Aggregate Wi-Fi Telemtery Use Cases

A Technical Paper prepared for SCTE•ISBE by

Colleen Szymanik
Sr. Principal Engineer
Comcast
1800 Arch Street Philadelphia, PA 19103
215-970-1953
Colleen_Szymanik@comcast.com

Table of Contents

Title	Page Number

1. Introduction 3
2. Data Sources 3
3. Metrics 3
   - 3.1. Interface Statistics 4
   - 3.2. Radio Statistics 4
   - 3.3. Client Station Statistics 6
4. Collection Interval 7
5. Active vs. Passive Measurements 8
6. Conclusion 8

Abbreviations 9

Bibliography & References 9

List of Figures

Title	Page Number

Figure 1 - 2.4GHz and 5GHz available channels 5
Figure 2 - Interference 6
Figure 3 - Client capabilities 7
Figure 4 - Aggregate Trending 8

1. Introduction

Assessing WiFi performance can be challenging, especially at scale. Most operators and service providers need provisions to understand the quality of delivered service to ensure customer satisfaction. In order to assess the service, we need to understand the requirements for consistent and reliable service offering over WiFi. Since WiFi can be used over both 2.4 GHz and 5GHz frequencies, they should be measured and assessed separately given there are different challenges per radio frequency. In this paper, each router that services a customer home will assume 2 connected radio interfaces. Once the 6GHz band is widespread, this will be an added (3^{rd}) interface to quantify. Each connected client device, such as laptop or smart phone, will also have different bandwidth requirements to obtain the intended service, regardless of the connected interface. In this paper, we will talk through data sources and collected metrics. From those metrics, is it possible to make assessments at a aggregate level to infer reliable and consistent quality of service from poor or unusable service? Conversation pairings will refer to the transmit and receive between an access point (AP) and client device, such as a laptop. The idea here is to refer to this as a simple conversation. In order to have effective communication between both the AP and the station, the rules put in place by the protocols and the WiFi specification need compliance. Assuming compliance is met, measuring quality is the next step. WiFi quality depends on the specific conversation pairing requirements. Each pairing does not have the same time and bandwidth requirements, such as an IoT device compared to a 4K streaming device. They should not be treated the same when appraising the connection quality.

2. Data Sources

The Access Point (AP) to client station connection is typically a one to many relationship. As an operator, measuring data from the AP is typical since it is the central point of service in the home. It is also usually managed at the service provider level. The other main element in the equation is the client station device. While client station is not under the service provider control, it is important to recognize that the important role this device component has in this exercise. These measurements are solely to provide consistent and reliable connections to these client devices. While client devices have varying capabilities, it is fundamental to consider what is needed for each of the individual client types. Some of the most successful ways to rate the level of service is to be as individual as possible for all the requirements. Clients on a network have specific needs. As an service provider or operator, can we measure those needs at scale?

3. Metrics

Most measurements for WiFi quality of experience combine the same basic elements. The primary components consist of interface statistics, radio metrics and client connection statistics. The challenge with these elements are the fact they are constantly changing. WiFi is a very resilient protocol, as with most networking elements. These elements are constantly changing to meet the current conditions for service delivery. An example of the changing conditions is the rate adaptation mechanisms in video delivery to drop from HD service to SD service if HD service levels are not met. Adaptability is a key method for resiliency, but it poses challenges to understand when the quality of service compromises the customer experience. The components can be measured and assessed to provide feedback on that customer experience.

3.1. Interface Statistics

Interface statistics are useful in understanding what was successfully transmitted and/or received between the AP and the client station. How often were those transmission successful, filled with errors, or needed to be repeated? When we are quantifying applications, such as video, some retransmissions can be tolerated due to provisions such as buffering. While video is mentioned, each application for each client device type can be tailored to the the corresponding applications. The result is tailored to the specific service level requirement for a good user experience. Having a good ratio for interface statistics can help properly assess whether or not the current level of interface statistics are viable for a consistent user experience. Interface statistics consist of sent and received, as well as errors, retries and retransmissions. One of the key differences could be fine tuning ratios that can prove service level offerings. If it is known what is transmitted and received. Looking at the proportion of frames sent and received with respect to retransmissions and errors can yield insight into the quality of the connection. As such, measuring round trip time to assess latency is another potential method to assess the service quality between the AP and the client station. In gaming, most gamers will use latency as a measure of whether or not they are able to achieve consistent service.

3.2. Radio Statistics

One can argue that gaining access to the medium is the biggest challenge with WiFi today. In fact, from 802.11's beginnings, access control still presents a challenge. WiFi has always been a "listen before you talk" service. MU-MIMO was the long awaited feature in WiFi5. Most of the excitement for WiFi6 is OFDMA. Efficiency in admission control to the medium is sought after in WiFi implementations. If the AP can successfully group devices to transmit and receive at the same time within the group, this can free up more airtime to be used for even more transmissions. The number of concurrently connected devices in a home are growing as more and more devices are WiFi capable. Practices such as band steerring, pushing a client device on the 5GHz band is one way to make better use of a less congested frequency. The unlicensed 2.4GHz band is already over used and over crowded, especially in densely populated areas. Measuring the radio quality is key to understanding how much time is effectively being used for client connections. In 2.4GHz, there are only 3 non overlapping channels to use.

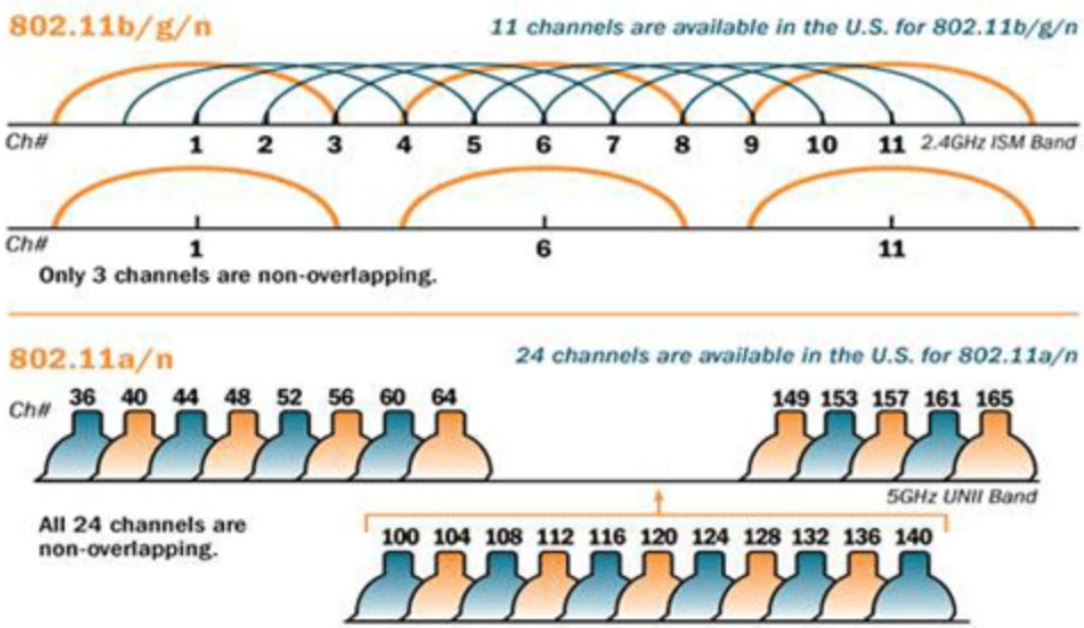

Figure 1 - 2.4GHz and 5GHz available channels

Given that WiFi operates in unlicensed spectrum, there is a lot of congestion in this space. As mentioned, it is typically worse in the 2.4GHz band, but congestion can still be an issue in 5GHz, especially in dense urban environments. This usually results in a lot of extra interference in this band that makes the signal to noise ratio for those connected clients lower than that in the 5GHz band. The noise from WiFi will results in interference.

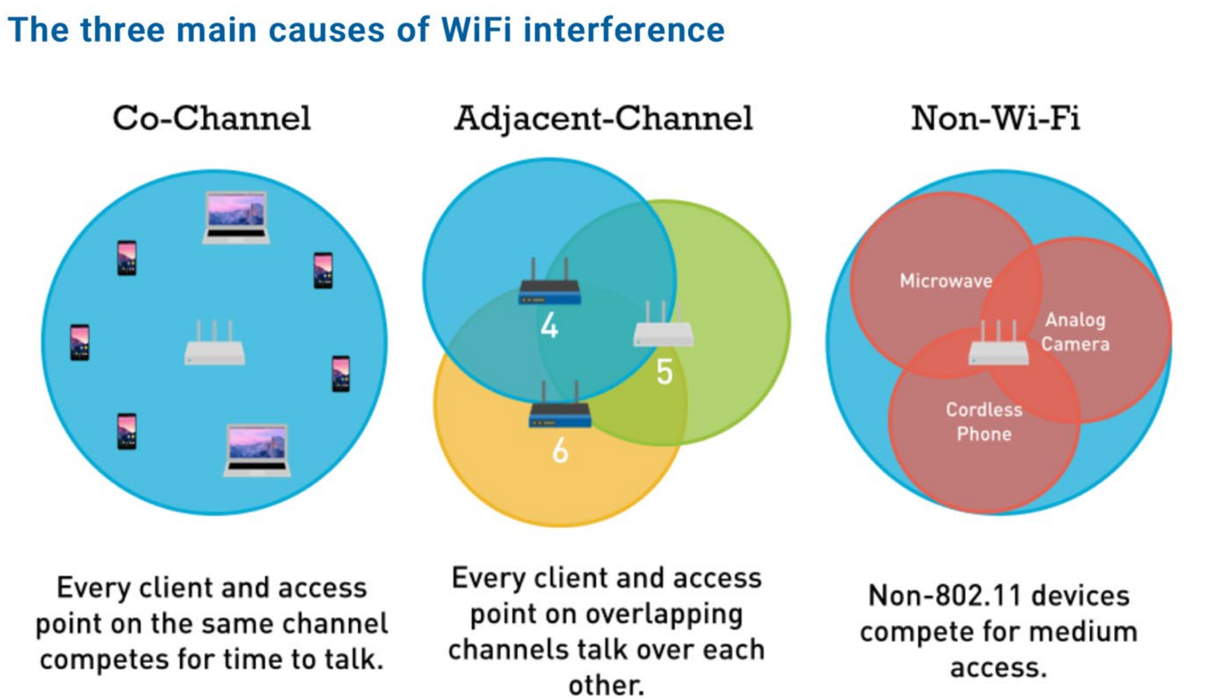

Figure 2 - Interference

In both bands, measuring the channel utilization where we can understand the amount of time used for servicing connected client stations is important for overall radio measurements. While there are some ways to do this in a standardized method, there are slight nuances that make normalization between APs necessary. Depending on the device, operators may see measurements such as transmit opportunity (TXOP) or free airtime (FAT) that begins to characterize the radio interface. These differences in the measurement will require normalization in a multi vendor environment. If we treat the radio measurement in the same fashion as the conversation pairing, it is a natural progression to normalize the measurement for aggregate trending.

3.3. Client Station Statistics

Client connection statistics that reflect signal levels of the current connection are always present in most quality of experience assessments. RSSI, PHY rates, MCS values, SNR values all provide feedback to estimate the AP to client station connection. Once again, there are nuances for these measurements. It is important to know what exactly is being reported. In WiFi, the management frames are transmitted and received at the lowest mandatory rates. If the reporting includes these management frames, this can skew the estimate of the connection quality. It is also worth noting that each client type is limited from the physical capabilities.

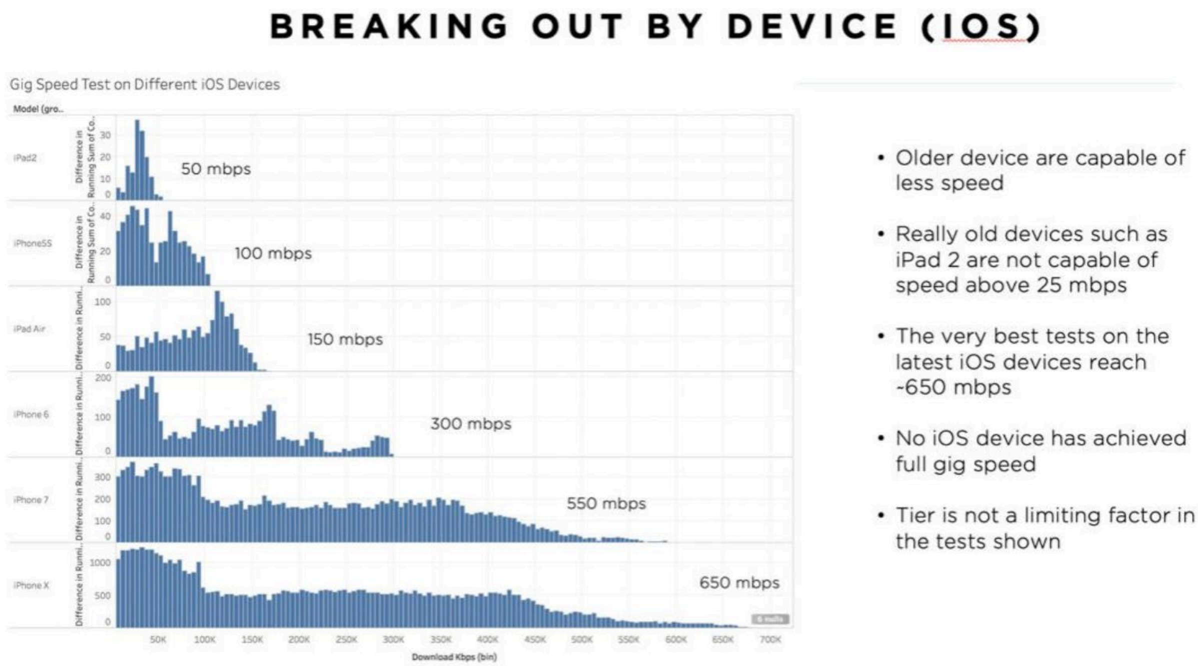

Figure 3 - Client capabilities

Not all client stations are created equal and should not be measured as such. As mentioned in Data Sources, operators can use the APs to take all these measurements so they know what is being measured and how often. Typically, APs have more transmit and receive antennas than the connected client devices. That imbalance of radio power needs to be recognized when evaluating these signal levels. Thresholds for these client types need testing feedback to ensure the levels are reflective of user experience.

4. Collection Interval

We mentioned earlier that WiFi is adaptive and constantly changing for current conditions. Collecting data at scale from moving targets should be considered when choosing collection intervals. Even if data storage and infrastructure is scaled well to collect these metrics often, it might not be necessary to collect data every minute or even every second to get reasonable measures of quality. In this conference, the idea of limitless possibilities is challenged. If that theme is extrapolated to the point where data storage and processing power is not an issue, do we want to collect data frequently? Does that make sense? In all computing systems, once mechanisms are built, efficiency is the next logical step, even if scale isn't an issue. Collecting data at the edge and aggregating into a cloud element is only meaningful if the right building blocks are in place. Beginning with the conversation pairing is fundamental to then build groups to count and ultimately trend and aggregate. Using that order of operations as a rule set can yield consistent aggregation points. If one can target the right thresholds for the individual conversation requirements, how often is necessary to check that connection? A good start would be hourly, if the infrastructure can support the individual pairing data model. A more granular set of metrics has been found to roll up 15 minute intervals. There are some methodologies that will have localized collection agents at the WiFi interface to aggregate the data to be collected less often by centralized infrastructure.

Strategies for this type of data is important make sense of this at scale. If one keeps averaging or aggregating data, it will just make the measurements less meaningful and will not capture the experience properly. If we can divide the data into different compartments for each measurement, we can count those data occurrences while quantifying scale at that individual level. As a result, collection intervals can be done less often and aggregation can be scaled while still providing a meaningful experience gauge. The idea here is to expand vertically rather than horizontally. Limitless possibilities doesn't always mean expanding in the same direction.

5. Active vs. Passive Measurements

All of the data metrics and measurements discussed at this point are assuming that these measurements are passive. This means that we are going to watch active client station connections to a given AP and measure the primary components and report those measurements accordingly. Given that applications and client stations change hardware and software often, there is a need to pulse ever changing conditions. Having active measurements can aid in trending those changing conditions. If the endpoints are static to provide a baseline for each AP type and client station device type, that baseline can be trended for deviations. This will be instrumental in keeping the thresholds for the individual conversations meaningful.

6. Conclusion

Aggregating quality of experience data at scale can be done multiple ways. The interesting effect WiFi has on this quality of experience is that one to many or 2 way conversation is where the first operation is computed. Measurement data from specific thresholds for endpoint pairings will be the fundamental step before rolling the data into aggregate trends. Anything that is connected to the AP that has a cataloged resource is considered an endpoint pairing, such as AP to gaming device or AP to IoT device. The basic element to measure all of the following remain consistent in WiFi quality of service: interface statistics, radio metrics and client connection statistics. Counting the amount of occurrences where an acceptable service delivery was not met for those threshold pairings is paramount.

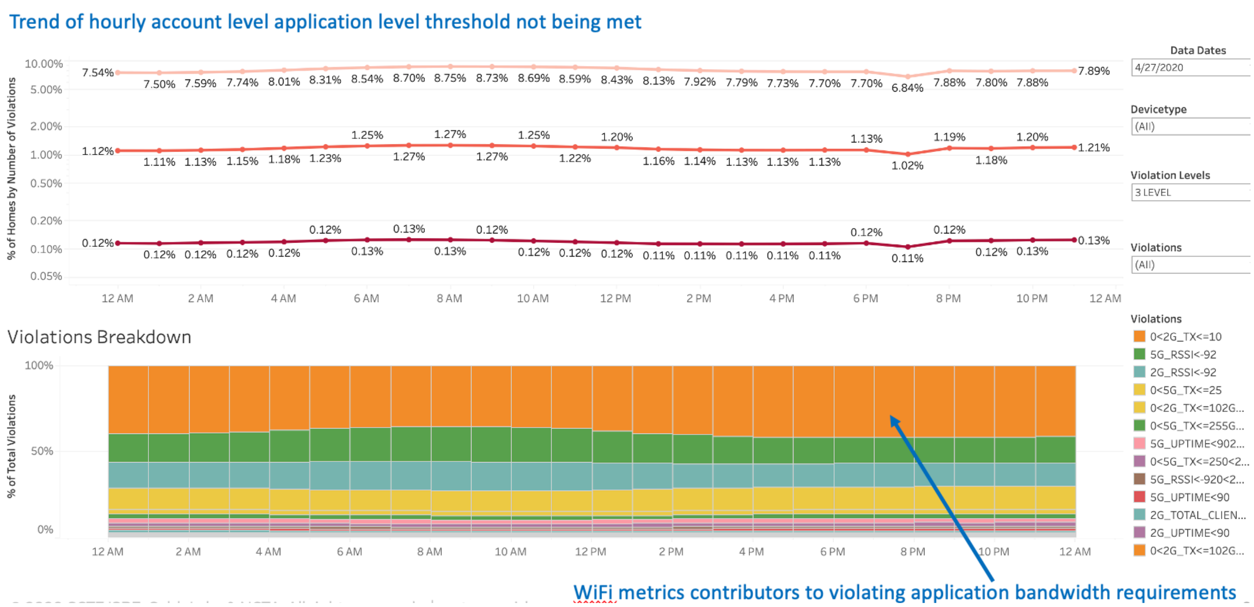

Figure 4 - Aggregate Trending

Once those counts are listed, it is important to give those counts context. Using those as a percentage in the overall service delivery can setup a baseline. In the above referenced figure, thresholds are applied across conversation pairings and counted as a percentage to assess impact on the wireless customer experience. The percentage of time polling systems captured the signal and bandwidth requirement below the configured threshold is trended. The bottom chart shows the distribution for the contributing reasons why the service delivery was not satisfied. If those endpoint pairings, such as service delivery for video streaming devices are contextualized, the ability to filter on those specifics can understand the trends for capacity management can be shown at an aggregate level. If that order of operations is not met, aggregate data will remain less meaningful.

Abbreviations

AP	access point
bps	bits per second
Hz	hertz
ISBE	International Society of Broadband Experts
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

Figure 1: Plashal, Andres. "Differentiating 5GHz and 2.4GHz Frequency Bands" *July 2020*
https://andres.plashal.com/2018/blogging/differentiating-5gvs2g-frequency-bands/
Figure 2: Metageek Training Documentation. "The Thress Main Causes of WiFi Interference" *July 2020*
https://www.metageek.com/training/resources/why-channels-1-6-11.html