

Characterizing Wireless Network Performance



Ruckus Wireless | Black Paper

Accurate performance testing for wireless networks requires understanding how to test for worst case scenarios

As expensive and inconvenient as wires may be, when it comes to transmitting data, they are generally quite reliable. Once working, they tend to stay working, offering the same steady performance level until someone cuts them, unplugs them or digs them up.

However, modulated electromagnetic waves propagating in free-space (aka radio waves or Wi-Fi signals), are anything but steady. They interact with the environment through exotic physical mechanisms such as reflection, refraction, fast fading, slow fading, attenuation and ducting. Even with the best wireless protocols, the best chipsets, the best RF design, the best software and the smartest antennas, wireless performance is going to vary — and it's going to vary a lot.

There will always be some location where, if the stars align, a user can achieve that magical, maximum physical layer throughput number of 300 mbps, in the case of two-stream 802.11n, for instance. But go far enough away from an access point (AP) and performance is guaranteed to eventually drop to 0.000 mbps. And everywhere else, where people are actually located, there will be performance levels of everything in between.

Performance can also vary significantly, even in a fixed location, due to motion in the environment, interference and the random background noise that came with this universe. Wireless performance is inherently statistical in nature, and accurate performance testing must account for this random component.


The goal of wireless performance testing

When testing for wireless performance, there are generally two scenarios:

1. testing a single system to determine if it meets some minimal (usually application-centric) performance criteria or;
2. comparing two systems to determine 'which is better'


In both cases, the intent is for the testing to predict the real-life performance of the systems, once fully deployed.

The following graphs illustrate 'normal' probability distribution, describing at least approximately, any variable that tends to cluster around the mean. Shown as the familiar 'bell curve,' the first graph (Fig. 1) depicts four different versions, each




Characterizing Wireless Network Performance

Figure 1
Bell curve graph depicting different average and standard deviation parameters.



with different average and standard deviation (variability) parameters. This type of plot is known as a ‘Probability Density Function’ (PDF), because it shows the relative probabilities of getting different values. The Y axis legend is Greek for ‘relative probability of seeing the values on the x-axis’. It answers the question, “What is the chance I will see a certain result?” If you were examining some random process represented by the red curve, one would expect outcomes with a value around ‘0’ to be twice as prevalent as outcomes of around 1.25 (40% versus 20%). In many cases, the more interesting question is, “What is the chance I will see a result less than or greater than a particular value?” A different, but related graph (Fig. 2), the “Cumulative Distribution Function” (CDF) helps answer this question. Take a look at the following graph that shows the corresponding CDF plots for the previously shown bell curves. (Note: a PDF graph,

Figure 2
Typical cumulative distribution function (CDF) graph.




can always be converted into a CDF graph, or vice versa. They are just different ways of plotting the same information).

This chart shows, for example, that the probability a process (represented by the red curve) will produce a result of less than -0.75 is about 20%. The probability of the red curve producing a result less than 0 is 50% (as you would expect from looking at the bell curve), and the probability of producing a result less than 2 is about 95%.

To characterize wireless performance, CDF graphs and the information that goes into creating them are immensely useful. Formulating a CDF graph for a given wireless product helps predict the percent of time, or the percent of locations, at which performance for that product will be above or below a certain threshold.

Consider the following plot (Fig. 3).

Figure 3
Example of a complementary CDF graph that helps to predict the percent of time different wireless throughput can be achieved in a given location.




Technically this is a CCDF (the extra ‘C’ standing for ‘complementary’) plot. This graph shows the probability of getting different wireless throughputs with a certain Ruckus AP product in a particular location (remember that wireless performance is statistical).

Characterizing Wireless Network Performance

Based on this graph, one would expect a 95% chance of getting 80 mbps or better. If the application demanded 40 mbps or better, this curve predicts that the AP would achieve this data rate over 99% of the time. This is useful in planning any wireless infrastructure, especially when certain traffic types, such as streaming video, require a constant bit rate at a particular speed.

Network planners can also use the CDF plots to compare two systems. Consider the following plot (Fig. 4):

Figure 4
CDF chart comparing a competitive video adaptor versus the Ruckus MediaFlex 7111.




This CDF plot predicts, for instance, that the WNDHE111 video adaptor has about a 50% change of doing 20 mbps or better while the Ruckus 7111 has better than 95% chance of doing the same.

Vendors can also use CDF plots to test and improve the performance of a given product. This chart actually shows the performance of two different antenna options on the same hardware.

In the following graph (Fig. 5), the two antenna options both perform well in 'good' locations — about 50% of the time the performance is very similar. At 'harder' locations, however, the

Figure 5
CDF chart comparing the performance of different antennas on the same product.



difference between the two antennas is dramatic. This critical information would be lost without a way of looking at the full wireless performance statistics.

Given that these CDF plots are so useful for comparing and predicting the performance of wireless systems, the question becomes, how to generate them? Unfortunately, network planners or product developers cannot calculate or simulate these plots because the real world of wireless is too unpredictable given the effects of a constantly changing environment, movement and interference.

To understand wireless performance, real-world testing is essential. However, it must be performed in a way that exposes the underlying performance statistics. Ruckus Wireless has developed a wireless performance testing tool, called "Zap," precisely for this purpose.

What is Zap and how does it work?

Zap is the culmination of several years of experience, field-testing wireless performance for IPTV-over-wireless applications. IPTV is a demanding application where knowing



Characterizing Wireless Network Performance

the 'average' throughput of a wireless system just isn't sufficient enough information to predict its performance. Any dip in performance, no matter how short, can result in dropped packets that cause visual artifacts or pixilation of the image. Consequently, the viewer experience is completely dependent on the wireless system's 'worst-case' throughput (e.g., the 99th percentile).

Prior to Zap, the focus of existing tools in the market has been on measuring average throughput, not worst-case throughput. Ruckus engineers originally designed Zap to measure and predict what type of performance they could expect most of the time (not just some of the time), using a large number of samples.


Sampling is the key to recovering the statistical performance and drawing the CDF curve of a wireless system. For the CDF to predict real-life performance accurately, network planners must conduct sampling tests across all relevant dimensions and variables. In most cases, planners must sample three dimensions to characterize wireless performance accurately: time, space, and frequency.

As mentioned earlier, wireless performance can vary significantly as a function of time. Even when wireless endpoints are stationary, performance will vary from moment to moment due to:

- random physical layer errors
- motion in the environment
- antenna or physical layer data rate selection algorithms
- environmental noise or interference, and
- other low-level radio changes related to temperature fluctuations or ongoing internal chipset calibration

The following plot (Fig. 6) shows an example of this effect. Time moves to the right across the x-axis and the y-axis shows 'instantaneous' throughput. As shown in the graph below, there are significant variations in throughput during this two second test. At one instant 140 mbps of throughput might be achieved while 100 milliseconds later only 90 mbps is possible. A longer test would reveal even more variability.

Figure 6
Throughput over time comparison of two different APs.



Zap works by time-based sampling of the wireless channel performance. Zap sends controlled bursts of UDP packets, measuring the time between the arrival of the first packet in the burst and the last packet in the burst. Based on the size of the packets, Zap can calculate the instantaneous throughput for that burst. It continuously repeats this process with additional packet bursts for the duration of the test.


A typical test duration would be 30 seconds, but a particularly noisy or highly variable environment might require a longer test to reveal the complete link statistics. A 30-second test would reveal about 600 throughput sample data points. The main caveats to this approach are that all involved network devices (e.g., Ethernet switches) must have sufficient buffering to absorb the packet bursts and only a single wireless link can be tested at a time. Obviously, the wireless link needs to be the actual performance bottleneck in order for the results to be valid.

At the end of an individual test run, Zap sorts the sample throughputs from highest to lowest and saves them to a file. These 'percentile results' can then be easily plotted as a CDF or combined with results from other locations to plot an 'uber CDF'. The graph on the following page shows the exact same data from the previous "Throughput versus Time" (Fig. 6) test result in CDF format (Fig. 7).

With the time dimension known, the question becomes, how does performance vary as a function of location?

Characterizing Wireless Network Performance

Figure 7
Throughput over time depicted in a CDF chart.

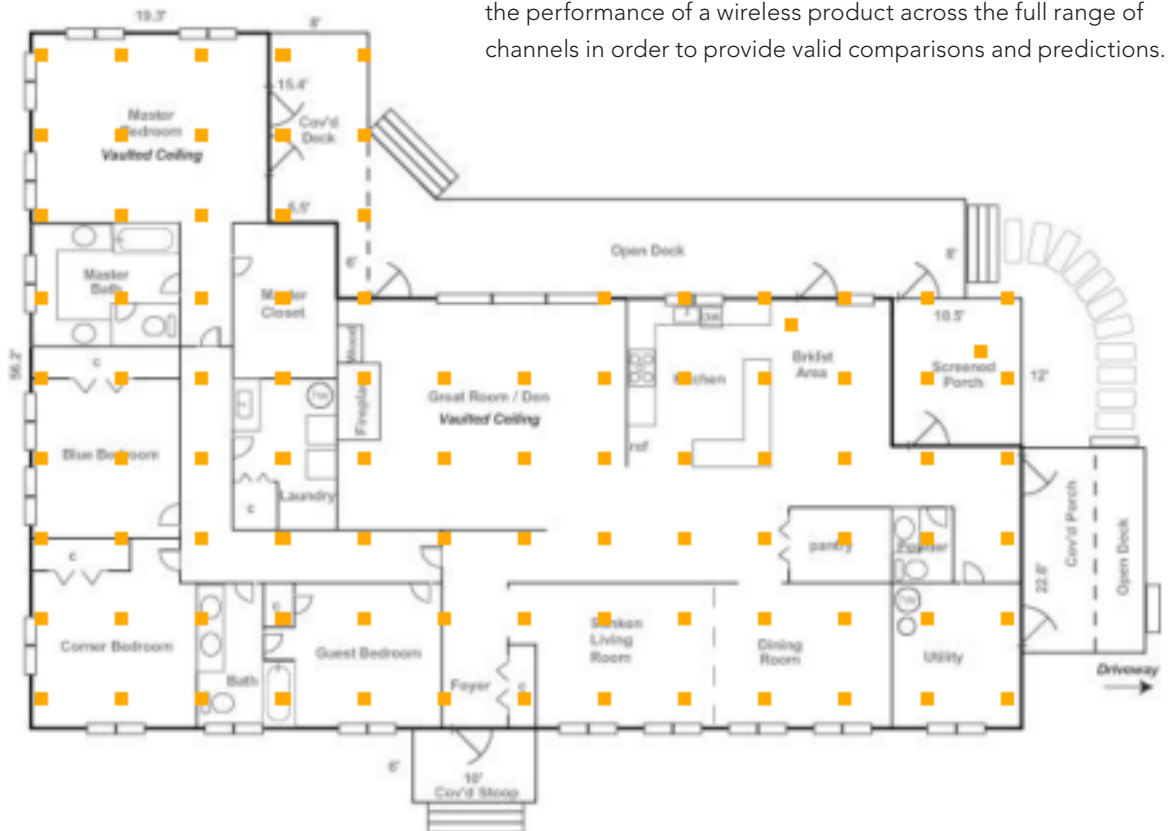


The key to incorporating spatial variability into wireless performance statistics is to get sample performance information at a large number of locations. For instance, if a wireless AP product is targeting whole-home coverage, one approach would be to conduct performance tests at regular intervals throughout a home and combine the results into a single CDF. The floorplan below (Fig. 8) illustrates this concept with orange squares indicating client test locations where planners could perform individual runs of the Zap testing tool.

While this represents a lot of testing, experienced planners can eliminate many of these test locations. However it is critical that the redacted set of test locations have the same distribution of easy, medium, and hard locations as would have been found in the original full set. Furthermore the final set of locations must still be fairly large if the resulting CDF is to be valid.

Frequency (e.g., the operating channel) represents the final performance test dimension. Even in the absence of external interference, the performance of wireless products can vary dramatically from channel to channel due to a variety of factors such as regulatory power limits, local digital noise, and RF component variation. This is especially true in the 5 GHz spectrum. It is critical that network planners sample the performance of a wireless product across the full range of channels in order to provide valid comparisons and predictions.


Figure 8
Floorplan showing the different client test locations required to account for spatial variability.



Characterizing Wireless Network Performance

Figure 9

Chart depicting throughput as a function of a given operating channel.

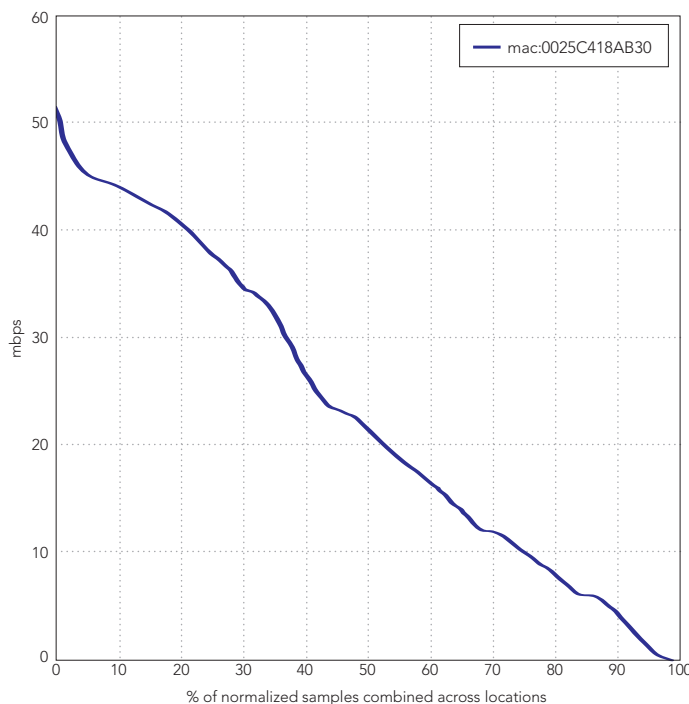


The chart above (Fig. 9) illustrates this throughput variation based on actual measured results from a long distance point-to-point link.

In practice, this channel-to-channel variation means that planners must test the devices on each channel at every test location. The multiplicative implications of this on test cases are severe, so most people with traditional network testing tools usually give up before this point as the large amount of results data is too cumbersome to absorb.

Figure 10

CDF chart, depicting combined percentile results across all test runs.



With Zap, sampled output can be easily combined across multiple locations. Essentially, all the temporal, spatial, and channel dimensions are thrown together into the same pot via the Zap output file. Combining the percentile results across all of the test runs allows Zap to characterize the wireless performance as a CDF (Fig. 10), providing an accurate and easily digested representation of the wireless performance across the sampled conditions.

Summary

Even with the best wireless gear, characterizing Wi-Fi performance is difficult at best. Reflections, refractions, signal fading and attenuation all wreak havoc on throughput causing performance to vary widely.

Because wireless performance is inherently statistical, accurate performance testing must account for this random component.

Ultimately, real-world wireless testing is essential, but this testing must be performed in a way that exposes the underlying performance statistics, looking beyond average throughput.

Sampling is the key to recovering the statistical performance and must be conducted across all relevant dimensions. Time-based sampling of the wireless channel, sampling at a large number of locations and sampling across the full range of channels are the keys to providing valid comparisons and predictions.

Zap is a wireless performance testing tool developed precisely for this purpose. Zap allows any organization to better understand the statistical throughput distribution of a wireless system in order to characterize performance. With Zap, organizations can now easily test sustained throughput of an existing system and predict the real-life performance of a planned system before deployment.

By enabling an accurate determination of the true, sustained and worst-case performance that a wireless network can deliver 99.5 percent of the time, companies can become more confident in knowing that their wireless network will adequately support the more stringent application requirements that exist and the quality of service that users have come to expect.

Ruckus Wireless, Inc.

880 West Maude Avenue, Suite 101, Sunnyvale, CA 94085 USA

(650) 265-4200 Ph \ (408) 738-2065 Fx

