

# *Advances in Wireless Infrastructure Control:* Optimizing Client Behavior for Improved Network Performance

*A Farpoint Group Technical Note*

Document FPG 2008-341.1  
September 2008




While it's very clear that the core objective for wireless overall is to duplicate the capabilities and performance of wire, such also serves to highlight the essential differences between these two as interconnect and access technologies. Wire and cable (and especially fiber) represent a nice, neat, contained electromagnetic domain that, when properly deployed, is (a) resistant if not immune to externalities like RF interference and electromagnetic-wave propagation challenges, and (b) has a high likelihood that a signal placed on one end will appear at the other. Radio, on the other hand, is fundamentally statistical in nature, and is consequently much more challenging - signal fading, interference, limitations on transmit power, ever-greater application requirements for higher throughput and network capacity, noise (as measured by the signal-to-noise ratio (SNR) at any given moment in time), and client motion all affect how well wireless can in fact duplicate, if not at least partially replace, wire, especially in high-demand environments.

Or do they? We've previously written on the benefits of RF Spectrum Management (*RFSM*; see Farpoint Group White Paper 2003-201.1, *Beyond the Site Survey: RF Spectrum Management for Wireless LANs*) techniques used to improve overall capacity and to support increasingly-dense client-device populations as well. These include, at a minimum, the ability to automatically set and dynamically adjust infrastructure transmit power levels and channel assignments. But dense AP deployments (see Farpoint Group White Papers 2004-193.1, *Rethinking the Access Point: Dense Deployments for Wireless LANs*, and 2005-083.1, *Wireless LAN Dense Deployments: Practical Considerations*), a key strategy (if not a requirement) to increase capacity, can clearly complicate *RFSM*. This creates the requirement for more sophisticated solutions that, while more complex to implement, can (as we recently learned) yield significant improvements in the performance of high-density, mixed-mode (802.11n and legacy) deployments.

While it can be argued that 802.11n corrects some of the problems inherent in earlier 802.11 standards in that it provides much greater throughput, reliability, and improved spectral efficiency (perhaps simply as a brute-force argument reminiscent of the transitions to fast and gigabit Ethernet), operating 802.11n in mixed-mode environments, which will be all but inevitable until the installed base completely rolls over, can result (despite or perhaps because of co-existence mechanisms within the standard itself) in poor performance for everyone. Add in increasing requirements for raw throughput, data volume, and time-bounded traffic (most notably voice and streaming video), and, clearly, more sophisticated RF Spectrum Management techniques are required. We will refer to these new developments under the general heading of *Wireless Infrastructure Control (WIC)*.

We recently had the opportunity to gain some exposure to an advanced *WIC* implementation in the lecture hall of the Architecture Building on the campus of the University of Washington in Seattle (see Figure 1). Specifically, Aruba Networks invited us to observe testing of their Adaptive Radio Management (*ARM*) 2.0 (also called *Adaptive Infrastructure Control* by Aruba), as embodied in their 3.3.2.5 software release, in this case for an Aruba 3600 Multi-Service Mobility Controller connected to four Aruba AP125 802.11n Access Points. Presenting the load in the testing performed were a mixture of up to 101 (depending upon the specific test) client notebooks from a variety of vendors (including Acer, Dell, Gateway, HP, IBM, Lenovo, Toshiba, all running XP or Vista, and even a few Macs) with a variety of radios



**Figure 1** – The lecture hall on the University of Washington campus. Two APs were placed on each side of the hall, with clients on the desks in alternating rows. *Source:* Farpoint Group.

(including Atheros, Broadcom, and Intel), as would typically be found in a university or, indeed, almost any corporate setting. Note that we used a combination of .11b/g/a/n clients, and that each radio implementation has its own unique behavior due to radio chip-set vendor design decisions and the client-centric nature of 802.11 itself. We observed both quantitative (using the IxChariot 6.5 performance-evaluation tool) and qualitative (using actual applications found in educational environments) tests, and the results are described below.

## Wireless Infrastructure Control – The Basics

All successful wireless-LAN installations take advantage of parallelism in

the form of multiple simultaneous radio channels. The key to WIC, as we saw with Aruba’s ARM 2.0, is in controlling multiple APs and channels simultaneously – with, in this case, surprisingly good results.

As we noted above, basic RFSM includes channel and power settings and adjustments. Aruba’s ARM 2.0 extends this feature set with four additional functions, as follows:

- *Traffic Shaping and Airtime Fairness* – In traditional WLAN installations, all clients are allocated access to the airwaves under the same set of conditions. Aruba noted (from simple queuing and scheduling theory) that it makes sense in some cases to allocate more airtime to devices with greater capabilities (particularly those based on 802.11n) and/or traffic requirements. The result in this case is lower throughput for 802.11b clients, but much greater throughput for .11g and .11n, and much improved aggregate throughput given a particular workload. Aruba uses an airtime allocation ratio for b:g:n that assures equal access to the air on a per-client basis (which they call *Fair Access*), or

which can be weighted towards favoring higher speed clients (which they call *Preferred Access*).

Traffic shaping provides not only throughput optimization given fast and slow clients, but also assures that client starvation does not take place when a high-density user population is supported on a single AP, as we observed in the testing reviewed below.

- *Band-Steering* - Given dual-band infrastructure and that a dual-band WLAN client usually has an equal likelihood of associating with either the 2.4 or 5 GHz. band, a less-than optimal allocation or balancing of clients is the norm in dense deployments today. This results in inappropriate channel allocations, greater contention, lower signal-to-noise ratios, and overall suboptimal performance for any type of traffic. Aruba's band-steering capability steers dual-band clients to the 5 GHz. band; this functionality, like all of ARM 2.0, requires no client-side configuration or modification and is completely transparent to the user. Farpoint Group always recommends the use of spectrum above 5 GHz. whenever possible, and that buyers always obtain WLAN infrastructure and client adapters that can operate on either band.
- *Co-channel interference management* - This interesting feature allows multiple APs operating in close proximity on the same channel to coexist with very little degradation in aggregate throughput. Co-channel interference is often responsible for severe performance degradation of 50% or more.
- *Spectral Load Balancing* - Analogous to Band-Steering, this function balances traffic load across Wi-Fi channels within a given radio band. Rather than attempting to establish simply a fair distribution of clients across APs in a given location irrespective of channel (which can be ineffective in dense deployments), clients can be moved to a different channel based on a per-channel load balancing algorithm that takes into account the AP density, client density, traffic load, SNR (as reported by individual Wi-Fi chips) and channel conditions as reported for each channel.

Our interest, of course, was in seeing if these capabilities would yield improvements in real-world performance, particularly in high-demand, dense-user environments. As one might assume in an education environment, the University of Washington site proved perfect for this. We should note, however, that school was not in session at the time (the reason we got to use the lecture hall in the first place) and that, while there were other WLANs operating in the building, the general lack of traffic and the nature of the very solid structure we were in minimized external sources of interference to a great degree. We were thus looking at what ARM 2.0 could do in a fairly controlled setting, the question being what meaningful benefits could be realized from its application in high-demand applications. But note that we did not attempt to document the actual RF environment or run rigorous benchmark comparisons – we were simply looking to see if these features worked in what everyone agreed was a dense and thus challenging environment.

## Evaluating Next-Generation Radio Management: Testing and Results

Results for each of the tests we observed were as follows:

- **Test 1a - Traffic Shaping for Improved Mixed-Mode Client Operation**

In this test, we used three clients, one each on 802.11b, g, and n, and all on the same 2.4 GHz channel on a single AP. We ran a thirty-second, four-stream IxChariot script simulating a TCP download using both default settings, and again with Preferred Access enabled. The results are shown in Figure 2.

By more carefully scheduling traffic, a net combined increase in throughput of 18% was realized, with both .11g and .11n benefiting from Preferred Access. .11b throughput, while more than cut in half, was very low to begin with, and it is unlikely that a .11b user would have noticed more than a minimal perceptual change in performance - while .11g and .11n users would likely be very pleased indeed.




Figure 2 - Results of Test 1a, in Mbps. Source: Farpoint Group


- **Test 1b – Traffic Shaping with b/g Clients Only**

This test was identical to Test 1a, but used only two clients, one each on .11b and .11g, again on the same channel with a single AP. In this case (see Figure 3) we saw a 42% jump in combined throughput, and .11g throughput improved by a whopping 136%.


- **Test 1c – Traffic Shaping with a/n Clients Only**

This test was identical to test 1b, but the two clients were on a single 5 GHz. channel. The

results here (see Figure 4) were similar, showing a 51% decrease in .11a throughput , but an 18% improvement for .11n, with a combined 7% improvement.



**Figure 3** - Results of Test 1b, in Mbps. *Source:* Farpoint Group



**Figure 4** - Results of Test 1c, in Mbps. *Source:* Farpoint Group

- **Test 2 – The Effects of Band-Steering**

It stands to reason that balancing clients across physical channels, if done properly, will result in improved performance. This should be particularly valuable when moving clients and

thus their traffic to the 5 GHz. bands, which remain relatively open in most settings. Traditional concerns about limited range at 5 GHz. are no longer relevant, at least with respect to our test within a single room, and regardless when operating over very limited range, but also due to the many benefits of 802.11n with respect to rate-vs.-range performance improvement.

This test involved four clients, one each on 802.11b, g, a, and n. Only the .11n client was allowed to float between bands. The results (see Figure 5) were quite dramatic – moving (again, automatically, and with no client modifications or re-configuration of any form) the .11n client from 2.4 to 5 GHz resulted in a 216% improvement in .11n throughput, and a corresponding improvement in .11b and .11g of 108% and 162%, respectively. The .11n improvement did come at the expense of .11a performance, however, but this is to be expected. But note that aggregate throughput increased in this case by 70%.




Figure 5 - Results of Test 2, in Mbps. Source: Farpoint Group

- **Test 3 – Fairness with Traffic Shaping**

This test involved 25 clients across a mix of client OSES, radio chipsets, and drivers, all associated with a single .11g AP. The objective was to evaluate fairness, as implemented with the Preferred Access option. We found that aggregate throughput in the default case was 20.498 Mbps, and 20.07 Mbps with Preferred enabled. But these numbers do not tell the real story: the standard deviation among all results (in Mbps) in the default case was .6 and the difference between maximum and minimum Mbps was 2.7. With Preferred Access enabled, the standard deviation fell to .4, indicating smoother traffic flow, and the max-min difference fell to 1.8, showing greater uniformity. While we do not expect that clients with fundamentally bursty traffic flows would benefit much from this feature, those with more uniform or homogeneous flows across all clients would at least see no clients starved.



Aruba later noted that the algorithm used here actually learns over time, and thus would produce more refined results over time. We expect that high-density environments, such as the one in which we were testing, would benefit the most from this capability.

- **Test 4 - Co-channel Interference Mitigation**

This test involved the same configuration of units as Test 3, but this time with two APs on the same channel placed roughly 20 feet apart, with wireless clients connecting to both of the APs. As we noted above, this situation would normally be expected to result in severe contention and interference, and indeed, Aruba reported that this configuration using a previous software release without Fairness and Traffic Shaping yielded only 11.2 Mbps – again, about as expected. But this test instead yielded 18.9 Mbps aggregate throughput, showing that the controller could indeed tame traffic flows in a dense deployment, in this case to the benefit of net client throughput and thus all users.

We then ran four subjective tests using applications that would typically be found in an educational environment like a university. These tests involved the use of up to 101 clients simultaneously, with exceptions as noted below. The evaluation of performance was typically made by standing in the back of the lecture hall and observing the amount of time required for all screens on all clients to change in response to direction from a central console, which would typically be under the control of an instructor or lecturer at the front of the hall. We were looking for uniform changes in screen context, indicating uniform performance and throughput across all clients, and noted any laggards.

- **Test 5 – Video Furnace**

Video Furnace [<http://www.videofurnace.com/>] produces integrated IP video distribution systems used extensively in education and government applications. The system is designed to provide secure, high-performance multicast video. Our application involved three different video streams consuming approximately 2.5 Mbps each, accessing a single dual-radio AP. Any of the 101 PCs could tune to any of these three streams, and we did switch streams on a few PCs during the test. But evaluation consisted primarily of walking around and observing video quality, which was overall excellent. Errors observed (video artifacts and occasionally brief freezes) were ultimately traced to source material. The presence of legacy clients did not disrupt delay-sensitive multicast video transmission.

- **Test 6 – DyKnow**

DyKnow Vision Software [<http://www.dyknow.com/>] is an application used extensively in educational environments, and features the display of lecture material (with annotations as entered in real-time by the instructor), along with collaborative note-taking, live chat, and student polling/response tools. Our testing here involved 95 PCs (DyKnow does not run on Macs), 75 on .11a and 20 on .11g, again with a single dual-radio AP. Standing in the back of the room, we observed a delay of less than two seconds on all client PCs as the instructor changed slides. There were no obvious performance issues and we think students and teachers alike would have been pleased with this performance.



- **Test 7 – Classroom Presenter**

Classroom Presenter [<http://classroompresenter.cs.washington.edu/>] is a tool developed at the University of Washington, and similar to DyKnow Vision, that allows the sharing of visuals in classroom settings. This test was the ultimate exercise in our time at the University of Washington, involving all of the PC clients (again, in this case, not Macs) and all four AP125s. Spectral Load Balancing was enabled, and system logs revealed that a few clients were in fact moved between channels during this test, resulting in greater uniformity of the distribution. Spectrum Load Balancing and Band-Steering were enabled here as well, and system logs revealed that clients were in fact moved across channels and bands during the test run. In this case we noted essentially instantaneous screen changes; latency was less than one second in every case.

- **Test 8 - ExamSoft**

Finally, we used the ExamSoft Flex-Site [<http://www.examssoft.com/main/index.php>] application to simulate the administration of a classroom student-testing exercise. We had assistants moving through the room submitting exam results to the system from each client PC. While we could not synchronize more than a few submissions at once, we noted no latency or other issues sending roughly 250 KB of data from each PC.

## **Conclusions: Real Progress**

While this was, again, not a formal benchmarking exercise but rather a general test of advanced wireless infrastructure control capabilities, we were impressed at how far we have come in the ability of wireless LANs (properly equipped, of course) to handle large numbers of users with time-critical traffic. We concluded that, while it is indeed very difficult to manage Wi-Fi at the RF level, real benefits can be gained from Wireless Infrastructure Control based on the automated adjustment of key parameters in response to real-world radio and traffic conditions. And this testing was especially valuable in that high-density environments have always been challenging because the close placement of APs and densely-located clients represent a worst-case scenario.

It is quite clear that Aruba's ARM 2.0 has real benefits in educational settings and applications. But we expect that the benefits noted will carry over to any densely-deployed (infrastructure and users) environment, which we believe will clearly become the norm as users migrate from wired connections to wireless (especially 802.11n) installations. And we also believe that advanced wireless LAN architectures – and the applications that depend upon them - will continue to benefit from further advances in this area of technology.



Ashland MA 01721  
508-881-6467  
[www.farpointgroup.com](http://www.farpointgroup.com)  
[info@farpointgroup.com](mailto:info@farpointgroup.com)

The information and analysis contained in this document are based upon publicly-available information sources and are believed to be correct as of the date of publication. Farpoint Group assumes no liability for any inaccuracies which may be present herein. Revisions to this document may be issued, without notice, from time to time.

**Copyright 2008 — All rights reserved**

Permission to reproduce and distribute this document is granted provided this copyright notice is included and no modifications are made to the original.