WLAN Performance Analysis

by Richard O'Dea, product manager, Spirent Communications
The growth of wireless networks is being driven by the need to lower costs associated with network infrastructures and to support the mobile networking applications that offer gains in process efficiency, accuracy, and lower business costs. Worldwide wireless LAN (WLAN) hardware revenue exceeded $390 million in the second quarter of 2002 with WLANs showing strong penetration into enterprise and consumer sectors. But the really big news could be in the service provider sector. As the popularity of public access Internet hot spots grows, the debate as to whether WLAN competes with or compliments 3G deployment is still being hotly contested.

This being said, the advent of WLAN technology brings new challenges to the forefront for semiconductor chip manufacturers, network equipment manufacturers (NEM), service providers and network administrators. The risk of failed networks due to untested or unproven devices or equipment performance is high. Before shipping new WLAN products or deploying live WLAN networks, certain questions need to be answered. For example, how effectively is interference handled in the 2.4 GHz RF range? What sort of roaming hand-off delay will there be as mobile users move from access point cell to access point cell? How will the implementation of various MAC and upper layer security mechanisms impact product and network performance? How effectively does the WLAN station and access point handle power management? All of these issues, along with traditional end-to-end performance, latency and packet loss, can be addressed through comprehensive testing.

The focus of this article is on addressing the importance and benefits of WLAN performance testing and will attempt to specifically address the value and benefits of testing some of the key areas of concern relative to WLAN device development and network deployment.

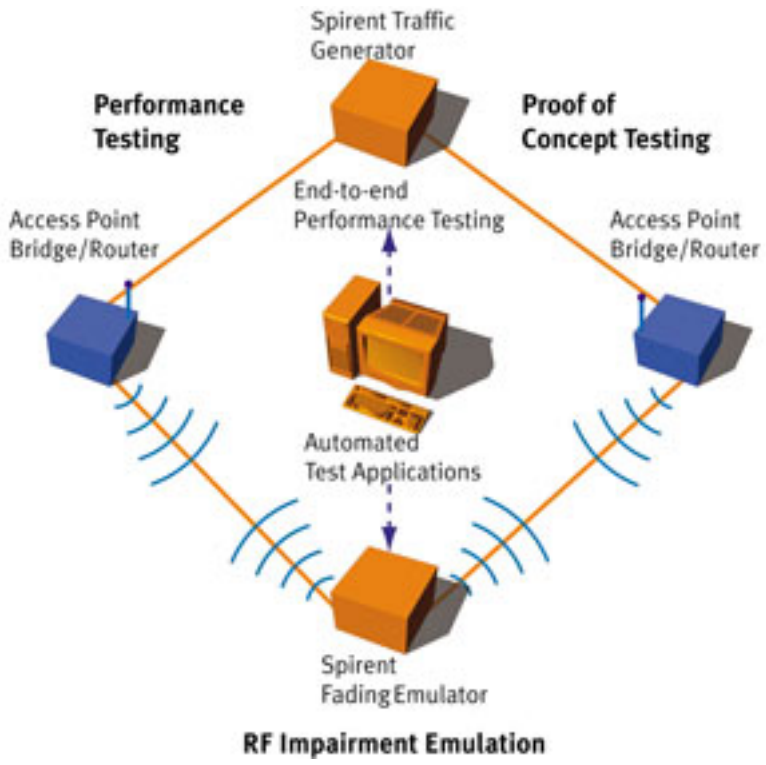

Traffic Generation and Impairment Emulation for an AP-to-AP Testing Environment

Standardized, real-life RF propagation modeling

The biggest challenge facing WLAN chip and equipment manufacturers will be the development of robust, high performance, physical layer implementations using 2.4 GHz and 5 GHz radio frequency signaling. Although 802.11b and 802.11a are rated at 11 Mbps and 54 Mbps of potential throughput respectively, actual throughput will vary greatly depending on how well the physical layer interface deals with RF channel characteristics such as multi-path fading, path loss, delay spread and log normal shadowing. Laboratory recreation of real world radio frequency impairments will be required by developers and service providers so that true product performance can be determined. This can be a significant challenge as the environmental variables that can impact radio frequency transmissions are far more difficult to recreate than a simple twisted pair, copper wire connection. Take, for example, the issue of multi-path interference. Radio waves can reflect off of solid objects and show up at the WLAN AP or W-NIC receiver slightly delayed from the original signal causing inter-symbol interference and degrading the signal quality. Typical reflected signal delay in an office environment is 50 nsec. and 300 nsec on a manufacturing floor. Inter-symbol spacing can be used to counter the effect of multi-path interference but at a cost to throughput. Additionally, flutter may occur as the radio tries to lock on to the reflected signals as well as the original.

Multi-path fading emulators can reproduce these types of RF characteristics, recreating real-world impairments against which the System Under Test (SUT) can be evaluated in the lab. The connection to the RF channel emulator and interference emulator requires either a RF isolation chamber or direct connection to the RF impairment emulator equipment through a coaxial cable. RF circulators may also be required to isolate the transmit and receive signal for connection to the

emulator. The SUTs performance should first be calibrated by testing it in a impairment free environment. Once this benchmark is established, progressive physical layer impairment testing can be performed by introducing various channel models such as JTC, or an exponentially decaying Rayleigh model, for example. One advantage of the JTC channel model is that it was agreed upon by the Joint Standards Committee as an acceptable model for wireless indoor communications and is therefore appropriate as a way to compare various WLAN systems' performance.

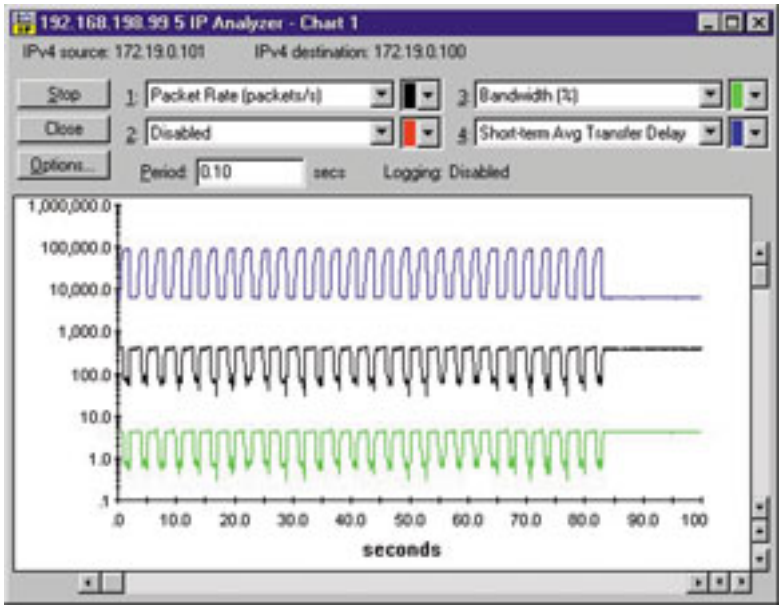

[JJ1] The graph illustrates: average transfer delay, packet rate and bandwidth percentage test results with an alternating 90 degree phase shift impairment

Performance when interference is present in the 2.4 GHz RF range

RF interference involves the presence of unwanted, disruptive RF signals that disturb normal system operations. An interfering RF signal of sufficient amplitude and within the receiver's frequency spectrum can appear as a bogus 802.11 station transmitting a packet. This can cause legitimate 802.11 stations to wait for indefinite periods of time until the interfering signal goes away.

This is of particular concern in the crowded 2.4 GHz range where microwave ovens, radiophones and Bluetooth devices and other non-802.11 devices share the spectrum with 802.11b devices. Noise and interference emulators can simulate these types of impairments in a lab and should be include as a component of any fully equipped WLAN test lab. Standardized WLAN interference models are being developed for the WLAN market but non have been released at the time this article was being written.

Measuring roaming hand-off delays and RF rate adaptation performance

When a WLAN user is moving about the AP will try to maintain the connection as the user moves out of range by reducing the speed of the connection first. Then, if the

signal continues to degrade, the client has the option of looking for another AP with a stronger signal. Intuitively, one might think that roaming hand-off delay would not be a significant issue for WLAN users since they are not typically on-line while walking around with their laptops. But there are a number of scenarios that must be considered where poor performance of the rate adaption function and roaming hand-off to another AP could impact the effectiveness of the wireless application.

Roaming can create a scalability issue for APs in certain applications. Consider how roaming typically happens in a campus or large enterprise environment. In both of these cases the AP may be required to deal with huge spikes in roaming activity.

In the IEEE 802.11 specification, there are three classes of client mobility:

&#149 1. No-transition: movement may or may not occur and the client never changes its AP association.

&#149 2. BSS-transition: A wireless client moves between multiple AP’s that are a part of the same wireless infrastructure.

&#149 3. ESS-transition: A wireless client moves between multiple AP’s that are not a part of the same wireless infrastructure.

When the wireless client changes its AP association, any traffic flows destined for the wireless client must change their physical path. Such an exchange is unlikely to be instantaneous and is likely to trigger various errors, including lost, misinserted, duplicate, and out-of-sequence packets as well as increased latency.

A wireless client that is moving out of the range of an AP can constantly rate adapt between a lower speed that is passing data and a high speed rate that is not. Additionally, if it is in range of two APs it may constantly switch its association between the two APs. Depending on the performance impact of re-association and the rate of association changes, this ping-pong scenario can have a dramatic impact on performance.

The 802.11 specification also states that service interruptions are likely during ESS-transitions, but implies that BSS-transitions can be completed without interruption. However, the documentation provides no procedure for implementing either the client re-association or AP hand-off functions. Since there is no standardized method for performing this function, each vendor implementation provides a unique testing opportunity.

Scalability of various security solutions; WEP, 802.11i, 802.1x, IPSec, SSL, Firewalls

Every standard has its “gotchas” - X.25 was too slow, IP didn’t have enough addresses, and 802.11’s “gotcha” is its security definition, Wired Equivalency Protocol (WEP). WEP fails to be specific about how Initialization Vectors (IVs) for the RC4 stream cipher are to be implemented. Some implementations reset IVs to zero every time they are restarted, and then increment them by one for every use. This

results in the likelihood that key-streams will be reused. Combine this with the fact that the encryption key and IV were not hash coded using MD5 or SHA-1 and WEP is open to simple crypto-analytic attacks against the cipher and decryption of message traffic.

The IEEE standards group is working on correcting the short comings of WEP with its 802.11i working group, but equipment vendors and service providers have already turned to other security solutions like 802.1x, IPSec, SSL, integral firewalls and other proprietary solutions. This creates a playing field in which WLAN performance can vary significantly from solution to solution depending upon the security mechanism chosen and the efficiency of its implementation. To understand the performance issues associated with a WLAN security solution, realistic testing of the control plane with IP Virtual Private Network (IP-VPN) protocols, as well as the data plane with connection-oriented (HTTP) and connectionless (UDP) traffic is key.

It is important to determine IP-VPN tunnel creation capacity using IPSec protocols and also generate UDP or HTTP traffic over each tunnel and measure data performance characteristics such as packet loss, latency, and response time. Results should be displayed in an efficient layout for quick performance benchmark comparisons and should function as an aid to troubleshoot interoperability issues.

Firewall testing should determine the maximum application transaction capacity and measure application throughput with TCP acting as the transport agent. It should gauge the performance of firewalls executing Network Address Translation (NAT) as well as the impact of multiple packet filtering rules being set up. Firewall test equipment should also evaluate a firewall's ability to deal with Denial of Service (DoS) attacks by allowing different ratios of attack traffic to be generated in parallel with other application traffic. This permits bench marking of legal traffic performance while attack traffic is being directed at the device under test–consuming its resources.

Qualifying firewall performance requires exposure to a high volume of Web transactions sufficient to model peak Internet conditions and the ability to simulate an Internet mix of traffic as stateful firewalls require real application traffic generation in order to forward traffic properly. Scalability is a key component in choosing the right test equipment for firewall testing as thousands of HTTP transactions per second and millions of concurrent TCP connections may be needed to effectively test a firewall’s performance.

Testing will accelerate deployment of WLAN technology

Because WLAN technology brings real benefits in cost savings, mobility and increased productivity to the work place, campus, home and hot spots in between, it is being accepted and deployed at an increasingly rapid rate. Triple digit growth rates are projected out through 2005, for 802.11a, 802.11a/b dual mode and 802.11g devices. With rapid acceptance and deployment, though, comes the specter of unrealized performance and less than acceptable quality of service (QoS). The early buyer will tolerate some shortcomings in these areas just to get the

latest technology on the block, but as the solution moves into the mainstream the buying public will become less tolerant of performance and QoS issues. Users aren't happy to be able to just roam, for example, they want to move about freely without interruptions in service as they dis-associate and re-associate with APs. Realizing this, chip and equipment manufacturers and service providers are beefing up their test labs with test equipment that will allow them to emulate real world RF network conditions and to test the robustness and scalability of their products and network solutions. WLAN technology has enormous revenue generating potential for those who can deliver a solution that meets the performance demands of the mainstream buyer.

Source URL (retrieved on *07/24/2014 - 10:19am*):
http://www.wirelessdesignmag.com/product-releases/2002/12/wlan-performance-analysis