.3\% \times 3 + 0.5\% \times 4 + 3.2\% \times 5 + 32.3\% \times 6 + 28.7\% \times 7 = 4.5$$

Similarly, the average group size of the MU-MIMO transmission in the uplink direction is:

$$11.5\% \times 1 + 0\% \times 2 + 88.5\% \times 3 + 0\% \times 4 + 0\% \times 5 + 0\% \times 6 + 0\% \times 7 = 2.8$$

These two values are shown in the Additional Statistics table, in the Average MU-MIMO Group Size – Data row.

The multiplexing gain, which is the next row in the same Additional Statistics table, uses the same information used to calculate the average group size but with more granularity.

When a group is formed, the length in time (number of symbols) allocated to the group is determined by the SM in the group with the most data to send. If other SMs do not have enough data to send to fill the symbols allocated to the group, they repeat their data to increase the probability of correct reception at the receiver. The resources are still filled by all SMs, but when SMs do not send new data and repeat the data, the corresponding symbol is not counted in the multiplexing gain calculation, as it does not correspond to net new throughput.

Let us look for example at the uplink direction. 11.5% of the time data is sent in SU-MIMO mode; 88.5% of the time data is sent in MU-MIMO mode with a group of 3 SMs. However, not all SMs are equally active.



Figure 3 - Traffic load in uplink group

Figure 3 shows that the 88.5% of the time allocated to the group of 3 SMs is not fully utilized by SM 2 and SM 3. SM 2 is transmitting new data for 70% of the time, while SM 3 is transmitting new data for 51% of the time. For the remaining of the time within the group, SM 2 and SM 3 repeat some of the data they have sent during this uplink transmission. Note that here, 100% corresponds to the total time allocated to user data transmission only.

To calculate the multiplexing gain the repeated data is not counted. The uplink multiplexing gain in this case is:

$$11.5\% \times 1 + (88.5\% + 70\% + 51\%) / 3 \times 3 = 2.2$$

Similarly, the downlink multiplexing gain is calculated as 2.9, lower than the average group size of 4.5.

When the multiplexing gain is significantly lower than the average group size it typically means that some of the SMs in the group have much more traffic than the other SMs in the group, and the group time is not fully utilized by some SMs.

## Sector Utilization Table

The last table in the MU-MIMO Utilization section is the Sector Utilization table. This table shows the whole Frame Utilization, not just the portion used for user data transmission, like the two previous tables.

In this example, the downlink is 91% utilized. 31% of the downlink is used for SU-MIMO transmissions, 57% of the downlink is used for MU-MIMO transmissions, and 3% of the downlink is used for ACK transmission. The uplink is 100% utilized, and the table shows the breakdown of the types of uplink utilization.

Sector Utilization		
	Downlink	Uplink
SU-MIMO	31%	8%
MU-MIMO	57%	58%
ACK	3%	29%
MU-MIMO ACK		5%
Broadcast & Multicast	0%	
Total Utilization	91%	100%

Figure 4 - Sector utilization

## How to Understand if a Sector Is Reaching Capacity

When trying to understand if a sector is reaching its maximum capacity, the first parameter to look at is the Frame Utilization. In a case without MU-MIMO support, like in a 450 or 450i AP, the Frame Utilization parameter reflects the percentage of capacity used by the sector.

However, because the 450m AP supports MU-MIMO, this parameter by itself is not sufficient to get the full picture. In this example, the Frame Utilization is very high in the downlink and at the maximum level in the uplink. But the average group sizes and the multiplexing gains show that there is additional capacity in the sector, especially in the uplink direction where the maximum group size is 3, while up to 7 SMs could be grouped.

Let us review the information given by the MU-MIMO utilization tables.

First, we look at the Total Utilization. If this number is low, then there is extra capacity in the sector. If this number is high, then this number alone is not sufficient to determine if the sector is at capacity or not. As the scheduler schedules both over time and over the grouping dimension, it is possible that a sector with high total utilization could have plenty of capacity left in the grouping dimension.

Second, we look at the average group size. If this number is low, then there is extra capacity in the sector. If this number is high, it is an indication that the sector is reaching its capacity.

More granularity is given by the multiplexing gain. A sector with a high average group size but a low multiplexing gain indicates that there is more capacity in the sector, but it may be difficult to use the extra capacity as some SMs are much more active than others, and they tend to occupy more of the time resources.

Assume now that a sector has a high Total Utilization but a low average group size. We know there is extra capacity in the sector, but can this extra capacity be used? In other words, can the scheduler schedule more data in the unused resources?

The answer depends on the spatial distribution of the traffic, which is shown in the Spatial Utilization tab under Statistics.

## Spatial Utilization Statistics

The Instantaneous Distribution and Spatial Utilization tables in this tab show the most active SMs, and their corresponding spatial distribution.

### Spatial Utilization Table

Let us look at the Spatial Utilization table first.

Spatial Utilization						
Spatial Frequency	LUIDs in Range	Downlink				Azimuth (degrees)
		Instantaneous (%)	Total (%)	Max (%)	Min (%)	
0 - 31	2	100	100	100	99	(-50.7 - -49.3) , (0.0 - 1.4) , (50.7 - 52.0)
32 - 63	5	100	100	100	99	(-49.3 - -47.9) , (1.4 - 2.8) , (52.0 - 53.4)
64 - 95	9	100	100	100	99	(-47.9 - -46.5) , (2.8 - 4.2) , (53.4 - 54.8)
96 - 127		100	100	100	99	(-46.5 - -45.1) , (4.2 - 5.5) , (54.8 - 56.2)
128 - 159	14	100	100	100	99	(-45.1 - -43.7) , (5.5 - 6.9) , (56.2 - 57.6)
160 - 191		100	100	100	99	(-43.7 - -42.3) , (6.9 - 8.3) , (57.6 - 59.0)
192 - 223	8	100	100	100	99	(-42.3 - -40.9) , (8.3 - 9.7) , (59.0 - 60.0)
224 - 255	19	100	100	100	99	(-40.9 - -39.5) , (9.7 - 11.1)
256 - 287	11	100	100	100	99	(-39.5 - -38.1) , (11.1 - 12.6)
288 - 319		100	100	100	99	(-38.1 - -36.7) , (12.6 - 14.0)
320 - 351	22	100	100	100	99	(-36.7 - -35.2) , (14.0 - 15.4)
352 - 383	18	100	100	100	99	(-35.2 - -33.8) , (15.4 - 16.9)
384 - 415		100	100	100	99	(-33.8 - -32.3) , (16.9 - 18.3)
416 - 447	3	100	100	100	99	(-32.3 - -30.9) , (18.3 - 19.8)
448 - 479	6	100	100	100	99	(-30.9 - -29.4) , (19.8 - 21.3)
480 - 511		100	100	100	99	(-29.4 - -27.9) , (21.3 - 22.7)
512 - 543	4	100	100	100	99	(-27.9 - -26.4) , (22.7 - 24.3)
544 - 575		100	100	100	99	(-26.4 - -24.9) , (24.3 - 25.8)
576 - 607	17	100	100	100	99	(-24.9 - -23.3) , (25.8 - 27.3)
608 - 639	10	100	100	100	99	(-23.3 - -21.8) , (27.3 - 28.9)
640 - 671		100	100	100	99	(-21.8 - -20.2) , (28.9 - 30.5)
672 - 703	7	100	100	100	99	(-20.2 - -18.5) , (30.5 - 32.1)
704 - 735	20	100	100	100	99	(-18.5 - -16.9) , (32.1 - 33.8)
736 - 767		100	100	100	99	(-16.9 - -15.2) , (33.8 - 35.5)
768 - 799	12	100	100	100	99	(-15.2 - -13.5) , (35.5 - 37.2)
800 - 831		100	100	100	99	(-13.5 - -11.7) , (37.2 - 38.9)
832 - 863	16	100	100	100	99	(-11.7 - -9.9) , (38.9 - 40.7)
864 - 895	13	100	100	100	99	(-60.0 - -58.7) , (-9.9 - -8.1) , (40.7 - 42.6)
896 - 927	21	100	100	100	99	(-58.7 - -56.8) , (-8.1 - -6.2) , (42.6 - 44.5)
928 - 959		100	100	100	99	(-56.8 - -54.8) , (-6.2 - -4.2) , (44.5 - 46.5)
960 - 991	15	100	100	100	99	(-54.8 - -52.8) , (-4.2 - -2.1) , (46.5 - 48.5)
992 - 1023		100	100	100	99	(-52.8 - -50.7) , (-2.1 - 0.0) , (48.5 - 50.7)

Figure 5 - Spatial Utilization Table

This table shows 32 rows (bins), each corresponding to 32 Spatial Frequencies (SFs).

For each row, the table shows:

- Spatial Frequency: Range of 32 SFs for this row.
- LUIDs in range: List of LUIDs with spatial frequency falling in the bin.
- Instantaneous (%): Frame utilization for the SF bin updated every 500 ms. The frame utilization percentage accounts for all traffic, sector mode, beamforming mode and MU-MIMO mode.
- Total (%): Average utilization in the SF bin for the past 1/5/15 minutes, as selected in the Statistics Display interval.
- Max (%): Maximum instantaneous utilization in the 1/5/15 minute interval.
- Min (%): Minimum instantaneous utilization in the 1/5/15 minute interval.
- Azimuth (degrees): Azimuth range in degrees corresponding to the spatial frequencies of the bin. The zero-degree azimuth is boresight.

Note that some SF ranges correspond to multiple azimuth ranges. This is because for some spatial frequencies the AP generates beams in multiple azimuth directions. The SM can be physically located in any of the azimuth ranges.

## Instantaneous Distribution Table

The color scheme used in the LUIDs in Range column is the same used in the Instantaneous Distribution table shown above the Spatial Utilization Table.

Instantaneous Distribution		
Group	Downlink	
	Median Slot Count	LUIDs in Group
Top LUID	3584	19
Top 2 LUIDs	3584	18
Top 4 LUIDs	3584	15, 16
Top 8 LUIDs	3584	4, 5, 7, 9
Top 16 LUIDs	3428	2, 3, 6, 8, 10, 12, 21, 22
Top 32 LUIDs	3368	11, 13, 14, 17, 20
Top 64 LUIDs	0	
Top 128 LUIDs	0	
Top 256 LUIDs	0	

Figure 6 - Instantaneous Distribution Table

The most active SM is the SM with LUID 19. It is shown in yellow in the Spatial Utilization table and in the top row in the Instantaneous Distribution table. The next most active SM is the SM with LUID 18; then SMs with LUIDs 15 and 16, and so on.

The Spatial Utilization table in this example shows that:

- SMs are well distributed in spatial frequency: there are no SF bins with many more SMs compared to the other SF bins.
- The active SMs are well distributed in spatial frequency: the most active SMs are not in the same bin nor are they in adjacent bins.

Looking at the utilization percentages, all bins show a utilization of 100%, which means that this sector has reached its maximum capacity and additional traffic cannot be scheduled.

However, the tables above show traffic over a limited amount of time: 1, 5 or 15 minutes, based on the configuration. To have a better view of what happens over a longer period of time, the Spatial Utilization Heatmaps can be used. They capture the same data of the Spatial Utilization table, but they display it for the last hour.

## Spatial Utilization Heatmap

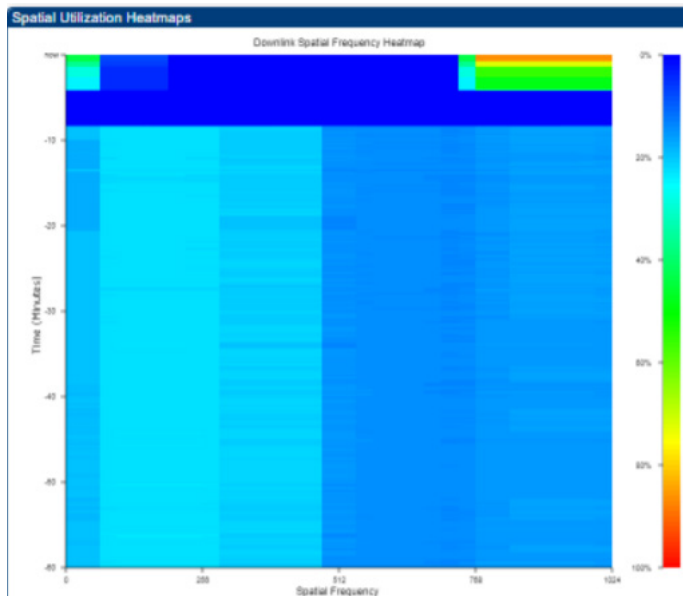


Figure 7 - Downlink Spatial Utilization Heatmap

Figure 7 shows the downlink Spatial Utilization Heatmap. In this heatmap, time is shown from top to bottom, where the topmost row is the latest utilization sample, while the bottom row is the utilization sample from one hour ago.

The SF bins are shown left to right, where the first bin starts with SF 0, which corresponds to boresight.

The color used in the heatmap is shown in the legend to the right. Bins with low utilization are shown with shades of blue and green, while bins with high utilization are shown with shades of yellow and red.

In this example, for most of the last hour, traffic was very light in every direction. In the last several minutes, some SMs became active, mostly in the last bins to the right and the very first bins to the left.

This heatmap shows there is still a lot of capacity, even when there is traffic, traffic is not filling the grouping dimension, and it is not constant over time.

In Figures 8 and 9 below, let us consider an example where the spatial utilization and heatmap look like this:

Spatial Frequency	LUIDs in Range	Downlink				Azimuth (degrees)
		Instantaneous (%)	Total (%)	Max (%)	Min (%)	
0 - 31		0	1	0	0	(-50.7 - -48.5), (0.0 - 1.4), (50.7 - 52.9)
32 - 63	20	0	1	0	0	(-48.5 - -46.5), (1.4 - 2.8), (52.9 - 55.3)
64 - 95	13	0	1	0	0	(-48.5 - -44.5), (2.8 - 4.2), (55.3 - 57.8)
96 - 127		0	1	0	0	(-44.5 - -42.6), (4.2 - 5.5), (57.8 - 60.0)
128 - 159	18	0	1	0	0	(-42.6 - -40.7), (5.5 - 6.9)
160 - 191		0	1	0	0	(-40.7 - -38.9), (6.9 - 8.3)
192 - 223	7	0	1	0	0	(-38.9 - -37.2), (8.3 - 9.7)
224 - 255	6	0	1	0	0	(-37.2 - -35.5), (9.7 - 11.1)
256 - 287	17	0	1	0	0	(-35.5 - -33.8), (11.1 - 12.6)
288 - 319		0	1	0	0	(-33.8 - -32.1), (12.6 - 14.0)
320 - 351		0	1	0	0	(-32.1 - -30.5), (14.0 - 15.4)
352 - 383	3	0	1	0	0	(-30.5 - -28.9), (15.4 - 16.9)
384 - 415		0	1	0	0	(-28.9 - -27.3), (16.9 - 18.3)
416 - 447	24	0	1	0	0	(-27.3 - -25.8), (18.3 - 19.8)
448 - 479	5	0	1	0	0	(-25.8 - -24.3), (19.8 - 21.3)
480 - 511		0	1	0	0	(-24.3 - -22.7), (21.3 - 22.7)
512 - 543	16	82	83	91	75	(-22.7 - -21.3), (22.7 - 24.3)
544 - 575	19	82	83	91	75	(-21.3 - -19.8), (24.3 - 25.8)
576 - 607		82	83	91	75	(-19.8 - -18.3), (25.8 - 27.3)
608 - 639	22	82	83	91	75	(-18.3 - -16.9), (27.3 - 28.9)
640 - 671	8	82	83	91	75	(-16.9 - -15.4), (28.9 - 30.5)
672 - 703		82	83	91	75	(-15.4 - -14.0), (30.5 - 32.1)
704 - 735	2	82	83	91	75	(-14.0 - -12.6), (32.1 - 33.8)
736 - 767		82	83	91	75	(-12.6 - -11.1), (33.8 - 35.5)
768 - 799	10	82	83	91	75	(-11.1 - -9.7), (35.5 - 37.2)
800 - 831	21	0	1	0	0	(-9.7 - -8.3), (37.2 - 38.9)
832 - 863		0	1	0	0	(-8.3 - -6.9), (38.9 - 40.7)
864 - 895	11	0	1	0	0	(-6.9 - -5.5), (40.7 - 42.6)
896 - 927	14	0	1	0	0	(-6.0 - -5.7), (-5.5 - -4.2), (42.6 - 44.5)
928 - 959		0	1	0	0	(-5.7 - -5.3), (-4.2 - -2.8), (44.5 - 46.5)
960 - 991	9	0	1	0	0	(-5.3 - -5.2), (-2.8 - -1.4), (46.5 - 48.5)
992 - 1023	4	0	1	0	0	(-5.2 - -5.0), (-1.4 - 0.0), (48.5 - 50.7)

Figure 8 - Example with traffic mostly in one direction - spatial utilization table

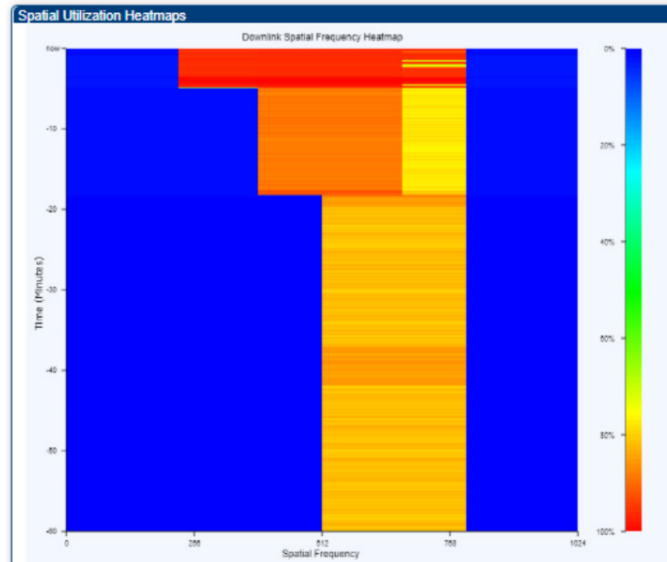


Figure 9 - Example with traffic mostly in one direction - heatmap

In this example, there is one SM that is very active, the SM with LUID 8. The heatmap shows that the bins adjacent to the location of this SM have very high utilization (yellow and red colors). The heatmap also shows that at one point, a second SM became active and then a third. The bins affected by traffic of the additional SMs also show very high utilization in the top portion of the heatmap, which corresponds to the most recent data.

Here, only three SMs are active, but the same effect can happen if multiple SMs are active, and they are all located in the same direction.

In a scenario like this, capacity is available but not in the congested direction. If other SMs in this direction become active, the scheduler will have to reduce the throughput on existing SMs in this direction to allow scheduling of some of the traffic on the newly active SMs. The total throughput in this direction cannot increase. However, if SMs in other directions become active, then the scheduler would be able to group them with the SMs that are already active, and the sector throughput would increase.

## Example of Non-Evenly Distributed SMs

Figures 10 and 11 show the Spatial Utilization table and the heatmap corresponding to a deployment where SMs are not evenly distributed across the SFs.

Spatial Frequency	LUIDs in Range	Downlink				Azimuth (degrees)
		Instantaneous (%)	Total (%)	Max (%)	Min (%)	
0 - 31	18, 19, 52, 67	22	29	73	7	(-39.4 - -38.0), (0.0 - 1.1), (39.4 - 40.9)
32 - 63	4, 38, 53	22	29	73	8	(-38.0 - -36.6), (1.1 - 2.3), (40.9 - 42.5)
64 - 95		22	31	73	8	(-36.6 - -35.2), (2.3 - 3.4), (42.5 - 44.0)
96 - 127	15, 36, 42, 44, 61, 70	22	31	73	8	(-35.2 - -33.8), (3.4 - 4.6), (44.0 - 45.6)
128 - 159	37, 62, 71	22	30	73	8	(-33.8 - -32.4), (4.6 - 5.7), (45.6 - 47.3)
160 - 191	25, 30, 32, 34, 40, 41	29	29	70	8	(-32.4 - -31.1), (5.7 - 6.8), (47.3 - 49.0)
192 - 223	12, 63, 65	29	21	58	7	(-31.1 - -29.8), (6.8 - 8.0), (49.0 - 50.7)
224 - 255	39, 51	29	21	58	7	(-29.8 - -28.5), (8.0 - 9.1), (50.7 - 52.6)
256 - 287	22, 59	29	11	50	3	(-28.5 - -27.2), (9.1 - 10.3), (52.6 - 54.5)
288 - 319	60	29	11	49	3	(-27.2 - -25.9), (10.3 - 11.5), (54.5 - 56.5)
320 - 351	35	29	10	49	2	(-25.9 - -24.6), (11.5 - 12.6), (56.5 - 58.6)
352 - 383		29	8	31	2	(-24.6 - -23.4), (12.6 - 13.8), (58.6 - 60.0)
384 - 415		29	8	31	2	(-23.4 - -22.2), (13.8 - 15.0)
416 - 447		29	9	30	4	(-22.2 - -20.9), (15.0 - 16.1)
448 - 479	20	46	33	46	18	(-20.9 - -19.7), (16.1 - 17.3)
480 - 511		46	33	46	18	(-19.7 - -18.5), (17.3 - 18.5)
512 - 543	26	47	36	52	20	(-18.5 - -17.3), (18.5 - 19.7)
544 - 575	8	50	36	52	20	(-17.3 - -16.1), (19.7 - 20.9)
576 - 607	3, 6, 11, 13, 21, 31	50	36	52	20	(-16.1 - -15.0), (20.9 - 22.2)
608 - 639		50	36	52	20	(-15.0 - -13.8), (22.2 - 23.4)
640 - 671	14, 28, 54, 56	50	36	52	20	(-13.8 - -12.6), (23.4 - 24.6)
672 - 703	10, 47, 58	50	36	51	20	(-12.6 - -11.5), (24.6 - 25.9)
704 - 735	5, 16	50	36	51	20	(-11.5 - -10.3), (25.9 - 27.2)
736 - 767		4	6	17	3	(-10.3 - -9.1), (27.2 - 28.5)
768 - 799	2, 9	4	6	17	3	(-9.1 - -8.0), (28.5 - 29.8)
800 - 831	24	3	4	17	2	(-8.0 - -6.8), (29.8 - 31.1)
832 - 863	66	3	4	17	2	(-6.8 - -5.7), (31.1 - 32.4)
864 - 895		6	5	22	3	(-5.7 - -4.6), (32.4 - 33.8)
896 - 927		7	8	29	3	(-4.6 - -3.4), (33.8 - 35.2)
928 - 959	7, 46, 49, 55	7	18	46	3	(-3.4 - -2.3), (35.2 - 36.6)
960 - 991	17, 23	7	18	46	3	(-2.3 - -1.1), (36.6 - 38.0)
992 - 1023	33, 50	11	29	72	7	(-1.1 - 0.0), (38.0 - 39.4)

Figure 10 - Spatial utilization table of non-evenly distributed sector

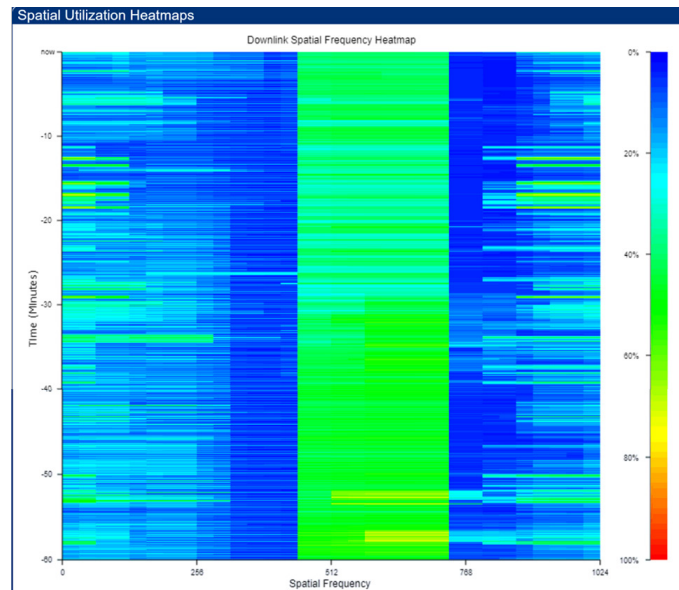


Figure 11 - Heatmap of non-evenly distributed sector

Looking at the Spatial Utilization table, we see clusters of SMs in some SF bins, while other bins do not have any SMs. This uneven distribution is not ideal, but in itself it is not necessarily problematic.

What we need to look at is the amount of traffic the SMs generate. If there are many SMs in one spatial direction, but they all generate little traffic, that direction may not be congested. On the other hand, if there are only a few SMs in one spatial direction but they generate a lot of traffic, that direction may be congested. Because of this, there is no rule on how many SMs are allowed in each bin. It all depends on the amount of traffic they generate.

Looking at the physical location of the SMs is only one aspect of the problem. However, if there is reason to believe that all SMs will more or less generate the same amount of traffic, then a more even distribution of the SMs is beneficial to the overall sector capacity.

The heatmap gives a better idea of how the traffic is distributed. For example, bin 96-127 has one of the largest numbers of SMs of any bins. However, the heatmap does not show excessive traffic in this bin.

On the other hand, bin 576-607 has the same number of SMs as bin 96-127, but these SMs generate a lot more traffic as shown in the heatmap.

In this specific example, the heatmap shows there is still some capacity in this direction, so action is not necessary by the operator. But what if this bin starts showing continuous usage with yellow and red colors? If the deployment allows SMs to possibly connect to more than one AP, a load rebalance exercise could be beneficial.

First, identify the bins that generate the most amount of traffic and appear congested over long periods of time.

Next, identify if some SMs can be attached to a different AP.

If that is possible, check first if the other AP has spare capacity. If it does, change the color code or the frequency of the SMs to force them to attach to the other AP.

Having fewer SMs in the congested bins decreases the amount of traffic in that spatial direction and allows the scheduler to schedule more traffic from other SMs in the same SF.

## More Examples

Figures 12, 13 and 14 show more examples of heatmaps taken from live sectors.

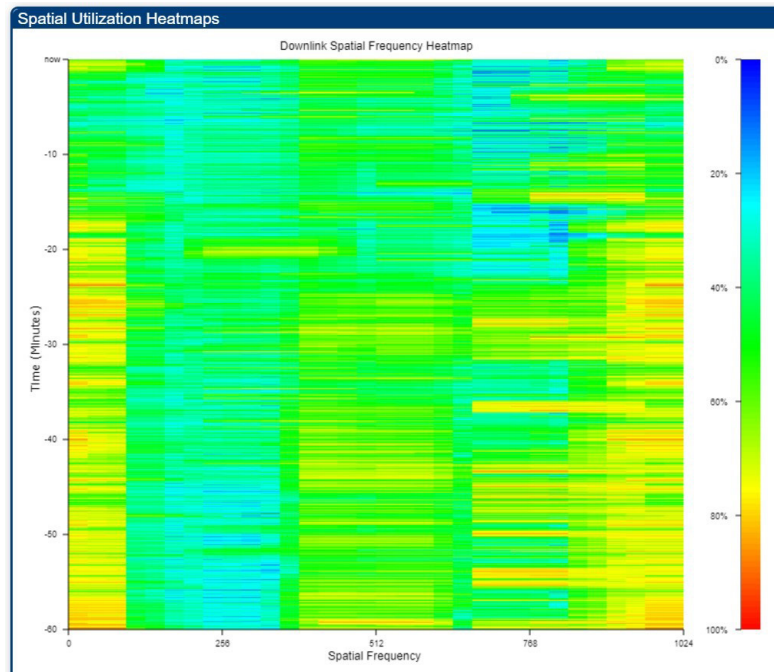


Figure 12 - Additional example 1

Figure 12 shows a sector with mostly uniform traffic, as most of the color is green. Some SF bins (the very first ones and the very last ones) are slightly more active, while some SF bins (around SF 200-300) are slightly less active. Overall, this sector shows a good traffic distribution and additional capacity in every SF direction.

Note that each update in the heatmap is itself an average over 6 seconds. Due to the bursty nature of internet traffic, there still can be brief periods of time where some SF bins are at full capacity, and it is not reflected in this graph. That can be observed from the MAX column of the spatial utilization table.

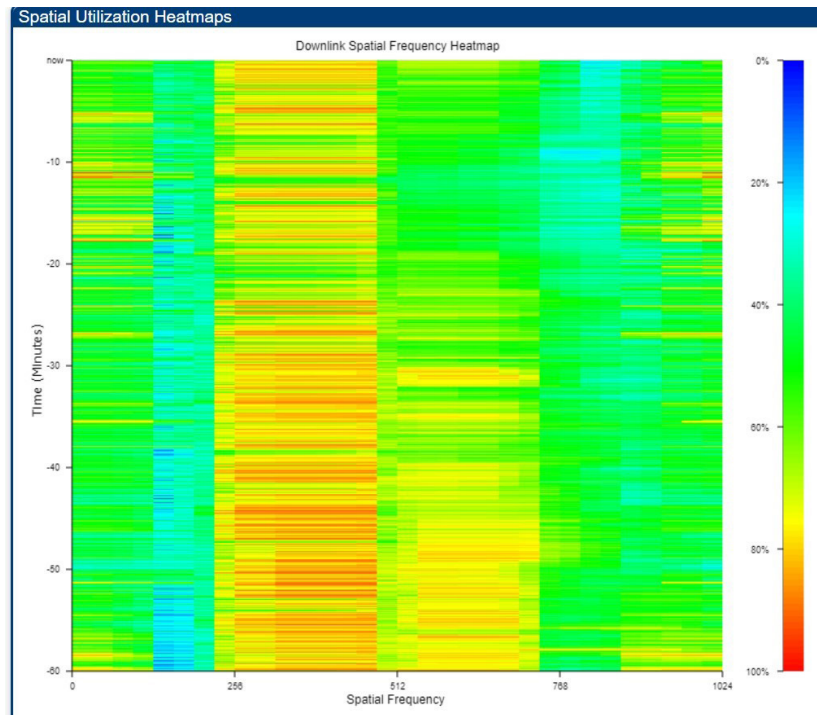


Figure 13 - Additional example 2

Figure 13 shows a sector that is reaching capacity in one SF direction, in bins with SF around 250-500. Other directions are significantly less busy.

Does this sector have spare capacity? If additional traffic is generated in the direction that is already busy, then the throughput cannot be increased, and the sector appears as having reached capacity. If the additional traffic is generated in other directions, then it can be grouped with the existing traffic and more data can be sent by the AP.

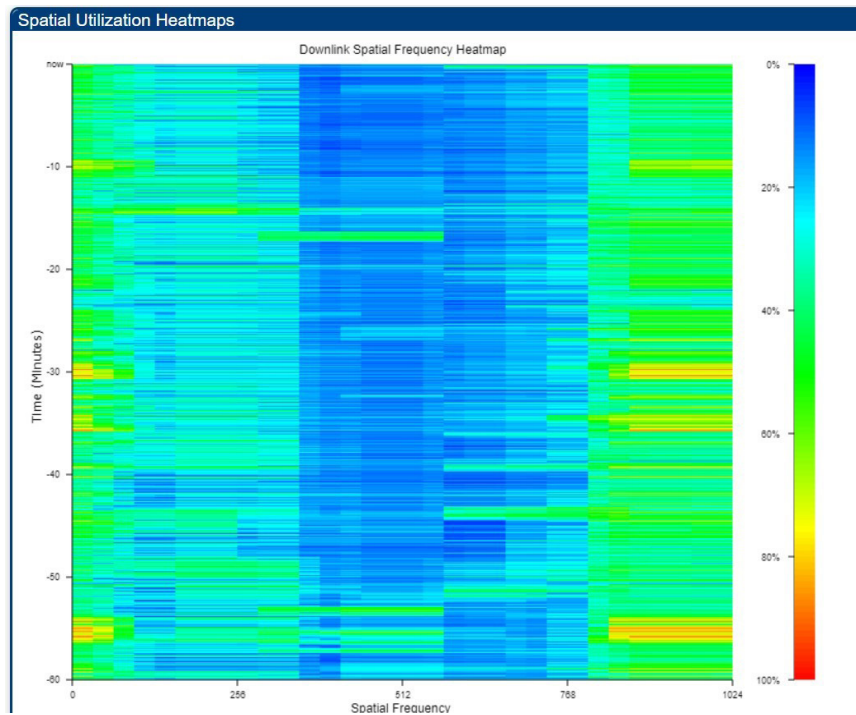


Figure 13 - Additional example 2

Figure 14 shows a sector which is not very active. Especially in the bins in the middle of the heatmap the utilization is constantly low over time. If the number of SMs in these directions is also low, an operator could target their marketing material to potential users in the corresponding areas, as there is clearly more capacity in this sector.

---

## Summary

In conclusion, support of the MU-MIMO feature makes the analysis of a PMP 450m sector's spare capacity more complex. A single parameter is not sufficient to give the full picture, as it happens with APs that do not support MU-MIMO.

In a 450m AP, follow these steps to have a full view of the sector's capacity:

- 1. Check the Total Utilization percentage:** If the value is low, the sector has spare capacity. If the value is high, further investigation is needed.
- 2. Check the average group size:** If the value is low, the sector has spare capacity, but the traffic pattern will determine if the extra capacity can be used by the scheduler or not. Information from the heatmap is needed. If the value is high, it is likely the sector has reached its capacity. The multiplexing gain gives more granularity than the average group size.
- 3. Check the multiplexing gain:** If the multiplexing gain is much lower than the average group size, then some SMs in the group generate much more traffic than others. If there is spare capacity in the sector, it may be difficult for the scheduler to make use of it.
- 4. Check the heatmap:** The heatmap gives the best view of the sector capacity, as it gives information over SF and over time. If traffic is high and evenly distributed, then it is likely that the sector is reaching capacity. If the traffic is high but concentrated only in some directions, then additional traffic in those directions is not possible, but traffic in other directions may be possible. If the traffic is low, but concentrated only in some directions, then additional traffic in those directions is not possible, but traffic in other direction is definitely possible. If the traffic is low and evenly distributed, then additional traffic in all directions is possible.