WHITE PAPER

The Importance of Performance Testing

for Residential Wi-Fi Deployment

Table of contents

Figures

1. Introduction

Somewhat surprisingly, the concept of standardized performance testing is a relatively new concept in the world of Wi-Fi®. (We will use the term "Wi-Fi" throughout this document with reference to the certification program and certified products from the Wi-Fi Alliance. Although we will not always include the trademark symbol, this is always implied.) When we talk about performance testing, we are distinguishing it from other product/device testing, such as regulatory testing and conformance testing. So, it will be useful to begin by defining those other types of testing that are not the focus of this paper.

1.1. Regulatory testing

Everywhere in the world, before wireless products can be used in an open-air environment, they must conform to certain regulatory requirements. In the US, for example, Wi-Fi products must conform to the Federal Communications Commission's (FCC's) Part 15 rules for Radio Frequency Devices. These tests are aimed at making sure that devices obey the operating rules for the spectrum they will use, and that they don't interfere with other users in that spectrum or in other spectrum bands. The rules say nothing about what technology should, or should not, operate in the spectrum. This is why, for example, a Wi-Fi device and a garage door opener can operate in the same band. The devices don't interoperate (that is, they don't communicate with each other) but they can share the spectrum. This is the focus of almost all regulatory testing – basic spectrum compatibility.

1.2. Conformance testing

From a Wi-Fi perspective, this is the type of testing that is most well-known/understood and has contributed to the incredible growth of Wi-Fi over the past 20 years.

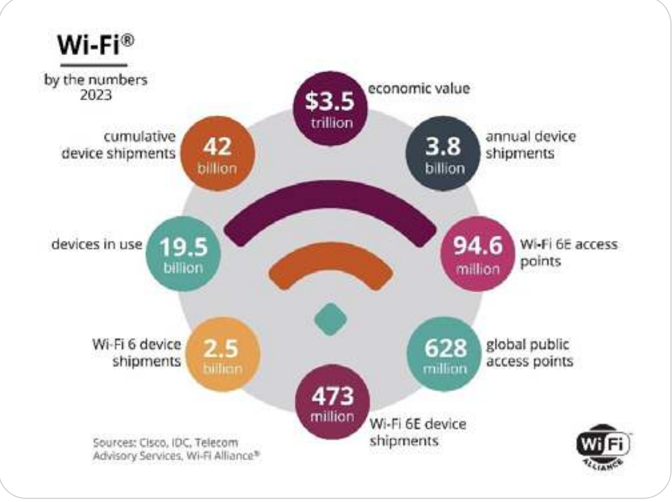

Figure 1: Wi-Fi by the numbers[1]

The success and growth of Wi-Fi, as demonstrated in Figure 1, is due in large part to the conformance testing regime created and managed by the Wi-Fi Alliance. Although it is the IEEE's 802.11 group that creates the specifications that describe the underlying Wi-Fi technology, it is the Wi-Fi Alliance that has imposed a conformance testing regime that awards certification to products declaring that they support some subset of the IEEE specification(s). It is this conformance regime that has allowed users of the technology, both consumers and service providers, to have confidence that certain devices will interoperate with other devices. When those devices bear the Wi-Fi Certified label, interoperability is all but guaranteed. People trust that products from one company will work with products from another and that it will be simple to create a multi-vendor deployment. It is partly for this reason that, for example, a survey in the UK declared that "Wi-Fi is the best invention of the past 25 years[2]".

1.3. Performance testing

Given the success that Wi-Fi has seen over the decades, why is even more (and different) testing important? That is, why on top of regulatory and conformance testing is there now a focus on standardized performance testing? And how has it been that, over all this time, performance testing has not been an industry focus? We will answer these questions and get into the details of performance testing in the next sections.

> Wi-Fi is the best invention of the past 25 years

2. Performance testing in Wi-Fi

The difference between conformance testing and performance testing is the difference between the question "will this device interoperate with that device," and "how well will this device perform when interoperating with that device?" It is, in many respects, the success of the Wi-Fi Alliance's conformance testing regime that has elevated the need for performance testing. If the conformance testing were not a success, there would be very few interoperable products, so then very few deployment options to consider. The most extreme example of this is for proprietary technologies with no conformance standards, in which all radios in the system must be manufactured by the same vendor.

Given that Wi-Fi products are so interoperable, when a service provider, for example, decides to deploy a Wi-Fi enabled router, they have a large number of options to choose between, all of which are, at a minimum, Wi-Fi certified. The question then becomes, do they all give the same performance? This is a question that conformance testing is generally not designed to answer, and this is where performance testing comes into the picture.

It would be misleading to say that throughout the history of the development and deployment of Wi-Fi products, no one has done any performance testing. This is clearly not the case. What is the case, however, is that no standards organizations have addressed the topic of performance testing until very recently. In the early days of Wi-Fi, much performance testing was done using walk-test methods, in which Wi-Fi devices were deployed into actual buildings, and measurements were taken with test tools to look at how the devices behaved under various conditions. (See Figure 2.)

This methodology has a number of drawbacks, in that:

a. The results are not very repeatable because in open-air environments, the RF environment can be constantly changing, which will affect the results.

b. The results are not easy to reproduce, among different testing labs, for example, because one real environment will likely be very different from another.

c. The testing can take a very long time, since networks need to be configured, and testers have to physically walk around the environment to collect data before it is analyzed.

d. Various test conditions can be very difficult to set up in a real environment, for example, many devices sending various kinds of traffic under different interference conditions, etc.

e. The testing is not amenable to standardization because it's very difficult to standardize a physical environment (like a home) in which devices were to be deployed and walk-test data was to be collected.

These are among the reasons that performance testing has not been standardized until recently, but there is one more reason that has limited the development of standardized performance testing.

Figure 2: Illustrative data taken during a walk test in a real building

2.1. The meaning of the word "performance"

Standards bodies are made up of participants from many areas, including academia, government, service providers, and equipment vendors. While all of these stakeholders need for equipment to conform to regulatory requirements (regulatory testing), and in many cases they all see advantages to products meeting conformance requirements (conformance testing), it has not always been quite as clear that there is a common understand of performance testing.

Equipment vendors, in particular, have historically been concerned that, once a test defines performance characteristics, certain products will be perceived as "better" than other products, even though there might be a complicated mapping of performance to price point, so that what is "better" for one user might not be "better" for another with different requirements.

2.2. The turning point for performance testing

Despite the difficulties in creating standardized performance tests for Wi-Fi, the past several years have seen a rapid uptick in the number of such standards, and the momentum for them is continuing. How has this happened?

a. The advent of compact Wi-Fi testing tools has removed many of the difficulties listed above. Wi-Fi testing can now be done more quickly, more repeatably, under more complex scenarios, and in a way that can be more easily defined by a standards organization.

b. An understanding has emerged across the industry that performance testing does not imply "better," rather, it looks at whether a product is "fit for purpose." There is a complicated mapping of a product's feature support, performance, price, and intended area of operation. Because of this, there is no real meaning to saying that one product is "better" than another, but it is possible to say that a product may be more fit for use in, for example, a small home environment, while a different product is a better fit for a larger, crowded, interference-prone office environment.

3. Survey of the performance testing landscape

The demand for standardized performance testing for Wi-Fi was ultimately driven by service providers (although other stakeholders have been active participants and supporters). The success of conformance testing led to a dramatic rise in the use of Wi-Fi, and as a result, most users experience their broadband connection via Wi-Fi as the final link. According to one study, 92% of US internet households use Wi-Fi at home[3]. This being the case, it became vitally important that broadband service providers manage this part of the connection since the entire broadband experience depends on it. With ISPs making Wi-Fi a part of their service offering, they need to be able to compare different Wi-Fi products to see which have the best fit with their offering, so they started to advocate for the position that standardized performance testing was increasingly necessary.

The first standards organization to act on this was the Broadband Forum (BBF), which in 2019 issued the first version of its TR-398 test plan, "Wi-Fi In-Premises Performance Testing." The focus of this testing was residential access points. (Client devices are not a focus of this test plan.) This was the first Wi-Fi test plan to focus on performance by including not only test cases and methodology but also pass/fail criteria to go along with each test. Obviously, this could only be done under the assumption/requirement that largely similar, repeatable testbeds could be created in which these tests could be run.

A few years later and at roughly the same time, both the Wi-Fi Alliance and ETSI also entered the Wi-Fi performance testing space with the Wi-Fi Customer Experience group's release of the Wi-Fi® Device Metrics test plan[4], and ETSI Broadband Radio Access Network's (BRAN's) release of the Multiple Access Points Performance Testing plan[5], both in mid-2022. As is already somewhat obvious from the names of all of these test plans, they do not all cover the same thing in terms of "performance." (This relates to our previous comment about products being "fit for purpose.") For example, while the BBF TR-398 test plan is explicitly called out as being for in-premises performance, the ETSI test plan is called out as being explicitly about multiple access point performance (so performance of APs when configured in a mesh, or extender, scenario). There are many areas of "performance" that can be examined, and in the next sections, we will summarize what each of these groups has covered so far, and then we will discuss what types of performance testing have yet to be addressed.

3.1. The assumption of compact, consistent testbeds

One thing that is immediately obvious when reviewing these three specifications is the impact that compact test systems have had on the ability of groups to even consider widespread, standardized performance testing. The BBF, for example, describes how the testing can be done in either single, or multiple-chamber testbeds (Figures 3 and 4). In some more advanced test cases (discussed below) only the multiple chamber testbed option is available.

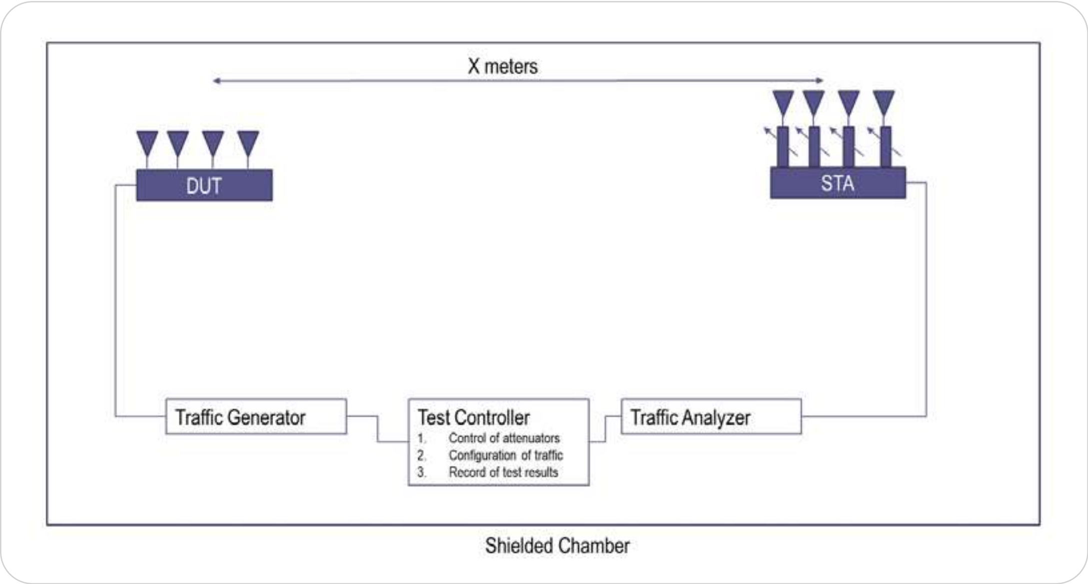

Figure 3: BBF Single-chamber diagram

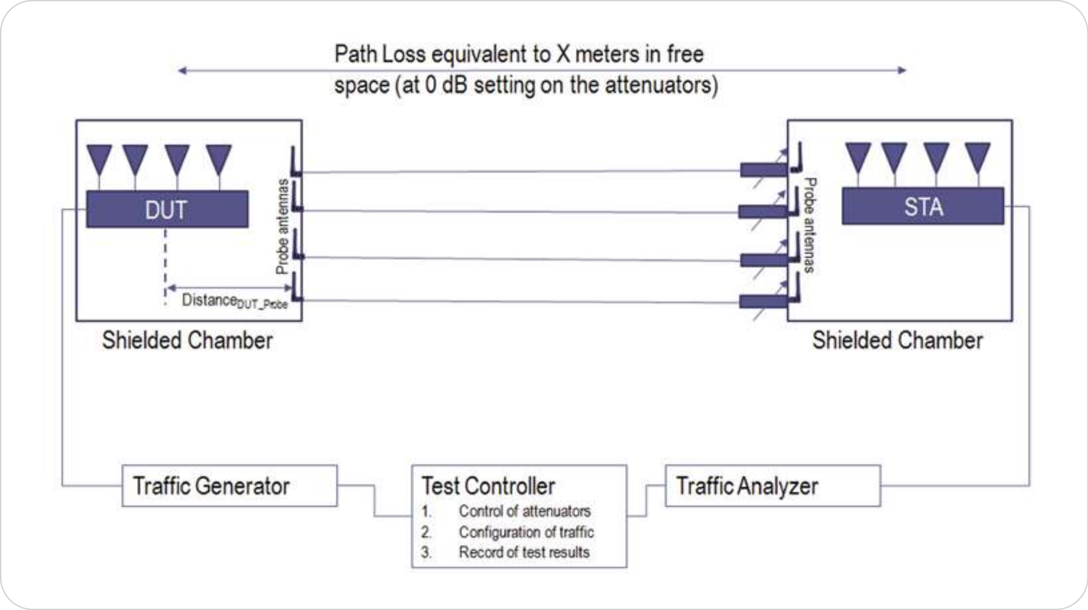

Figure 4: BBF Multiple chamber testbed implementation

The Wi-Fi Alliance shows, as an example, a multiple chamber implementation that can be used for their testing (Figure 5).

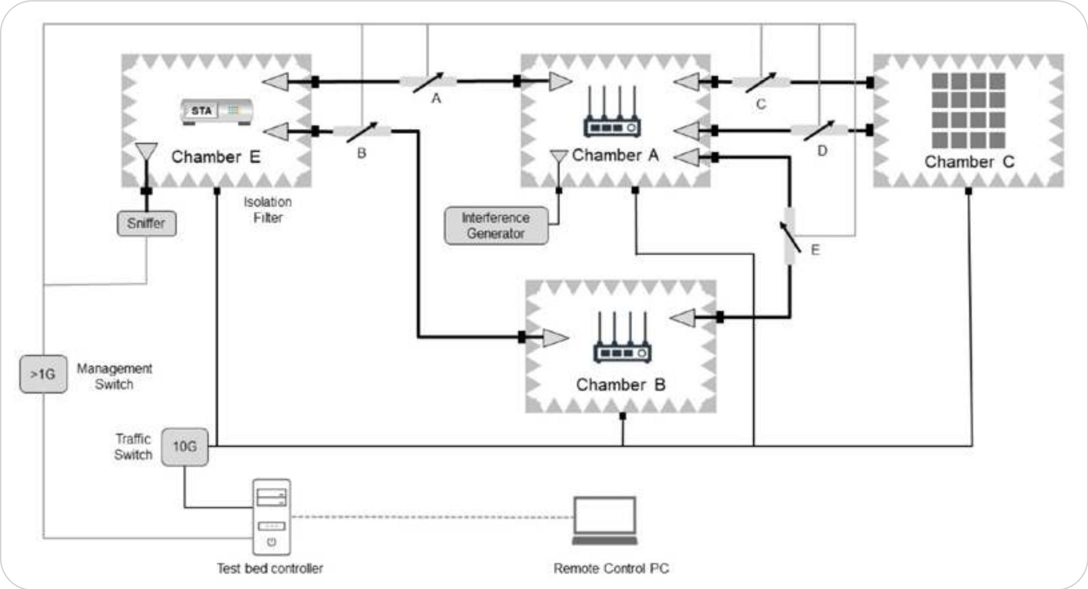

Figure 5: Wi-Fi Alliance example three-chamber testbed

And, similarly, the ETSI test plan describes how a multiple chamber testbed can be used to configure the complex, multiple access point scenarios described in that test plan.

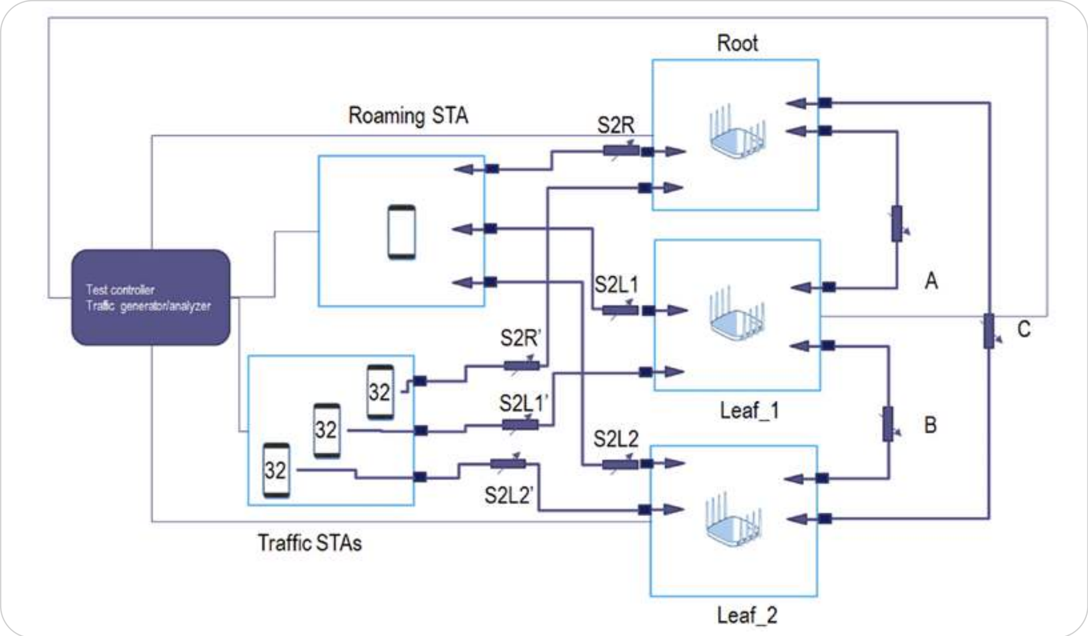

Figure 6: Multiple chamber testbed for multiple access point testing

There is commonality between all of these standards, in that they all require the testing be done in a controlled RF environment (not in open air). In addition, they all require (at least for some of the test cases) that multiple, independent, interconnected RF chambers be used to create more complicated testing topologies than would be possible in a simple shield room scenario.

3.2. BBF TR-398

As mentioned above, the first group to enter the Wi-Fi performance test space was the Broadband Forum with TR-398, which is why that group is currently nearing the release of the third revision (or "issue") of that specification. As these tests were primarily created to respond to the needs of service providers, they are focused on the performance of Wi-Fi Access Points, or APs, (not clients, or "STAs" in Wi-Fi terminology). In addition, as we'll see, the focus has been primarily (but no longer solely) on single access point deployments, which implies a fairly simple home deployment with a single AP. This is the use case for many, even most, home Wi-Fi deployments.

3.2.1. TR-398 Issue 1

As the first mover in the Wi-Fi performance testing space, the BBF group began with a fairly wide-open field of items to test. In Issue 1, they broke their plan down into 5 sections:

1. RF capability
2. Baseline performance
3. Coverage
4. Multiple STA performance
5. Stability/Robustness

We will not, in this paper, get into the details of each test, its configuration, etc. (The specification itself is freely available, and all of the details are spelled out there.) Rather, we will focus on how, even in this first set of tests, the complexity of Wi-Fi performance testing can be seen, and the idea of "fit for purpose" becomes ever more apparent.

In some of the tests above, the concept is quite simple, and it might seem obvious what a good result is and what a bad result is. Under "Baseline performance", for example, is a simple test of maximum throughput. Traffic is run from the Device Under Test (DUT), which is an AP in this test plan, to a test instrument, and the throughput is measured. It would seem obvious that in this test, a higher number is better.

But not all tests are quite so obvious. A fairly simple test (in the Coverage section) is the "Range vs. Rate" test, in which the test emulates moving a STA away from the AP so we can measure not only the maximum throughput (at the lowest attenuation point) but also the throughput as the attenuation between the endpoints is increased (emulating an increasing spatial separation). However, depending on the expected deployment scenario, the performance of an AP at high attenuation might be more important for some use cases than for others. Nevertheless, the test plan still needs to provide a baseline of performance in order to be able to qualify as a performance test plan.

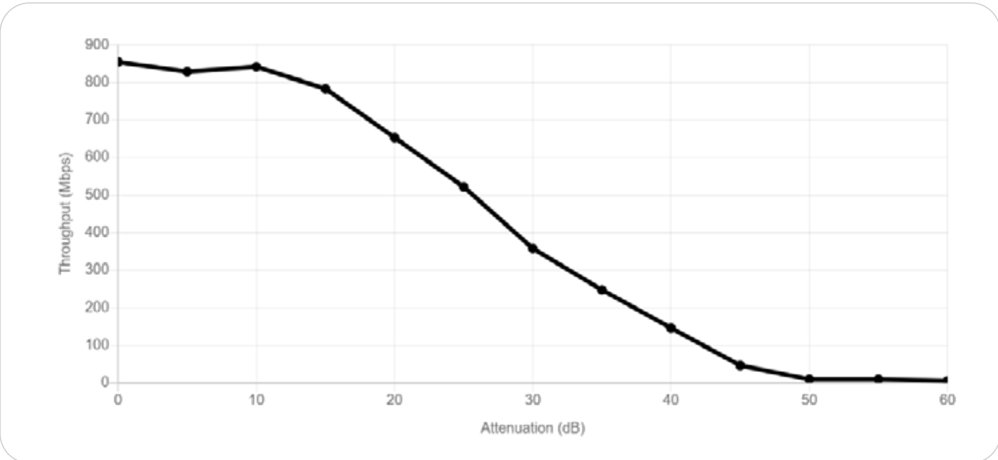

Figure 7: Example of a Range vs. Rate curve

And while even the Range vs. Rate test might seem to still imply an obvious result (more throughput is better), there are even more complicated tests where the ideal result can be less obvious. Take for example, the “Multiple STA performance” test. In this test, we emulate sending traffic first to three clients close to the AP. Then we add in another three clients that are a “medium” distance away. And finally, we add in three more clients that are “far” away (Figure 8).

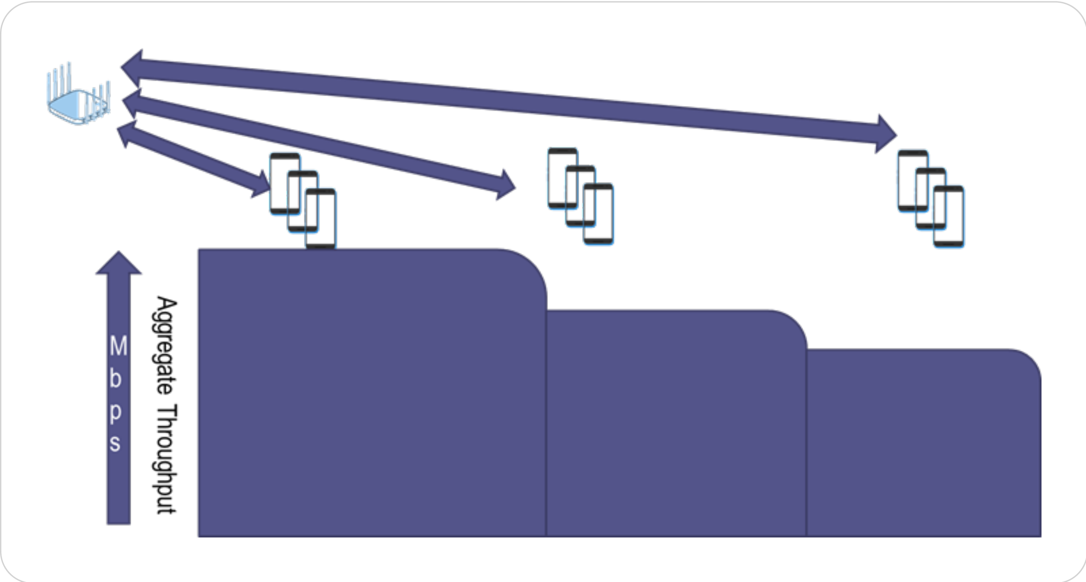

Figure 8: TR-398 Multiple STA performance. A visual representation.

While the test plan does call out the threshold throughput that must be achieved in each scenario, it is clear that in tests like this, we start to enter subjective territory about what is the best way of handling the situation. The highest throughput can always be achieved by handling only the close STAs and ignoring the far STAs, but this is, also obviously, not a good solution. How much should far-away STAs be allowed to impact the overall throughput? What is fair? These are questions that different AP vendors may handle differently, and what is fit for one application may not be fit for another.

3.2.2. TR-398 Issues 2 and 3

As the TR-398 specification has progressed through Issue 2 (released March 2021[6]) and Issue 3 (not released as of the writing of this paper, but likely to be finalized in late 2023 or early 2024), more complex scenarios have been addressed. For example, Issue 2 includes:

- Dual-band and bi-directional throughput tests
- Channel auto-selection tests
- A first attempt at "roaming" and "repeater" tests

In addition, since the primary Wi-Fi technology in the market at the time of Issue 1 was 802.11ac (Wi-Fi 5), there were no tests in Issue 1 related to 802.11ax (Wi-Fi 6.) So, Issue 2 added in a few new tests to begin to address Wi-Fi 6.

With Issue 3, the group has begun to address the new 6 GHz spectrum, as well as areas such as quality of service (QoS) and latency, and the accuracy of measurements reported by the AP that may be used for radio resource measurement (RRM) algorithms, such as received channel power indicator (RCPI), noise, and channel utilization.

3.2.3. Wi-Fi Alliance Device Metrics

The Wi-Fi Alliance Device Metrics test plan is similar to that of TR-398 but with a noticeably different focus. The test cases are familiar to TR-398:

- Rate vs. Range
- AP latency
- Channel switching [for various reasons, like interference or dynamic frequency selection (DFS)]
- Band steering
- Roaming
- AR/VR device performance

As with TR-398, each test case comes with a description of the test configuration and the methodology. The Device Metrics test plan has a heavy focus on how the collected data should be analyzed and presented to the user. A difference from TR-398 is that the Device Metrics test plan does not provide any pass/fail criteria for these tests.

The philosophy behind this test plan is even more along the "fit for purpose" concept, the thought being that if people are presented with the results of these tests in a clear and consistent way, they can decide for themselves (based on their own understanding of their market/customer requirements) whether any given device satisfies their requirements. It is the focus on the statistical analysis and consistent data presentation that distinguishes the Device Metrics test plan.

3.2.4. ETSI Multiple Access Point Performance

Another example of a standards body addressing Wi-Fi performance is the ETSI specification (TS 103 754) called "Multiple Access Points Performance Testing." The interesting thing about this specification is that it is an excellent example of the "fit for purpose" concept we have been discussing. This specification, as its name implies, is only related to multiple-AP scenarios (mesh/repeater installations). For a user/operator that is not interested in a multiple AP scenario, there would be no reason to use this test plan. That seems to be a good description of the direction that Wi-Fi performance testing is moving in general: "Performance" will be defined by whatever the group feels is important to the members of that group (and what they perceive to be important to the broader community of interest). Tests will then be defined to give consistent, repeatable, quantitative insights to tests of interest to address those requirements.

For the Multiple Access Points group, those tests include:

- Roaming time and throughput
- One-hop and two-hop throughputs
- Network configuration and self-healing
- Band steering

4. The future of Wi-Fi performance testing

Based on the above description, the reader would be forgiven for concluding that performance testing for Wi-Fi is a solved problem. There are at least three major standards bodies addressing it (Broadband Forum, Wi-Fi Alliance, ETSI), and each group is tackling the problem from a somewhat different direction.

However, Wi-Fi performance testing is anything but a solved problem. In fact, the efforts described above are a description of the early stages of grappling with how to define Wi-Fi performance. Some of the most interesting features of Wi-Fi have not been dealt with by any of these specifications, and more features are being developed all the time. We describe just some of these advanced features below.

4.1. New Wi-Fi features, and their area(s) of application

Before we get into a specific discussion about some of the newer (and more complicated to test) features of Wi-Fi, it is useful to describe why these new features were developed in the first place.

At its inception (in roughly the year 2000), Wi-Fi was considered mostly a "nice-to-have" feature for broadband users, allowing them to connect to the internet wirelessly "when possible." Coverage was often limited to specific locations (the living room, the study), and speeds were not high. At that time, the main applications were web browsing, email access, etc. Video streaming, IoT, voice-over-Wi-Fi, and the like were virtually unheard of.

The world has changed dramatically since that time. First of all, Wi-Fi is no longer a "nice-to-have" feature of a broadband service plan. As mentioned above, the vast majority of users are interacting with their broadband service over Wi-Fi as the last link in the connection. And the applications are no longer what they were in 2000. According to the graphic shown in Figure 9[7], nearly half of all mobile traffic is streaming video. And many of the other applications didn't even exist in 2000, like social networking, much of the online gaming industry, etc.

Figure 9: The world's most used apps by downstream traffic

So, Wi-Fi has moved from the periphery of the offering to being a critical component, and it needs to be able to deliver a lot more in a variety of areas, like higher throughput, lower latency, better overall coverage, and higher overall capacity. Take, for example, just the issue of coverage. In the early days of Wi-Fi, having some wireless connectivity was enough. Now, however, the Internet of Things (IoT) has become widespread, and many home devices (stoves, refrigerators, thermostats, alarm systems, water shutoff systems, sprinkler systems, electric vehicle charging systems, internet-connected photo frames, etc.) rely on available Wi-Fi coverage, so Wi-Fi has to be available “everywhere” in the home.

To address the flood of new requirements, the Wi-Fi standards bodies have added a host of new features to Wi-Fi to improve performance in these and other areas. In Figure 10, we show just some of these features (in grey) mapped along the four axes of throughput / latency / coverage / capacity, along with some relevant new technologies in the Wi-Fi standards (in yellow.) For example, while video calling is sensitive to both throughput and latency [and, therefore, may take advantage of the new higher data rates as well as the orthogonal frequency-division multiple access (OFDMA) functionality of Wi-Fi 6], something like residential IoT is sensitive mostly to coverage and therefore will be helped most by the newer mesh capabilities as well as the power save features such as targeted wait time (TWT). This figure is not intended to be exhaustive, and it’s certainly possible to argue that some of these features may be useful in other parts of the requirement space. The point of the figure is just to show that with exploding requirements for Wi-Fi networks has come an exploding list of new features to be tested.

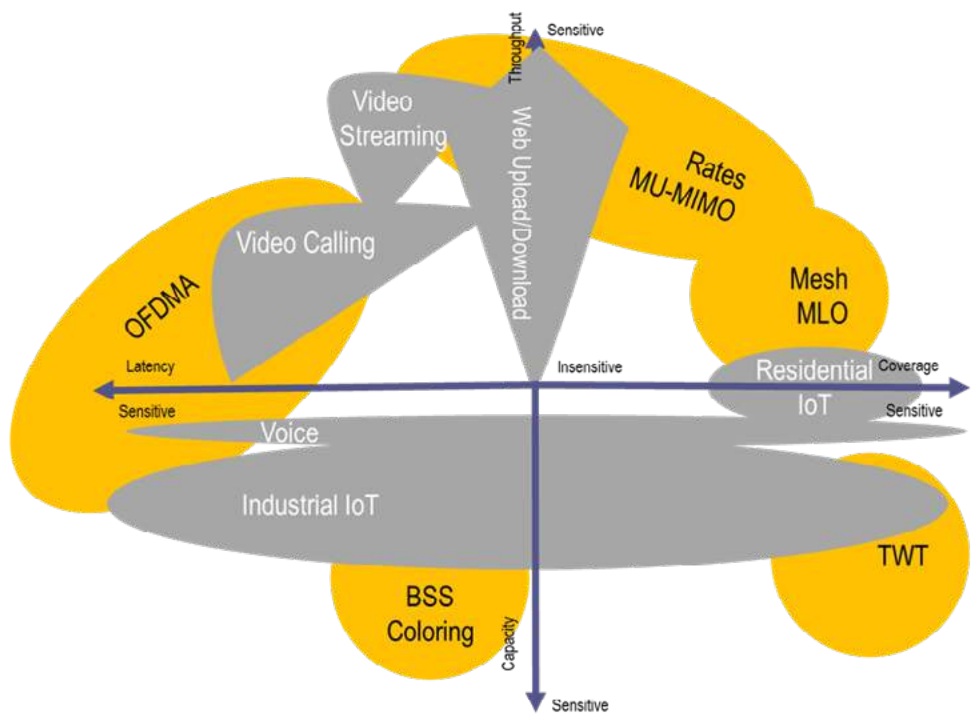

Figure 10: Some new Wi-Fi features mapped with applications

In the following sections we will discuss just a few of these new features in more detail.

4.2. Multi-user operation in Wi-Fi 6

The first multi-user Wi-Fi capability was introduced in Wi-Fi 4 with multi-user MIMO, and indeed some of the test plans described above do probe that functionality (e.g., TR-398). Even at that, Wi-Fi supports both downlink and uplink MU-MIMO, and so far, test plans only address the downlink capability.

In addition, one of the most talked-about new elements of Wi-Fi 6 is the addition of OFDMA as a channel access capability. OFDMA is also a multi-user feature, capable of allowing an access point to address and receive from multiple STAs simultaneously. One of the ways that this functionality can be useful is in terms of limiting per-STA latency since STAs would no longer need to wait for the AP to address each STA separately; rather, an AP could address many or all STAs at the same time.

Although, as discussed above, the BBF's TR-398 added in some Wi-Fi 6 testing functionality as part of Issue 2, that functionality was focused solely on the higher throughputs that Wi-Fi 6 offers. Looking at the performance of OFDMA has not been broached, and coming up with an implementable, repeatable scenario that can provide a quantitative measurement of an AP's OFDMA performance is a challenge that is, so far, for further study.

4.3. Advanced power save operation in Wi-Fi 6

Another highly touted feature of Wi-Fi 6 is what is known as Targeted Wake Time (TWT). This feature is an advanced version of the Wi-Fi's power save functionality that allows for very flexible sleep times for Wi-Fi devices. None of the performance testing has yet dealt with this functionality; indeed, none of the testing has yet addressed any of the power save functionality of Wi-Fi.

4.4. More features arriving in Wi-Fi 7

Wi-Fi 7 is right around the corner[8],and the number of new features in that specification will require whole new rounds of performance testing updates. For example:

- There are updates to basic modulations and bandwidths, so, higher throughputs.
- OFDMA has been upgraded to support more flexible spectrum allocations, so even though basic OFDMA functionality has not yet been addressed by these performance standards, Wi-Fi 7 will add more functionality to be addressed.
- TWT has been upgraded to provide restricted access to the channel for certain users and, so as with OFDMA, the TWT functionality has been upgraded before the performance testing specifications have even addressed the initial feature.
- One of the most highly anticipated features of Wi-Fi 7 is Multi-Link Operation (MLO), in which an AP/STA pair can communicate using a combination of channels across bands. There are a number of ways in which this functionality is likely to be useful, from increased throughput to interference robustness to low-impact band steering. Defining use cases, test methodologies, and metrics for MLO is going to be a major area of activity for at least some of the performance test specifications.

5. Challenges related to advanced feature testing

As should be evident from the above discussion, the difficulty with performance testing comes from the definition of "performance." There are some aspects of performance that can be understood fairly simply, and there are some that are much more complex. Compare the questions:

1. At high signal strength, what is the maximum throughput I can achieve between this AP and this STA?
2. How much more efficient is this product in Wi-Fi 6 mode than it is in Wi-Fi 5 mode?
3. What kind of experience does a client have, on a mesh system, while the mesh is supporting some set of other users, under specific traffic loads, with different interference conditions when roaming in a specific way?

The differences between these questions come down to the complexity of the question itself. The first question is fairly simple, and so leads to a fairly simple test: throughput is run between two points and measured. The second question is somewhat more complicated. What do we mean by "efficient"? Wi-Fi 6 is based on the IEEE 802.11ax specification, known as the "High Efficiency" enhancement[9], but what kind of test should be used to define "efficiency"? What metric should be evaluated for that case?

The final question indicates the level of complexity that these performance tests can elicit. While it is exactly the kind of question that a service provider would care about ("how well will this mesh system perform if I deploy it in the homes of my customers?"), the level of complexity is quite high. What do we mean by "well"? Is that throughput? Latency? A combination? Is it different for different applications? How do I create a mesh system in a compact testbed? How can I be sure that my emulated mesh system is a reasonable representation of what my customer will experience?

With the complexity of the question comes the complexity of the test itself. Compare the difference in testbed complexity between a testbed required to satisfy the first test (Figure 11)...

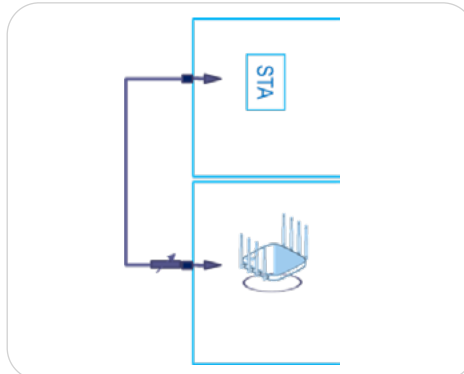

Figure 11: A testbed for running traffic between two endpoints

... and a testbed (Figure 12) that would be needed to answer the question about a mesh system with roaming stations (and possible interference conditions.)

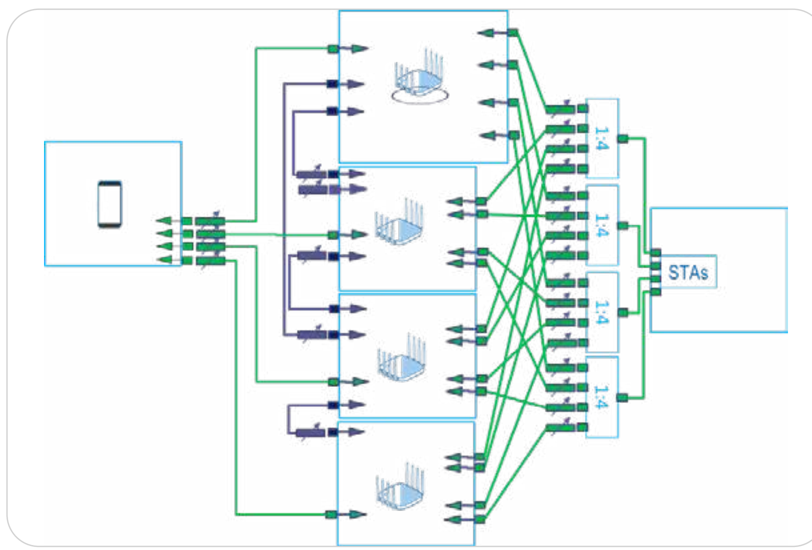

Figure 12: Example testbed for a 4-node mesh with multiple STAs

Also, increasing testbed complexity is not the only way in which the performance tests can become more and more complicated. The tests themselves get much more complicated and, as a result, rely more and more heavily on automation. To be clear, the assumption is that all of the tests can be automated (and usually are), but a simple throughput test (for example) can also fairly easily be run manually and therefore require fairly simple automation. The more complex tests, however, move beyond the stage for which manual testing can be reasonably expected to work.

5.1. Automation as a key component to Wi-Fi performance testing

To describe what we mean by testing complexity, and to illustrate how important automation becomes to enable these tests, let's consider the two more complicated tests described above: the OFDMA test and the mesh test.

5.1.1. OFDMA performance testing

The difference between OFDM and OFDMA is much discussed, but briefly, it is the difference between an AP being able to service only one user at a time in a given channel or multiple users. An example of a common figure is shown in Figure 13.

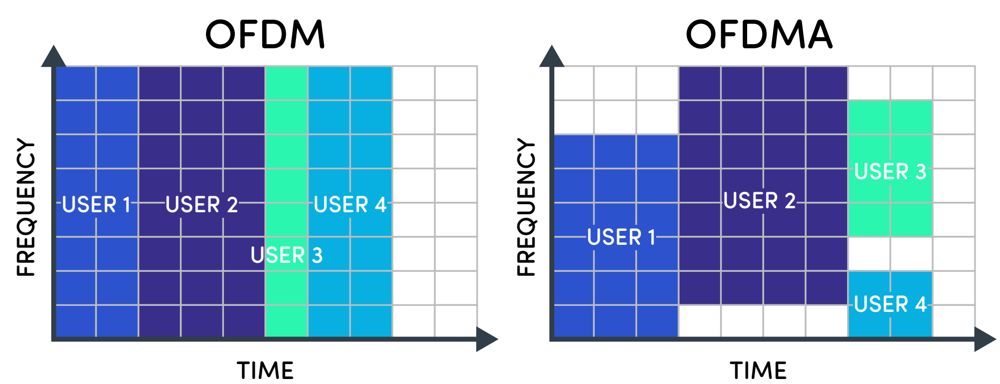

Figure 13: OFDM vs. OFDMA channel access

Notice that at any given time, in the OFDM case on the left, only a single user has access to any (actually to all) of the frequencies shown on the vertical axis. But in the OFDMA case on the right, there are times when both User 3 and User 4 are being served, with User 3 getting access to some of the frequencies and User 4 getting access to different frequencies.

These combinations of frequencies and times (represented by the small squares in the plots) are known as "resource units (RUs)" in Wi-Fi. To address the question asked above ("how efficient is the Wi-Fi 6 operation?"), it might be very useful to actually look at the RU allocations to see how much of the channel is actually occupied. So, for example, we may want to look at how the RUs have been allocated to the STAs.

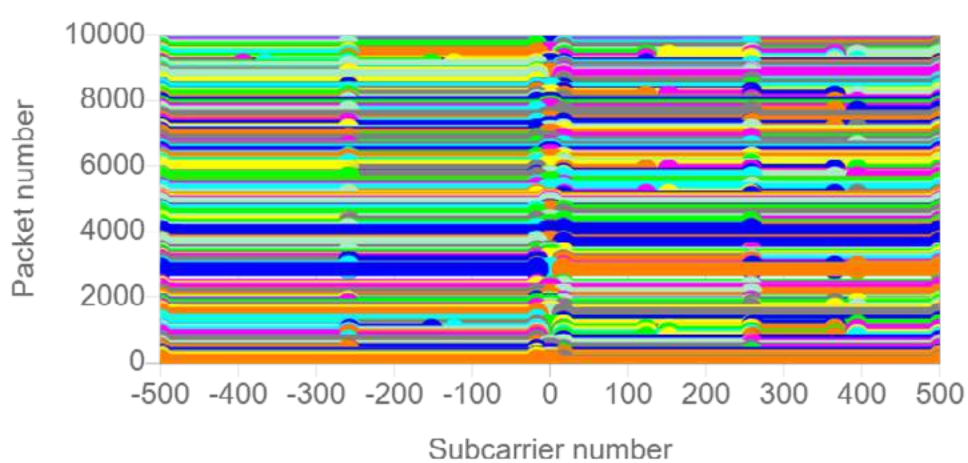

Figure 14: Example analysis of RU allocations in an OFDMA system

And we may care about, for example, how efficiently the channel was allocated to the different users. That is, how much of the channel is actually occupied.

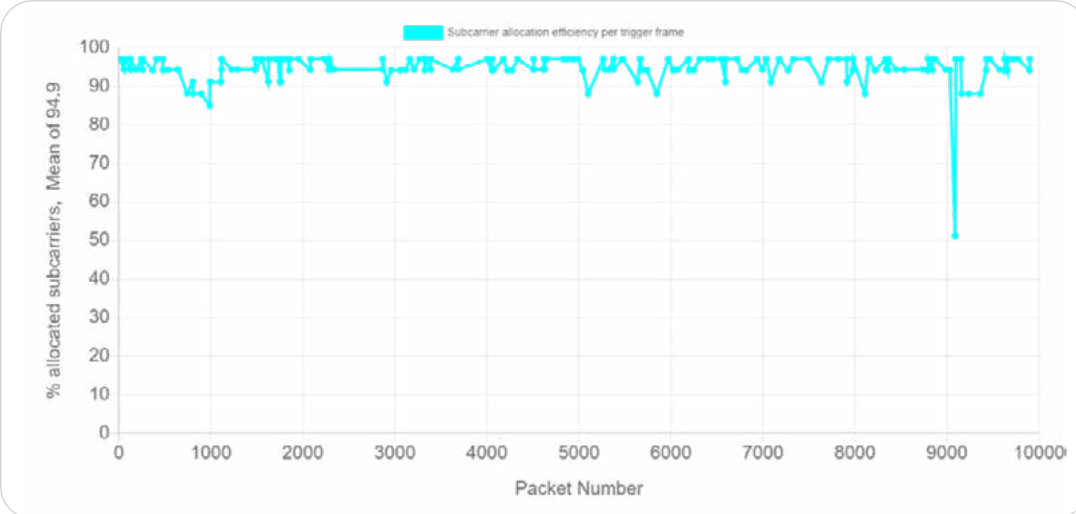

Figure 15: Example analysis of channel allocation efficiency in an OFDMA system

Doing OFDMA analyses can be much more complicated than doing, for example, throughput measurements. RU allocations are only visible "under the hood" by looking at the actual packet stream and decoding the OFDMA information. More than that, not all OFDMA information is visible to a standard Wi-Fi sniffer. A simple OFDMA setup is shown in Figure 16.

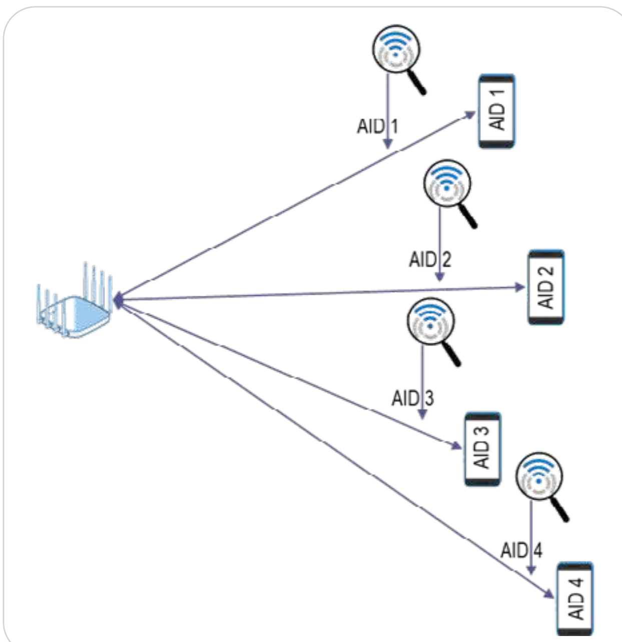

Figure 16: OFDMA communication in Wi-Fi, with sniffing indicated.

In this figure, we show an AP communicating with four STAs using OFDMA. In this process, the STAs have received an "association ID (AID)" at association time, and this is used to indicate the RU that it should use on a packet-by-packet basis. That AID is buried in the signaling information and needs to be pulled out in order to even do this kind of analysis. And the OFDMA traffic flowing between the two endpoints will not even be visible to a Wi-Fi sniffer unless the sniffer knows, and has been provisioned with, the AID being used. So, if we want to watch the traffic, we can no longer use a single sniffer. We will need to use four separate sniffers, each with a separate provisioned AID.

The complexity of this test should, by now, be obvious:

- Set up the OFDMA traffic
- Figure out the assigned AIDs
- Extract RU information, per AID, on a per-packet basis
- For OFDMA traffic, use separate sniffers with correctly provisioned AIDs to gather that traffic from that specific AID
- Combine all of this information into a coherent result in order to answer the relevant performance question

The sheer level of complexity involved in this test means that without a way of fully automating the test and analysis, this kind of test would not be practical or even possible.

5.1.2. Roaming Performance

The roaming performance test is another good example of how the more complicated the test, the more important automation will become.

As discussed above, the ETSI Multiple Access Points test plan focuses on the performance of a system designed for roaming. However, many of those tests are fairly straightforward from a conceptual perspective. Take, for example, the "two stage networking test" (Figure 17). This test is designed to look at the throughput achieved on a STA/AP link when the STA is separated from the AP by a two-hop mesh network.

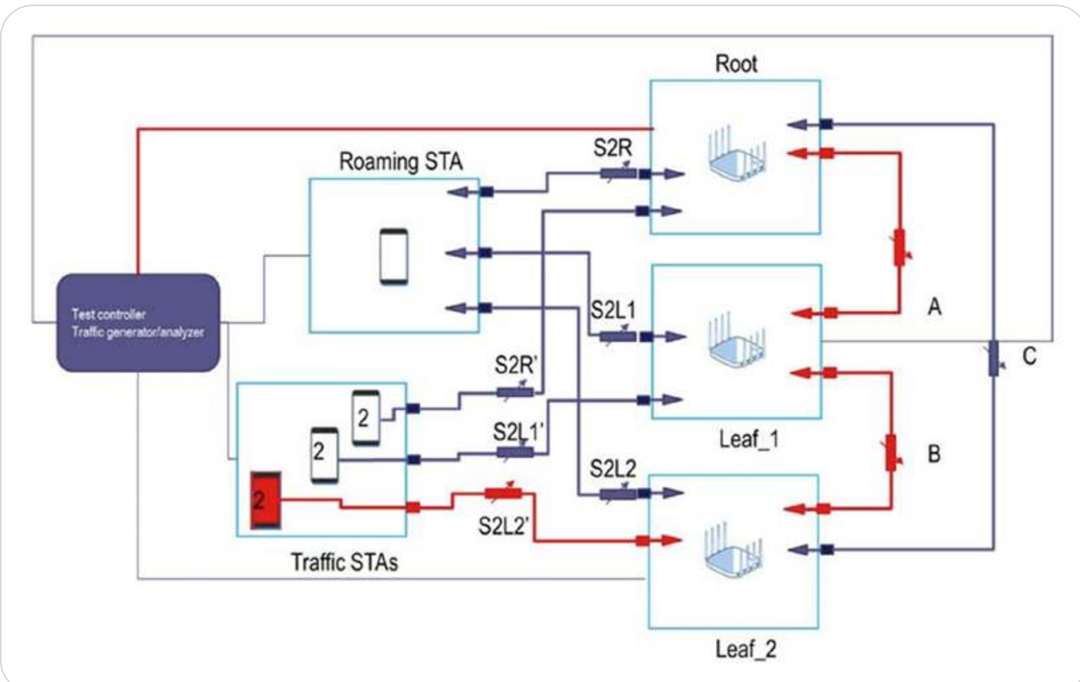

Figure 17: Multiple Access Points Performance, Two-Stage Networking Test

In this test, the red elements identify the test configuration. Traffic is run between a STA (the phone icon shown in red) and the Root AP, passing first between leaf nodes Leaf_2 and Leaf_1. This is a two-hop mesh throughput test. While this takes some work to configure, it is not difficult to imagine how this could be performed manually.

Now let's consider the Device Metrics "Roaming" test. This test does not fix the STA statically on one element of the mesh. Rather, it expects the STA to move between the elements of the mesh, thus "roaming" from node to node. The defined topology (Figure 18) is simpler than the above:

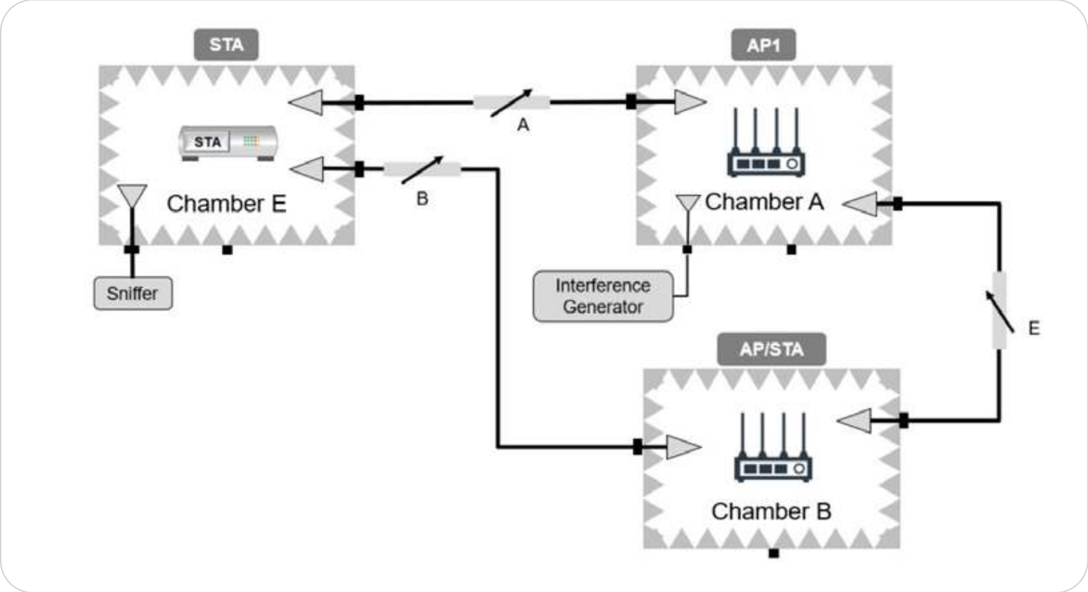

Figure 18: Wi-Fi Alliance Device Metrics Roaming Topology

However, the test itself is much more complicated. The STA is caused to roam between chamber A and chamber B by modifying the programmable attenuators A and B. As described in the test case, the attenuations used come from real-world measured data, and the "recording is then used to play back the walk through the attenuators creating a highly repeatable, lifelike test." Already, a test like that is virtually impossible manually because of the requirement that the attenuators be modified in this highly specific way. The test then also looks at the throughput achieved over multiple runs as a function of time. This data can be different from run to run based on the roaming algorithm and the probabilistic nature of the decisions it makes. An example from the test plan is shown in Figure 19.

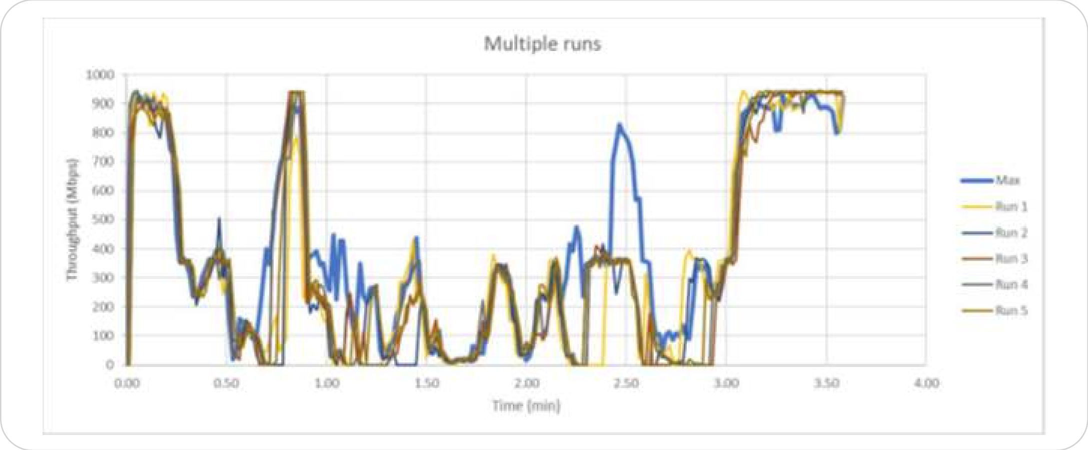

Figure 19: Maximum roaming throughput achieved over a set of runs

Adding on the statistical nature of these tests and the resulting analysis required, it becomes clear that without automation, the more useful Wi-Fi performance tests would not be possible.

6. Conclusion

The Wi-Fi community has finally embraced performance testing. Since 2019 three of the most important standards organizations (BBF, Wi-Fi Alliance, ETSI) have developed and published a specification targeting Wi-Fi performance. It remains, however, early days for Wi-Fi performance tests. Many of the existing tests are simple speed tests, albeit sometimes under more complex conditions.

But Wi-Fi is used in challenging environments under challenging conditions. Almost always used in shared spectrum, Wi-Fi must perform well even when interference exists from other sources (either Wi-Fi or not.) Wi-Fi is the main access network by which people experience their broadband connections, so Wi-Fi must be adept at handling high throughput connections for movies and downloads while also (and simultaneously) being able to handle low latency connections (for video calls and gaming). Creating test cases to accurately capture these conditions while identifying and measuring the relevant metrics will be the work of these groups, and undoubtedly others, as the need to understand Wi-Fi performance continues.

7. Abbreviations and Definitions

7.1. Abbreviations

AID	association ID
AP	access point
BBF	Broadband Forum
BRAN	Broadband Radio Access Network
bps	bits per second
ETSI	European Telecommunications Standards Institute
Hz	Hertz
IEEE	Institute of Electrical and Electronics Engineers
IoT	Internet of things
MLO	multi-link operation
OFDMA	orthogonal frequency division multiple access
QoS	quality of service
RCPI	received channel power indicator
RRM	radio resource measurement
RU	resource unit
SCTE	Society of Cable Telecommunications Engineers
STA	"station," a non-AP device in a Wi-Fi network
TWT	target wake time

7.2. Definitions

Downstream	Information flowing from the hub to the user
Upstream	Information flowing from the user to the hub

8. Bibliography

TR-398 Wi-Fi In-Premises Performance Testing, Issue 1, February 2019

TR-398 Wi-Fi Residential & SOHO Performance Testing, Issue 2, March 2021

Wi-Fi Device Metrics, v1.0, 2022

Broadband Radio Access Networks (BRAN); Multiple Access Points Performance Testing, ETSI TS 103 754 V1.1.1 (2022-06)

9. References

[1] https://www.wi-fi.org/downloads-public/Wi-Fi_by_the_numbers_2023.jpg

[2] https://wifinowglobal.com/news-and-blog/wi-fi-is-the-best-invention-of-the-past-25-years-says-uk-study/

[3] https://www.mediaplaynews.com/parks-92-percent-of-us-internet-households-use-wi-fi-at-home/

[4] https://www.wi-fi.org/system/files/Wi-Fi_Device_Metrics_Highlights_202209.pdf

[5] https://www.etsi.org/deliver/etsi_ts/103700_103799/103754/01.01.01_60/ts_103754v010101p.pdf

[6] https://www.broadband-forum.org/technical/download/TR-398_Issue-2.pdf

[7] https://www.visualcapitalist.com/the-worlds-most-used-apps-by-downstream-traffic/

[8] https://www.wi-fi.org/who-we-are/current-work-areas#Wi-Fi%207

[9] https://standards.ieee.org/ieee/802.11ax/7180/

© 2023 Spirent Communications, Inc. All of the company names and/or brand names and/or product names and/or logos referred to in this document, in particular the name "Spirent" and its logo device, are either registered trademarks or trademarks pending registration in accordance with relevant national laws. All rights reserved. Specifications subject to change without notice.