The Road to Gigabit+ Performance

Joe Zeto, Sr. Manager Market Development, Ixia

Chip and hardware developers have their hands full with the next generation of the 802.11 standard – IEEE 802.11ac. The deployment of 802.11ac promises to deliver higher bandwidth while retaining better quality of experience (QoE) for end users, as well as some significant technical challenges for developers. However, once they have overcome them, 802.11ac will be adopted in all major markets including residential, enterprise, carrier and large venue.

How Does 802.11ac Meet Demands?

There are several advancements that allow 802.11ac to achieve the targeted performance. The three driving initiatives are the enabling of multiple flows to use the medium concurrently, the increase in raw bandwidth, and the optimization of performance to specific clients. The caveat of these advances, however, is that 802.11ac must still be compatible with existing legacy and 802.11n client devices.

Moving forward, an 802.11ac access point (AP) will be expected to carry on conversations with multiple 802.11a phones, 802.11n tablets, and 802.11ac laptops. Each of these different devices supports a variety of features including QoS, power save, and multicast. The new 802.11ac chipsets will have to simultaneously excel at delivering blistering performance to other 802.11ac devices, while gracefully interoperating with several previous devices.

Enabling Multiple Flows

802.11ac can support more users because it no longer requires that just one 802.11 device transmit simultaneously. Multi-User MIMO (MU-MIMO) is a new technique where multiple client devices are receiving packets concurrently. In fact, it is the first time in Wi-Fi’s history that directed traffic can be delivered to multiple client devices simultaneously. MU-MIMO works by directing some of the spatial streams to one client while directing other spatial streams to a second client. Although there is some variation when it comes to this principle, MU-MIMO is critical to performance improvements in environments with high client counts.

Increasing Raw Bandwidth

A higher rate encoding scheme known as 256-QAM allows 802.11ac to increase the physical-layer transport rate, transmitting 33 percent more data than the 64-QAM used in the 802.11n standard.

802.11ac also introduces channel bonding approaches to further increase the amount of data transported per second. Although channel bonding approaches were

made popular in 802.11n, they have been expanded to provide 80 MHz-, and ultimately, 160 MHz-wide channels, allowing for more data to be transmitted simultaneously out of the same antenna. Compare that to 802.11n, when users could only select between a 20 MHz or 40 MHz channel operation. Importantly, these wider channel bandwidths – and the need for proper channel separation – mean that 802.11ac can only be used in the 5.0 GHz band (where more non-interfering bandwidth is available). While dual-band APs will still be produced, the 2.4 GHz band will be limited to 802.11bgn and will not be able to be configured for 802.11ac.

Adding spatial streams are also a key factor for increasing raw bandwidth of 802.11ac. The 802.11ac standard allows for up to eight spatial streams, a significant expansion from 802.11n's four spatial streams.

Lastly, packet aggregation, although also present in 802.11n, is worth mentioning because it is often the single biggest performance multiplier on a per-transmission basis. Once a high performance device obtains its transmit opportunity, the transmitter strings multiple frames together, and transmits them in succession without having to reacquire the medium.

Individual Client Channel Optimization

Finally, 802.11ac also gets a major performance boost from technologies that optimize communications when speaking to a specific client. A concept known as transmit beamforming (TxBF) allows the access point to communicate with the client devices to determine which types of impairment are present in the environment. After which the AP "precodes" the transmitted frame with the inverse of the impairment such that when the next frame is transmitted and transformed by the medium, it is received as a clean frame by the client. And since no two clients are in the same location, TxBF needs to be applied on a client-by-client basis and constantly updated to reflect the changing environment.

Putting it all Together – Technical Challenges for Developers

Developers will need to expand the complexity and precision of their designs in order to deliver breakthrough 802.11ac performance. What's more, it requires a rethinking of the traditional approach to testing due to the introduction of new technologies and the overall technical complexity. 802.11ac testing demands coordination and control between the different layers of the protocol stack. Without it, it becomes difficult to exercise functions and to quickly pinpoint performance issues to a specific function of the hardware.

802.11ac is pushing us into a new generation of testing. This testing should be able to decode every frame in real-time and determine each frame's RF characteristics, as well as their frame-level performance, and generate every frame without limitation in real-time to adequately test the receiver performance. It also needs to be able to tightly integrate RF and MAC functionality in 802.11ac, and include integral, real-time channel emulation to address TxBF performance.

Layer 1 Testing

246-QAM and wider bandwidths are both extremely challenging changes for radio designers to address. They must deliver performance advances in virtually every dimension that impacts digital modulation performance – phase noise performance, noise floor, and modulation accuracy. Best-in-class performance means verifying transmit and receive performance of literally hundreds of frame definitions, varied by modulation rate, frame length, bandwidth, frequency, channel model, and power level. Plus, the testing must include legacy frames, 802.11n frames and 802.11ac frame in various combinations, and must be representative of the actual diversity, rate, and complexity that actual devices will experience in the field. Yikes!

There are two critical components that developers need in order to test these highly-integrated and high performance 802.11ac radio systems on a timeline that meets the aggressive market pace. One, the ability to make RF measurements on every frame transmitted at line rate and two, the ability to generate any desired frame quickly and without limitation. The tables below list the measurements required before the radio transmit and receive performance can be declared complete.

Key 802.11ac RF Transmitter Tests

Area	Measurements
Power	Average, peak, power spectral density, power peak excursion, power-on / power-down
Frequency	Center frequency tolerance, Symbol clock frequency tolerance, preamble frequency error, RF carrier suppression
Spectral	Tx spectrum mask, Spectral flatness, Tx center frequency leakage, CCDF, Occupied Bandwidth
Modulation	Constellation error, EVM, Transmitter modulation accuracy
I/Q	Gain mismatch, Phase mismatch

Key 802.11ac RF Receiver Tests

Area	Permutations
Tone Generation	Frequency, amplitude
Frame Generation	Frame sequences, rate, length, encoding
Modulation	a/b/g/n/ac PHY rates, preamble, FEC, etc.
Impairments	Frequency offset, pre/post encoder bit errors, channel models, etc.

Layer 2 Testing

Adding to the difficulty, 802.11ac also includes a great deal of complexity at the MAC layer. Underperformance can be due to any number of potential causes including: slow ACK response times, poorly designed aggregation algorithms, internal buss limitations, poor rate adaptation algorithms, poor AP selection algorithms, power save implementations, poorly implemented legacy protection schemes, etc.

Similar to RF testing, layer 2 functionality testing should be simple in the beginning and gradually add more complexity. For 802.11ac specifically, that means progressing through the following series of steps, all run under cabled conditions:

· Ensure that a single client is able to reliably connect/disconnect;

· Conduct benchmark tests using a single client to understand any basic system bottlenecks in the upstream and downstream direction;

· Activate more features on the single client to ensure that the basic functions work properly (aggregation, power save, IPv4 and IPv6, QoS, etc.);

· When testing an access point solution, benchmark with multiple clients to ensure that the system can achieve the expected performance under ideal conditions at scale

· Test with mixed-mode clients to ensure that the introduction of legacy devices does not adversely impact your 802.11ac solution.

The last two steps are to execute the system level tests while simultaneously varying the individual RF channels of each client.

First, the test should be run in a cabled environment with a mix of clients in a variety of locations. It is important and necessary to create a wide variety of conditions without necessarily using numerous client devices in a physical layout. The cost and repeatability of performing this testing prohibits using actual facilities. There are a set of channel impairment models developed by IEEE members that can be used, along with adjustments to client power levels, to virtually create a layout and conduct this testing.

Second and lastly, the same tests must be rerun over the air in an RF-isolated environment to eliminate any unintended interference. This ensures that the integrated system, including antennas, cabling, etc., continue to perform as expected.

Testing Multiple Layers Concurrently

It is logical that more comprehensive testing can be accomplished much faster if one tests at multiple layers of the protocol stack concurrently. The underlying medium of wireless technologies will likely experience issues that will affect the performance of the overall system. Knowing what to attribute performance degradation to – an RF design issue, an upper layer protocol issue, or the result of impairment in the RF spectrum – is the real challenge.

In order to conveniently address these issues, a single solution should enable testing of each layer individually while simultaneously allowing visibility into metrics from both the RF layer and the upper layers. This approach allows rapid identification and isolation of any discovered issues, and provides the confidence that a product is really ready for production release. By using a common set of

tools, productive communication is facilitated between these two functions and test results are easily duplicated.

The Road to 802.11ac Performance

The enterprise, residential and carrier markets will undoubtedly benefit from 802.11ac's gigabit+ performance. But before realizing these benefits, chip and hardware developers have some technical challenges to conquer, many of which can be addressed with a re-thinking of the traditional approach to testing as described above.

www.ixiacom.com [1]

Posted by Janine E. Mooney, Editor

June 22, 2012

Source URL (retrieved on *04/24/2014 - 5:04pm*):
http://www.wirelessdesignmag.com/blogs/2012/06/road-gigabit-performance?qt-blogs=0

Links:
[1] http://www.ixiacom.com