Azimuth

September, 2005

Colliding Views on Call Capacity Measurement

By Fanny Mlinarsky and Jose Graziani

With the emergence of voice over WiFi and WiFi-enabled cell phones, call capacity of the WiFi infrastructure has come under scrutiny. What is the call capacity of WiFi? Can the WiFi infrastructure handle this capacity? How do we test it?

In pondering these questions some industry experts instinctively focus on WiFi channel capacity. The trouble is – channel capacity and infrastructure capacity are two different things. So the first question is…

What are we measuring?

If we are measuring the performance of an AP, we need to bombard this AP with the maximum theoretical number of calls and see how it handles such a load. If the AP crashes or otherwise misbehaves in the presence of the maximum call load, then it is not ready for deployment. With millions of frames traversing an average network each day, occasionally the traffic load will peak at maximum theoretical capacity. Thus, APs must be tested at peak traffic loads and this means: without collisions.

Collisions reduce the test traffic load

During collisions, the colliding stations back off and wait for some random period of time before re-transmitting their frames. This slows down the true throughput of the channel and results in some statistically average traffic load, which is insufficient for a proper test of AP capacity. The only way to create the maximum test load on the AP is by eliminating collisions.

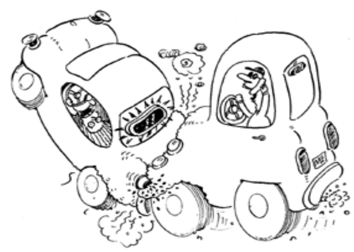

Collisions reduce network throughput, so test traffic with collisions cannot reach peak channel capacity.

The controversy

Some industry experts insist that a capacity test with no collisions is invalid because it does not measure 'realistic' channel capacity. This is true if *channel* capacity is what we want to measure. But AP capacity and channel capacity are two different things. For the network to be robust, the AP capacity must exceed the 'realistic' average channel capacity.

Test to ensure network robustness

Peak call capacity of the WiFi channel only occurs during periods with no collisions. To verify whether APs can sustain peak channel traffic we must test them with peak load.

Collision-free maximum test load can be emulated using a multi-client emulator such as the Azimuth testMAC module running the Call Capacity test, which measures the ability of APs, switches and other infrastructure devices to forward and prioritize voice traffic in the presence of background data traffic.

This test measures call quality as a function of call and background traffic load. It also verifies the ability of APs and other infrastructure devices to schedule calls so as avoid overload conditions that degrade voice quality.

and

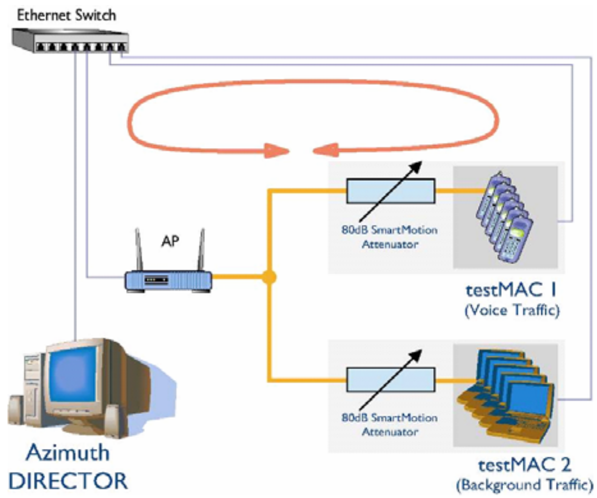

Figure 1: Call Capacity test uses Azimuth's testMAC traffic emulation modules to generate analyze high priority call traffic and low priority background traffic from multiple clients. The testMACs use the WMM protocol to prioritize voice traffic over background data traffic. While sending call traffic, the testMAC measures frame loss, delay and jitter on the voice streams going through the infrastructure under test. Admission Control is used to test the infrastructure's ability to limit the number of active calls when capacity is reached.

The testMAC module emulates multiple voice or data softClients using a single PHY interface. The transmissions from different softClients are interlaced on the medium to attain maximum theoretical throughput. Thus, in a controlled way, the testMAC can increase the number of active clients and the density of emulated traffic for a thorough test of the infrastructure.

Interlaced packets from multiple softClients

Figure 2: The traffic emulated by the testMAC looks like voice streams or data from multiple clients. The frames from different softClients are interlaced on the medium to attain maximum theoretical throughput.

The frames from different softClients are separated by inter-frame spacing (IFS) consisting of the DIFS period plus backoff (figure 3). The backoff is a random time period within a backoff "window", with a value between zero and CWmin[1]. The reason for randomizing backoff periods is to minimize the chances of different stations accessing the medium simultaneously and colliding.

The IFS determines the maximum test load offered to the AP under test. The smaller the IFS, the higher the frame rate, the more challenging the test. The greatest theoretical frame rate is reached with

Effect of collisions on test load

When two or more stations collide, each station involved in the collision has to back off with a backoff window that's double the size of the previous window. For example, for the first collision backoff window is increased from 15 slots to 31. With each consecutive collision the backoff period is doubled until it reaches CWmax (1023 slots for 802.11a). During periods of heavy load, multiple successive collisions are not uncommon and this significantly decreases the effective load. The only way to reach a theoretical maximum load is by eliminating collisions.

[1] CWmin is a time interval defined in the 802.11 standard for each PHY (802.11a,b,g).

minimum size frames separated by minimum size IFS. Some APs may not be able to keep up with short voice frames arriving back-to-back from different softClients. Therefore, a good test will include such a test load.

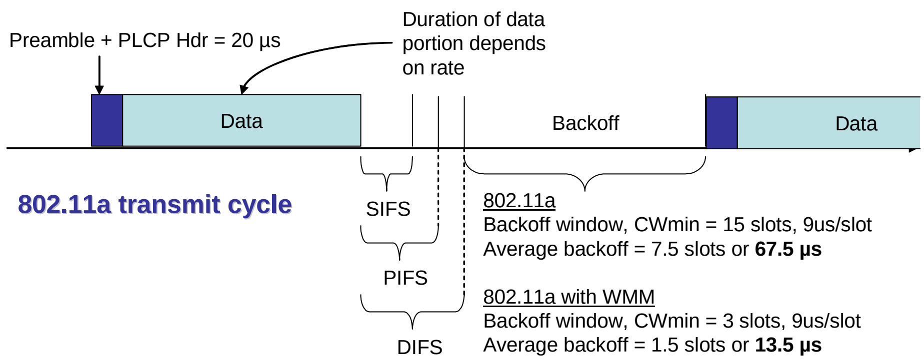

Figure 3: 802.11 transmit cycle. The inter-frame gap (IFS) for normal voice and data traffic is DIFS plus the backoff period. Theoretical minimum gap between frames is equal to the DIFS period.[2]

In addition to supporting average 802.11 and WMM backoffs, the testMAC has the flexibility to eliminate the backoffs entirely to produce the theoretical peak load on the network.

When measuring AP call capacity, you can configure the test traffic to use any of the back-offs listed in Table 1. You can also select a codec, which determines the frame size and frame rate of the voice stream traffic. Table 1 shows an example of the popular G.711 codec that uses 236-byte frames. G.711 does not represent the worst case call load – this distinction that belongs to the SRP codec. SRP uses 128-byte frames and can create aggregate worst case traffic at the AP of 10,000 fps[3] on 802.11a and 802.11g networks. Can your AP handle this?

Table 1: Test loads supported by the testMAC – example for voice codec G.711-20ms (frame size: 236 bytes, voice stream: 50 frames/sec, bidirectional). Peak frame rate[4] in fps (frames per sec) and call capacity are shown for each IFS.

	Inter Frame Spacing				
	standard backoff	WMM voice backoff	DIFS, no backoff	PIFS, no backoff	SIFS, no backoff
802.11a (fps, calls)	4,963 fps 50 calls	6,780 68	7,463 75	8,000 80	8,621 86
802.11g (fps, calls)	4,963 50	6,780 68	7,463 75	8,000 80	8,621 86
802.11b (fps, calls)	1,342 13	2,151 22	2,299 23	2,410 24	2,532 25

[2]DIFS = distributed (coordination function) interframe space; PIFS = point (coordination function) interframe space; SIFS = short interframe space
[3] 10,000 frames per second rate can happen with SIFS intervals, which are used in burst mode
[4] Peak frame rate shows a theoretical maximum frame rate that an AP has to forward. The SIFS column shows the frame rate when bursting (TXOP) is used – a load that can legitimately occur in 802.11 networks.

The test results produced by the Azimuth Call Capacity test include measurements of infrastructure frame loss, delay, jitter and the computed voice quality metrics, R-Factor and MOS.[5]

As the number of calls is incremented, the measurements of voice quality metrics are preformed at each step. A plot of voice quality metrics as a function of call load (figure 4) demonstrates at what point the AP is unable to handle the load and to shows whether an AP can support peak traffic. This test can be performed with background traffic from multiple softClients to verify the AP's ability to prioritize voice traffic over data.

Figure 4: The Call Capacity test increments the number of calls going through the infrastructure in a controlled way. It measures and plots delay, jitter and frame loss on each active call and then computes the R-Factor and MOS voice quality metrics for each call.

Summary

This paper has presented the difference between probabilistic capacity of contention-based networks such as 802.11 and maximum capacity of infrastructure components comprising such a network. Probabilistic capacity can be measured using contention-based traffic or it can be theoretically derived. Maximum capacity of network components must be measured using controlled contention-free multi-client traffic load.

---

[5] For more on Voice over WiFi requirements see Wireless Systems Design article, "Metrics And Methods Bring VoWLAN Success" (http://www.wsdmag.com/articles/print.cfm?articleID=10003)