Competitive Product Review

tom's hardware

Beamforming: The Best WiFi You've Never Seen

Forget 802.11n Draft 2.0. The future of video-capable WiFi depends on a signal-boosting technique called beamforming. We put the pioneers in this frontier through some real-world testing to find out which technology is going to change the wireless world.

more >>

Open-Mouthed Amazement

You should have seen my wife's face when she found me glued to the Victoria's Secret Fashion Show. "No, honey, come here!" I said, my face aglow with the bikini-clad pixels of Tyra and Heidi Klum. "You've got to see this!"

Arms crossed. Pursed lips. "Mm-hmm. Yes?"

I pointed at the laptop on the counter in front of me. "Not the models. The video. It's high-def with a 19.2 megabits per second stream rate. Looks perfect, like HDTV, right?"

"Mm-hmm."

"Now turn around." I pointed at the plasma screen on the wall pulling a different part of the same video, a second stream at 18.4 Mbps, through our Xbox 360 with an attached 802.11n bridge. "That's almost 40 megabits streaming over WiFi. I've never even been able to do one stream before, and now we've got two!"

My wife looked at the screens, looked back to me, and shrugged. "OK, then. I'll leave you and the girls to it. Have fun."

She walked away and slammed the front door. I don't think she actually cared if I was having fun. Strange. Clearly, she didn't understand that something amazing had fallen into my lap. Actually, let me rephrase that. Something incredible had happened to my network. With an access point clear across the house, transmitting through one floor and three or four walls, coping with literally a dozen interfering WiFi networks surrounding the house, I was getting wireless network performance unlike anything I'd ever seen before.

This was my first experience with beamforming, something I'd only seen vague mention of on long-term WiMAX roadmaps. But here it was in an 802.11n access point from a company I'd never heard of, and it blew away everything I'd ever seen a wireless product do before.

Interested? Then let's dig in. I may not have runway models to offer you, but I still think you'll be impressed.

Beamforming Basics

Think of radio transmitters as little stones dropped in a pool. You know from high school physics that a dropped object will send out waves across the water's surface. If you drop two stones, those waves will overlap with each other in a regular "interference" pattern. Changing the characteristics of a stone will change the amplitude and phase of the waves it emits, as well as the characteristics of the interference pattern generated with waves from other stones.

If you have enough control over the situation, you can have a sensor at the edge of the pool looking for just the right wave pattern, and you can keep changing the stone characteristics until that exact pattern arrives at that particular point. Elsewhere in the pool, the wave pattern will be different, and that's fine. You're only looking for that one pattern in that one place. Everything else can be ignored.

In a nutshell, this is the essence of beamforming. You're controlling the output characteristics of each transmitter within a transmitter array so that the overall signal is optimized to reach a given receiver in a given direction. With an antenna array in which each antenna is transmitting with slightly different characteristics, you have what's called a phased array. As we'll see, there are two primary forms of phased array used in wireless access points: on-chip and on-antenna, adopted by Cisco and Ruckus Wireless, respectively.

Inside On-Chip Phased Arrays

Let's get a little more specific. You may be familiar with MIMO (multiple-input, multiple-output) technology, first adopted with some 802.11g products and now incorporated into the 802.11n specification. Go back to our pool example. When you drop a stone in the left edge of the pool and the receiver is on the pool's right edge, some of the waves will travel in a straight path from left to right and arrive at the transmitter via the shortest route possible. However, some waves will bounce off the top, then arrive at the receiver a bit later than the straight-path waves. Some will bounce off the bottom, then up to the top, and then arrive at the receiver. All of these variations emerge from a single stone drop, or radio burst. To a simple receiver, this sounds a bit like confusing, overlapping echoes. This "multipath" effect has traditionally plagued the performance of radio communications.

But what if you used multiple antennas at each end of the pool, applying enough analysis intelligence on each side to turn those echo paths into conduits for different data streams? With multiple antennas on each end, you can send different data streams from different antennas and receive them at the other end in the same manner.

To drag in a different metaphor, think of a highway. If the highway is only one lane, you can have one big truck going full-speed to its destination. However, if you sub-divide that one big lane into three or four narrower lanes, you can have three or four compact cars going to the same destination at the same speed. They just happen to be going there along slightly different paths. When you took good ol' 54 Mbps 802.11g and its 20 MHz channel highway, divided that highway into multiple sub-channels, and increased the number of antennas, you got 802.11g MIMO.

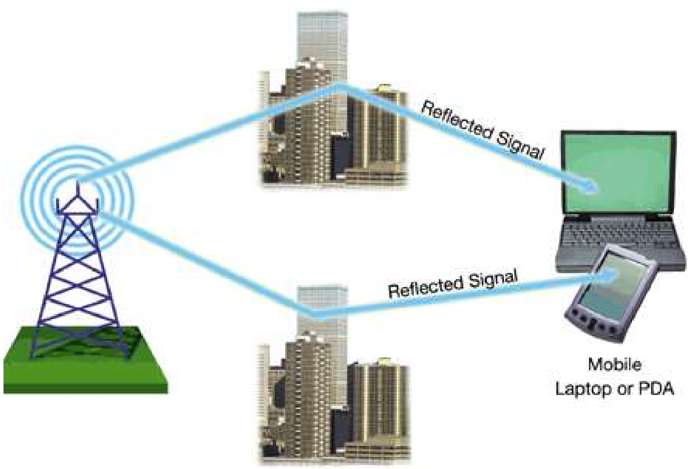

Specifically, 802.11n typically transmits three data streams and receives two, commonly referred to as a 3x2 antenna array. There are some 3x3 schemes in the works, such as the so-called 450 Mbps WiFi set out by Intel with the launch of Centrino 2, but no access points have arrived yet to support this. Like 802.11g before it, 802.11n can use channel bonding, turning two 20 MHz streams into a 40 MHz pipe. To be totally accurate, you should actually see antenna arrays noted with three numbers: the number of transmit antennas, the number of receive antennas, and the number of spatial streams (or data streams) to use our subdivided highway analogy. So a 3x3:2 (also noted as 3x3x2) array would have three transmit antennas, three receive antennas, and two spatial streams.

I mentioned earlier that on-chip beamforming was one of the two beamforming methods applicable in WiFi. This works by not only boosting total power gain by having multiple antennas in play, but also phasing the antenna signals so that a higher signal "beam" is cast in the receiver's direction while less energy can be expended in other directions. With two transmit antennas, you can expend less total energy while quadrupling the transmit signal sent in the beam's direction. The transmitter/access point only needs to receive a single packet from the client to get a lock on the signal path. Analysis of multiple packets can determine which of the multipath options is optimal at any given time.

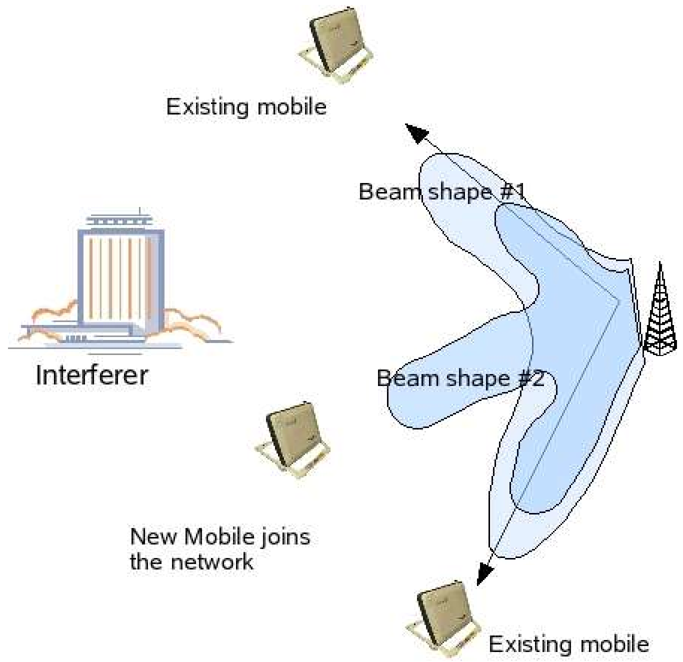

The incredible thing is that chip-based beamforming, like MIMO, has been compatible with 802.11a/b/g for years. In fact, the technology is an optional part of the 802.11n standard. Despite its benefits, though, Cisco is the first to deliver on-chip beamforming to market. The enterprise-oriented AIR-LAP1142N access point is Cisco's first and so far only product to feature beamforming, which it brands as ClientLink. It arrived in the first quarter of 2009, but the firmware that enables beamforming capability didn't arrive until July. We tested with this firmware literally within days of its release.

The Client That Could Be

Ever since the days when 802.11a/g grew a second antenna, we've had "transmit/receive diversity," which sends the same data stream out over multiple antennas and simply lets the access point select whichever antenna is receiving the best signal. Applied to 802.11n, transmit diversity used multiple antennas to help increase range and better deal with difficult locations. This is why 11n does a generally better job than 11a/g at eliminating dead spots.

However, 802.11n equipment got another jump in intelligence with the addition of maximal ratio combining (MRC). This technology combines multiple antenna signals in such a way that strong signals are multiplied while weak signals are attenuated. The signals you want get boosted, while those you don't have their power cut. MRC is built into all 802.11n chips.

Now, as you might expect, the receiving end can play an important role in optimizing chip-based beamforming. With 802.11a/g, access points could listen to the client and use rudimentary MRC analysis to boost power along the best-suited beam, providing a gain of roughly 1 to 2 dB. The catch here is that the access point was doing all the work. There was no active feedback coming from the 802.11a/g clients.

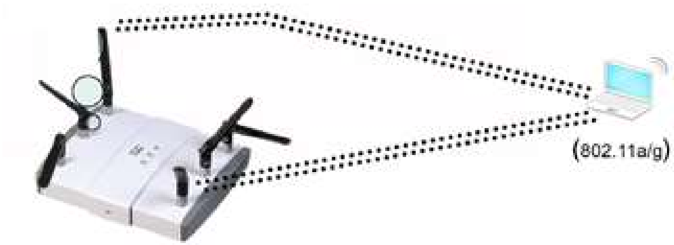

With "implicit beamforming," wherein an 802.11n AP is communicating with 802.11n clients, you can have some feedback. Rather than having the access point perform all of the signal analysis, it can query the client and see if it agrees that this or that particular beam orientation is optimal. Having this limited two-way communication yields a maximum of 3 dB additional gain, but the bad news is that there are currently no products on today's market supporting implicit beamforming.

With "explicit beamforming," feedback between the AP and client happens much more frequently. This way, if a client moves or an antenna gets adjusted or anything happens to alter the dynamics of the signal strengths, the system is able to adapt almost instantaneously to a new, optimized configuration. Again, having the client involved in this way can yield up to a 3 dB benefit with two radios, but there are no products available today offering this capability. Hopefully this will change.

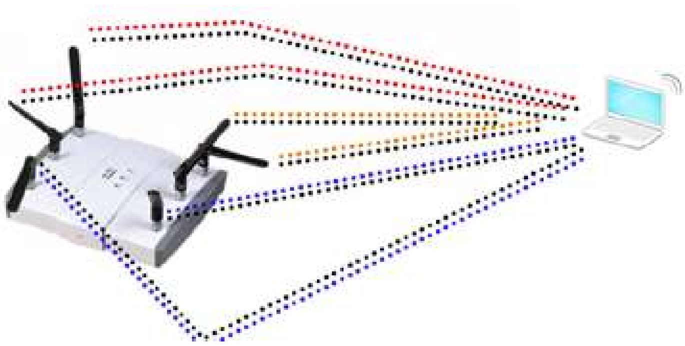

On-Chip Challenges

Now that you get the gist of how beamforming works, you're probably wondering why the technology never went mainstream. After all, in comparison, trying to optimize signal strength with several antenna jutting from a conventional 802.11n access point is a joke. These multiple antennas are, in a way, glorified rabbit ears. Even if you spend the time to fiddle with them and get what seems to be the best throughput in a given spot with the antennas set in just such a way, what happens if you have to move the access point or the end client? What if you add a second or third client? It's chaos. The fact is that proper signal optimization with current-gen products is futile.

Why hasn't intelligent beamforming, which has the ability to sense optimal phasing and orient beams for multiple clients, been widely adopted? It's a mystery—probably another one of those "we as an industry are still in the process of discussing various blah blah blah" things.

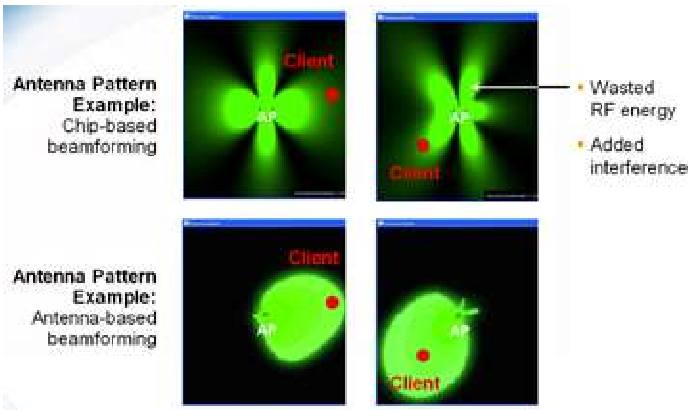

A skeptic might suggest that on-chip beamforming hasn't taken off because it's sounds better on paper than it is in real life. We know that, in theory, beamforming should save power. You only need to boost the signal in a certain direction and drop power for any signals that don't assist that beam. The problem here is that when you're dealing with omnidirectional antennas, there's only so much control you can have over your beams.

For an intriguing illustration, check out Falstad's Antenna Applet and be sure to choose Broadside Array from the top pull-down menu. You can increase antenna counts, play with the distances between them, and modify signal strengths. As you'll see, with two omnidirectional antennas you never get away from having a lot of beams, and therefore energy expended in unneeded directions (these unwanted beams are often called backlobes). Naturally, if you have beams going off in stray directions, these can cause co-channel interference and impede the signal you actually do want.

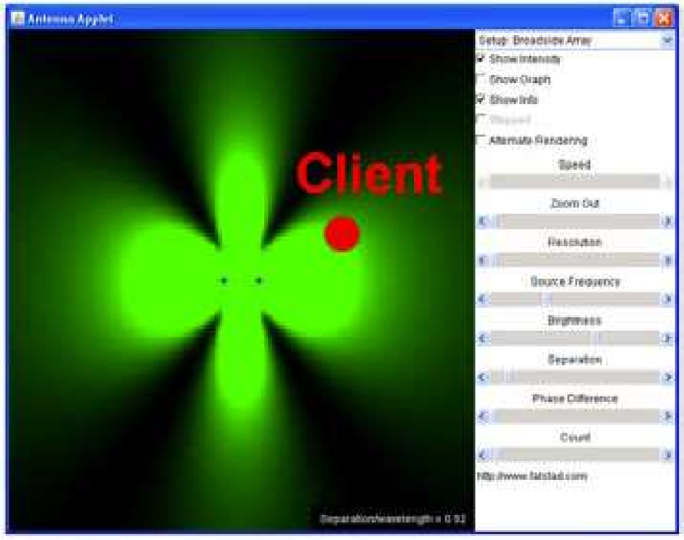

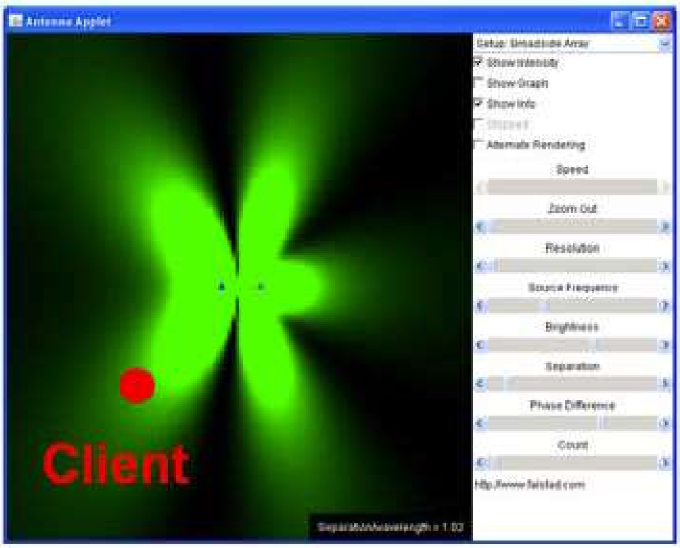

It seems likely that next-gen 802.11n will incorporate implicit and explicit beamforming at some point, as there are very few technical or cost barriers. However, which approach will vendors integrate? And we haven't even scratched the surface on options. For example, there are three sub-types of explicit beamforming. So if there's one culprit behind the lack of beamforming adoption, concerns over interoperability is probably to blame. For those of who find yourselves thinking, "Come on! I don't care about 100% interoperability. I just want crazy good wireless performance in my space," keep reading.

Ruckus and On-Antenna

Fortunately, an omnidirectional antenna is not the only way to obtain 360 degrees of wireless coverage. If you had enough directional antennas with overlapping cones of coverage, you could effectively cover 360 degrees. But the beauty of this sort of setup is that you don't have to run all of those antennas all of the time. Once you get a fix on the direction of the client, you simply have to determine which set of antennas (two or more) will create an optimal beam to the desired point.

Keep in mind that this isn't always a straight line. The client could be around a corner, and the better signal might be had by bouncing the stream off of a couple of walls rather than trying to punch straight through the obstructions. Or conditions may change. Doors can open and close. People move around. Someone might turn on a microwave oven and spew interference all over the spectrum. All of these things can alter signal paths and hamper traffic.

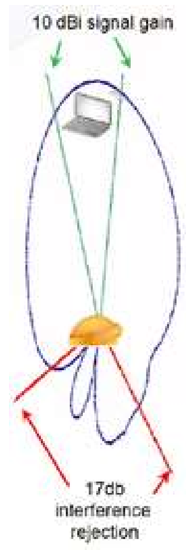

Traditionally, the only way an access point had to deal with dropped packets or CRC errors (corruption) was to drop transmit speeds. The PHY rate might ratchet down from 54 Mbps to 48, then 36, and so on until the client started acknowledging packet receipts. The slower the speed, though, the longer the radio stays on for a given data burst, and the longer the radio is on, the more susceptible it is to interference. So when conditions turn bad, you can get into this negative feedback loop that just slaughters performance. An intelligent antenna system will both dynamically change the beam orientation to a more optimal direction as well as forestall reducing the PHY rate until absolutely necessary.

Communication between the AP and client helps with these improvements, but it's not strictly essential. The bulk of optimization gets done by the access point. In the tests that follow, we did not use a Ruckus client adapter for two reasons. First, Ruckus told us that 75% of the performance improvement seen above standard 802.11n comes from the access point, so adding in a proprietary client would only yield a small benefit. Second, relatively few real-life scenarios can dictate the client. It's not like you can say, "Feel free to use my wireless LAN, but just make sure you use an XYZ adapter."

Ruckus uses "on-antenna" beamforming, a technology developed and patented by Ruckus under the brand name "BeamFlex." Essentially, BeamFlex uses an array of antennas and analyzes every packet to assess signal performance. Depending on the configuration, a BeamFlex access point can configure into any of thousands of possible antenna signal combinations. The access point monitors connections in real-time and modifies beams on the fly to fit dynamic conditions. Keeping with the MRC legacy, antennas that need signal boosting get boosted while those that don't are attenuated. This results in up to a 10 dB signal gain along the target beam as well as a -17 dB interference rejection in the direction of backlobes.

Can You Hear Me Now?

According to Ruckus, interference rejection can have an even more profound impact on performance than the target beam boosting. Imagine sitting in a crowded, noisy restaurant, and you're trying to have a conversation with the person across the table. Everyone, including your partner, is talking at the same volume, and you're having a very hard time hearing what the other person is saying. Seeing your listening problem, your partner talks a bit louder (boosting the signal a few dB), and this helps but nowhere near as much as when you cup your hands behind your ears to only let the "beam" of the other person's voice reach you while simultaneously muting a lot of that interfering background chatter.

With BeamFlex, software can steer access point beams dynamically, selecting the best path for each packet. The system will also automatically assemble a list of the 10 to 20 most commonly used antenna patterns. This functions a bit like the cache on a processor, keeping oft-needed data very close to the execution pipeline so it can be accessed more quickly. Ruckus has spent five years developing BeamFlex into its current form and fine-tuning the algorithms that help comprise its special sauce. Yes, BeamFlex is proprietary in that it doesn't adhere to IEEE 802.11n specs, but it's plainly interoperable with any standard WiFi client, and I think, if proven superior over competing approaches, Ruckus's approach to on-antenna beamforming could prove revolutionary in inspiring the next wave of wireless networking designs. We'll see in the following pages how well BeamFlex holds up against the competition, particularly against Cisco's chip-based beamforming option.

Beamforming Options

Legacy (802.11 a/g)	Implicit/Explicit (802.11n)	Beamflex (802.11a/g/n)
Inherently DSP (chip) based (tradeoffs required)	Inherently DSP (chip) based (tradeoffs required)	Additive to 802.11n chips at the physical layer
Doesn't use client feedback	Requires optional 802.11n client support (no client support today)	Feedback built into every 802.11a/b/g/n client, nothing needed
Provides up to 1-2dB gain with 1-2 radios	Provides up to 3dB gain with 2 radios	Provides up to 9dB gain
Can't focus RF energy in only one direction	Can't focus RF energy in only one direction	Focuses RF energy only where it's needed
No mechanism to avoid RF interference	No mechanism to avoid RF interference	Up to -17dB of interference rejection
Doesn't increase Wi-Fi "cell-size"	Doesn't increase Wi-Fi "cell size"	Delivers 2 to 4x increase in "cell size"

It's strange to say, but this may not be the first time you've run across BeamFlex. One of the very few times Ruckus ever poked its head into the mainstream press came in a roundup by Tom's Hardware sister site, Tom's Guide. This first-gen product had six antennas, each with about 60 degrees of coverage, arranged in a hexagonal design. You can still find this hexagonal layout in the company's 7811 access point, which managing editor Chris Angelini discusses shortly.

Our main testing, though, focuses on the Ruckus ZoneFlex 7962 access point. This is the enterprise version of the same BeamFlex technology, here upgraded to 19 antenna elements—10 horizontally polarized and 9 vertically polarized. Interestingly, though, according to Ruckus, the 7811 should perform very similarly to the 7962 when deployed in a single-story environment when only a few clients are present.

By now, you're probably wondering why, if beamforming is so amazing, Ruckus has kept such a low profile. The company says it's because the retail market sucks. In early 2005, Ruckus (then called Video 54) teamed up with Netgear to produce the seven-antenna RangeMax 824 router, which became hugely successful. But retail margins are precariously thin after considering support and marketing costs, and for whatever reason the love affair soured, culminating in 2008 with Ruckus suing Netgear for patent infringement over the third version of the 824. For the time being, Ruckus is choosing to pursue the enterprise and service provider markets, although it still keeps one or two etailers in the loop for consumers like us who want to get in on the goodness.

Test Gear: Ruckus 7962

Keeping in mind that the hardware you'd be most likely to buy would be consumer versions of the enterprise gear we tested, we wanted to approach this as a comparison of technologies more than a review of any given product.

By now, you know a bit about the Ruckus ZoneFlex 7962 access point we evaluated, but let's take a closer look. Unlike a consumer access point, which is typically built to look cool and enticing, the 7962 is the exact opposite. It's designed to blend into the woodwork—or, to be more exact, to the ceiling, where it will blend into the white tiling and look a lot like a dome light so as to go unnoticed by would-be thieves and vandals.

Naturally, with 802.11n radios, Gigabit Ethernet, and other electronics running full tilt, the access point can get pretty warm. This is why Ruckus build a heatsink into the bottom.

Under the hood, you can see how radically different the Ruckus design is from conventional access points. Note the circular arrangement of directional antennas.

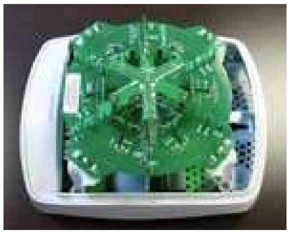

Once you get up into enterprise-class wireless equipment, deployments generally use a controller device to help manage traffic and coordinate multiple access points. We paired the 7962 with the mid-range Ruckus ZoneDirector 1000 controller running firmware 8.0.1.0 build 13.9

Test Gear: Cisco Aironet 1142 and Aruba AP 125

Ruckus' main competitor is the Cisco Aironet AIR-LAP1142N-A-K9 sitting alongside a Cisco 4402 controller running firmware 6.0.182.0. As mentioned before, this is our one possible candidate for chip-based beamforming to compare against's Ruckus's antenna-based beamforming (BeamFlex). The curious thing is that Cisco ships the 1140 with beamforming disabled.

Like all dual-band access points, Cisco uses two radios here, one for 2.4 GHz and the other for 5 GHz. The 1142 attaches three antennas to each radio, utilizing a 2x3 transmit/receive array. Cisco's beamforming uses two transmit antennas, so it follows that you can't have beamforming and spatial multiplexing operating concurrently.

So, it's curious that Cisco opts to prefer spatial multiplexing out of the box rather than its shiny new on-chip technology. Perhaps this is because the feature is so new that the company wants to ease into adoption. But could it have something to do with performance? We'll see.

Comparing the architecture of the 1142 against the Ruckus 7962 is intriguing. I couldn't dissect the Cisco unit, but I did find these photos in some online patent documentation.

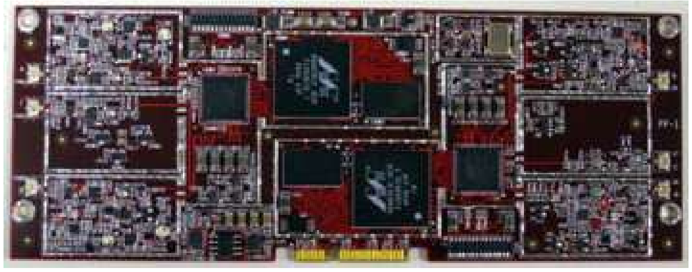

We also tested with an Aruba AP125 access point and Aruba 3200 controller with firmware 3.4.0.1 build 21611. There's not much to report here. This is a standard issue, enterprise-class, dual-band access point with three antennas you're never quite sure how to orient. The AP125 is pretty representative of "standard 802.11n," and I view it here as a performance baseline for comparison against Cisco and Ruckus.

Rounding out the test platform, I used a Dell 620 notebook with the Broadcom NetXtreme 57XX Gigabit Ethernet controller as the server and a Lenovo X61 with an Intel 4965AGN adapter (driver version 12.4.0.21) as the target client. The switch tying everything together was a 3Com 3CDSG10PWR OfficeConnect.

Test Environment

Knowing that beamforming should give 2x to 4x the signal distance of conventional 802.11n, I knew we needed a test environment bigger than a regular house could provide. Since I don't have access to any mansions, I went with the next-best thing: the 7,000-square foot headquarters office of Structured Communications in Clackamas, Oregon. This office takes up most of the 11,000-square foot top floor of its building. Special thanks go out to Structured for letting Tom's Hardware invade their space for two days of setup and testing.

I tested in five locations, seeking to get a decent mix of scenarios. For reference, here's the office floor plan with each of the five locations marked. The numbers noted next to locations 2, 3, 4, and 5 are the approximate distances in feet from these test sites back to location 1, where the access point sat.

Test Location 1. This is home base, where we set up out small mountain of equipment. No, not everything got used—we had some backup gear, just in case. In some of the pictures shown previously, I stacked the access point on top of the controller for convenience. During actual testing, these were kept separate to reduce interference, as you can see in this image of our Cisco setup. When testing in this location, the client was about two to three feet from the AP—essentially point-blank range.

Test Location 2. This is pretty much a straight shot down an aisle and across the office floor from location 1. You can see the meeting chairs in location one way off in the distance.

In retrospect, I should have maintained line of sight between the client and access point for this test. Instead, the access point is tucked just around the right corner of the doorway. So rather than line of sight, I instead have the straight path moving through a long wall line. Who ever gets straight line of sight on indoors wireless, anyway? And if you're curious, yes, that's me looking so out of place doing Zap testing that two people had to stop and stare.

Test Location 3. This spot made sense, as it was located straight across the building and entailed pushing the signal straight through several walls—no way to bounce your way out of this one. As you can see, the ThinkPad is only detecting four nearby access points in this fairly quiet location. On other areas around the floor, I picked up over a dozen.

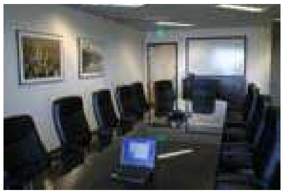

Test Location 4. Laterally, this was our longest and most isolated location, sealed off within another meeting room at the far end of the office. In the world of indoor WiFi, reaching through 100 feet of crowded, highly-sectioned space is insane. You'd normally want at least one more access point in operation here. In fact, Ruckus told us that four of the 7962s would be advisable...compared to 10 or more conventional access points. So, if we could get decent performance at this range, it would be a near-miracle.

Test Location 5. As you walk out the door of Structured's office, there's an open area largely filled by a glass-walled chasm that runs the height of the building. In order to test signal reach in three dimensions instead of the usual two, I set up shop on a table at ground level, down two floors. To determine the 95-foot signal distance, I had to use the Pythagorean Theorem, something I vowed in eighth grade Geometry class would never happen during my post-school lifetime. Never say never.

Test Apps and Methods

I used two applications during testing, Zap and Chariot. These examine UDP and TCP packet performance, respectively. You don't see UDP tested very often. Everybody simply loads up Chariot or iPerf, does some time tests, and that's about it. For conventional file transfers and similar everyday tasks, this is an appropriate methodology. However, UDP is what you use for streaming video. It's a faster protocol because the server system doesn't have to sit around waiting for receipt confirmation from the client. With UDP, you simply blast out a stream of high-speed packets and hope they get to their destination, come what may.

You've probably never heard of Zap because Ruckus developed it in-house for testing video streaming performance. To the best of my knowledge, this is the first extensive use of the tool in a mainstream review. As it was, I was sworn to not let the application out of my sight, so apologies in advance for not making it available to readers.

With that said, there's no dark magic to Zap. It simply takes a reference load of data and sends it between the server and client using UDP. The transfer is divided into percentages of the total work load, with each step being one-tenth of a percent. At each step, throughput rate is recorded and the number shown by the software is the lowest packet speed recorded up to that point in the transfer job. This is why Zap numbers look really fast at 1%, average at 50%, and very slow at 99 percent.

For our purposes, we're most interested in the average and lowest numbers. When it comes to video, you don't care what the fastest or average sustained rates are. You care about the slowest speeds, the weakest link in the wireless chain, because this will be the key factor in determining your video-watching experience. If you sustain a 70 Mbps connection 95% of the time but occasionally drop to 15 Mbps for whatever reason, then those drops are going to translate into dropped frames and hiccups if you're watching an HD stream with a 19.2 Mbps data rate. You can see a real-world example of this in the chart shown here, which (spoiler alert!) is the Chariot throughput data for Cisco's 1142 access point at short range.

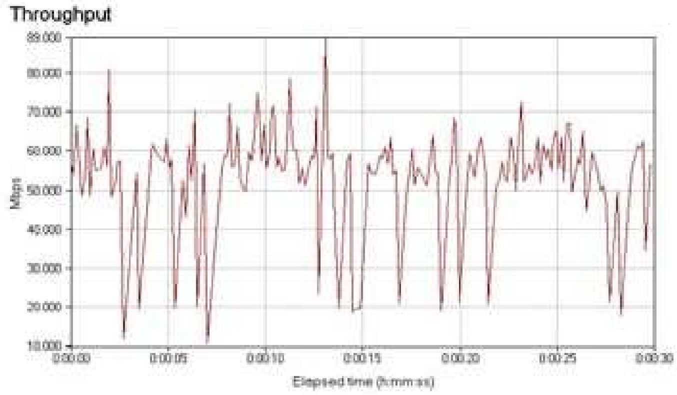

As mentioned previously, many things can impact wireless throughput, including the orientation of the client. There are three antennas in most 802.11n-equipped laptops, and in three dimensions these work (once again) a lot like rabbit ears. So I actually ran each test four times, rotating the laptop a quarter-turn for each test. The results were then averaged together.

Additionally, since each access point has the ability to run at either 2.4 or 5 GHz, I ran all tests on both radio bands. It's possible for a client that associates on one band to hop to the other if conditions deteriorate, but it's not common. Client sessions tend to stay loyal to whichever band they first associate with. Hence it's important to get a good idea of how both bands perform.

Not least of all, I made sure that power management in the Intel client driver was set to "highest." Otherwise, when running on battery power, performance can be more prone to fluctuation. If you're curious, that command line business sitting under the driver window shown here is Zap at rest.

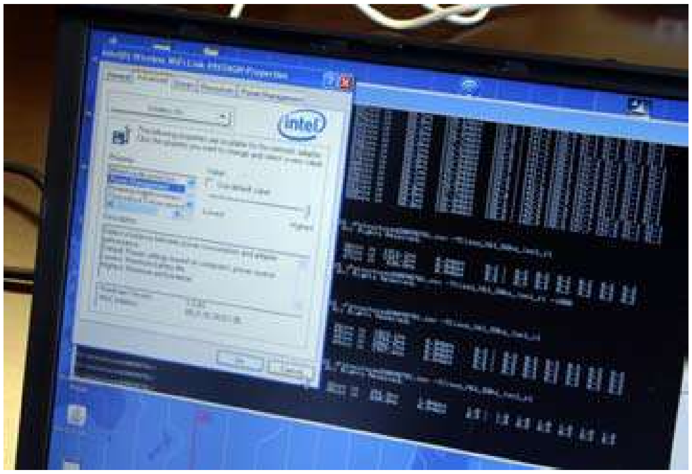

Zap in 2.4 GHz, Average

All right, let's get down to business. In location 1, with the client and APs practically kissing, we see very solid numbers across the board in our Zap 50% tests. I was actually a little surprised to see Ruckus pull so far ahead in location 1 because beamforming shouldn't provide much benefit at very close distances. We see this evident in the twin Cisco scores, which show beamforming only giving a 2 Mbps boost.

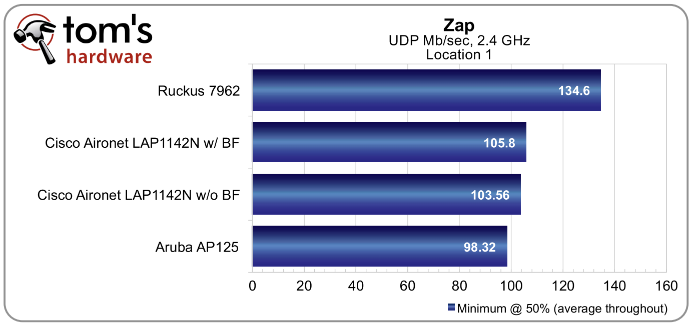

Our next two distance tests fall near expectations. I'm not surprised that Ruckus won these tests, but I am surprised that Aruba lagged so far behind Cisco, even without beamforming enabled. Speaking of which, location 3 shows Cisco's beamforming advantage, but it's interesting that location 2 does not, perhaps indicating that the arrangement was closer to line-of-sight than I might have imagined.

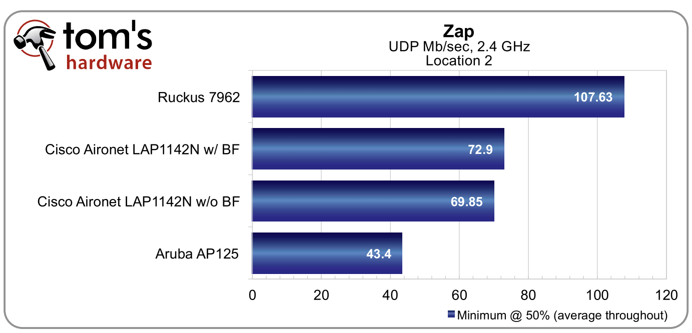

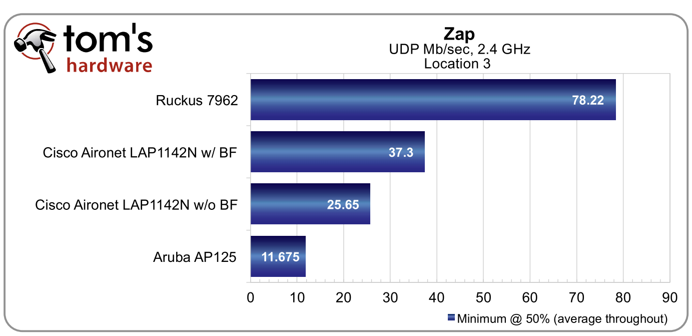

Secluded off in that meeting room at location 4, Aruba drops off a cliff, failing to average even 5 Mbps. And again, if anything, Cisco's beamforming appears to impair performance slightly. Very odd.

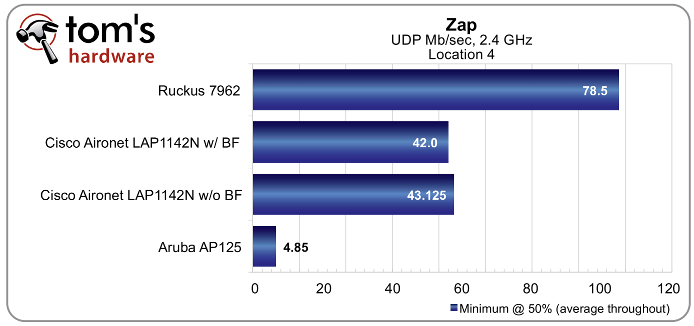

In location 5, it's more of the same, although Ruckus finally shows some signs of weakening. Aruba can barely hold a connection at less than 1 Mbps, but Cisco does relatively well—especially with beamforming disabled.

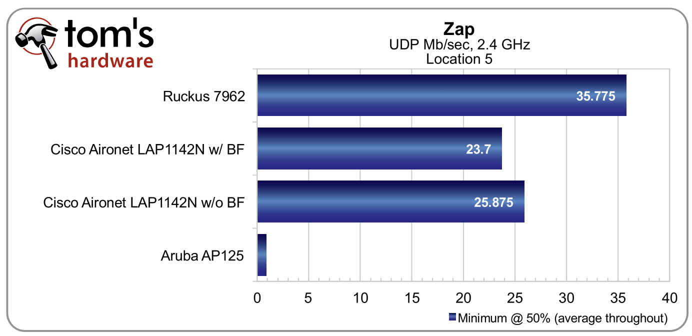

Obviously, the throughput number you expect for bare-minimum acceptable performance will vary based on application. If you want to hold two HD streams, then you need at least 40 Mbps and preferably more for possible sporadic interference. At 2.4 GHz, none of these access points could handle this, but who would really set up a scenario like this in real life? As we continue, keep in mind that these tests are meant to prove or disprove the viability of WiFi beamforming, not necessarily to show how equipment should perform in a given situation.

Zap in 5 GHz, Average

Switching to the 5 GHz band, we notice two things immediately. First, Aruba has a pulse! The AP125 handily blows past Cisco at close quarters. More importantly, though, both Ruckus and Aruba show 5 GHz numbers that are double the throughput seen on 2.4 GHz. So if you have a choice with your config, give 5 GHz a try, or at least do the throughput comparison. You may be stunned.

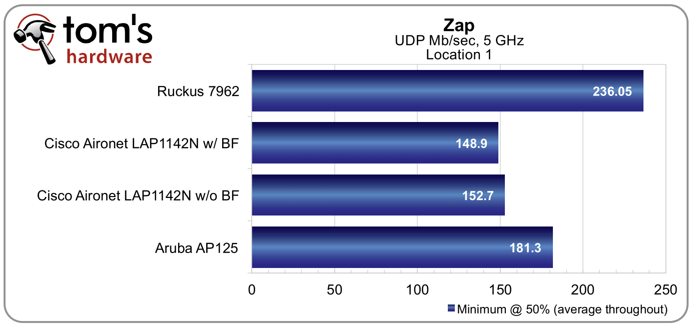

Life looks the same in locations 2 and 3, with Aruba still outpacing Cisco and the latter showing no benefit at all from its on-chip beamforming.

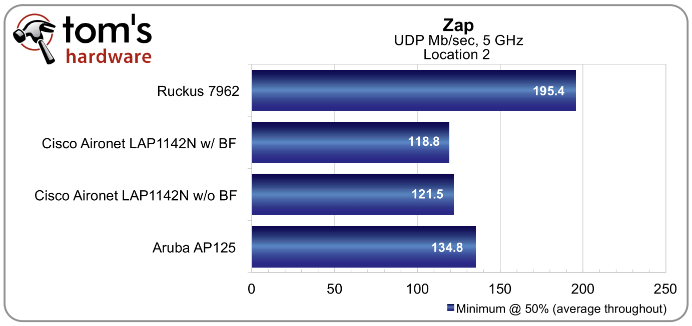

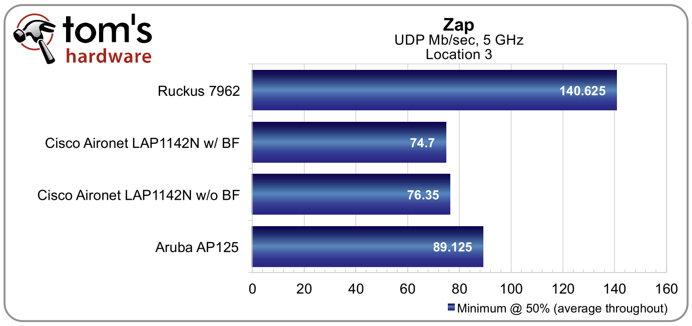

Finally, in our location 4 fortress of solitude, Cisco beamforming shows some utility, just nudging past Aruba's non-beamforming results. Ruckus still easily kicks dust in the others' faces.

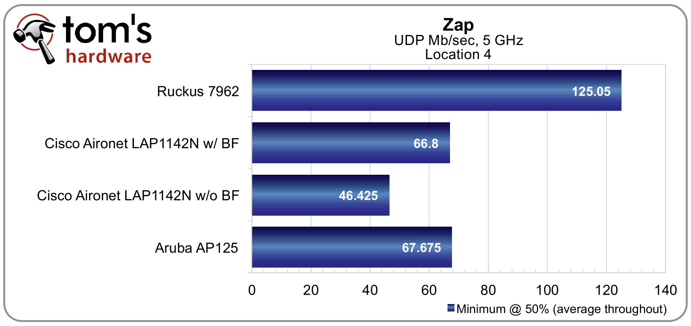

And here's the kicker. Down in location 5, Aruba clings by a thread to 0.1 Mbps. Cisco can't find a connection at all. Ruckus still nearly holds 25 Mbps. If ever there was a case for a technology to adopt in long-distance scenarios, this is it. Also, harkening back to an earlier point, now you start to see why Ruckus thinks BeamFlex can cover a large space with considerably fewer APs than alternative technologies.

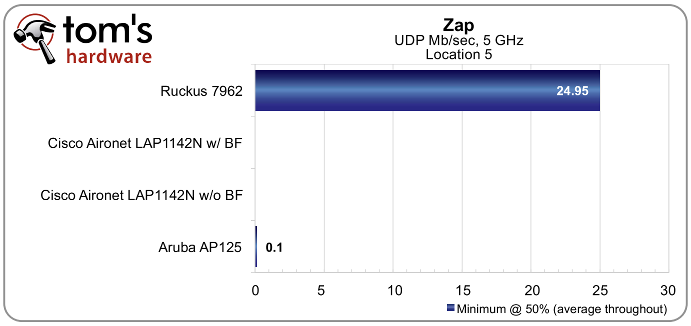

Zap at 2.4 and 5.0 GHz, Minimum

For all you video buffs, we come to the critical 99% Zap tests. We nearly decided not to include these results in this article because they're so ridiculously one-sided. Even still, there are some (disappointing) surprises here.

Right off the bat, the big question: What on Earth happened to Cisco? How do you have 1 Mbps performance pockets at arm's length? Aruba didn't. Ruckus sure didn't. This is embarrassing.

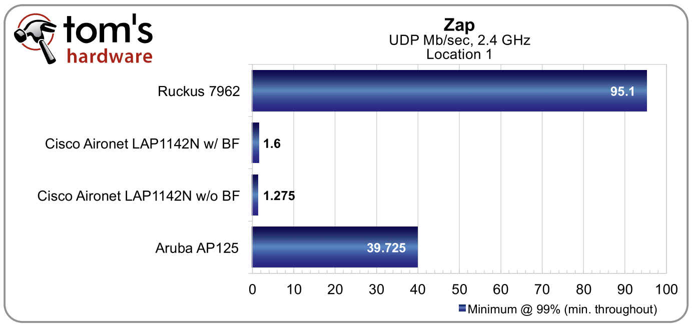

For the rest of our 2.4 GHz 99% tests, Aruba gives up its fleeting advantage, joining Cisco at the 1 Mbps Pity Party. But at least nobody lost a connection. That's something.

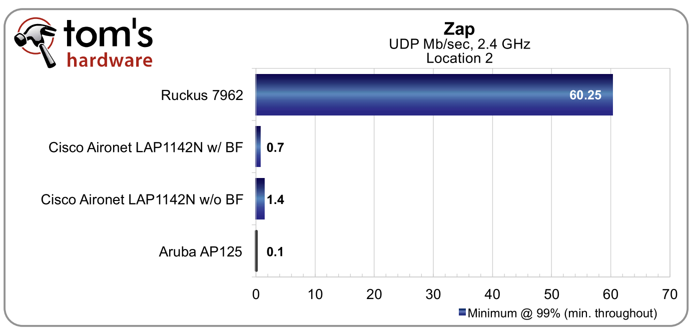

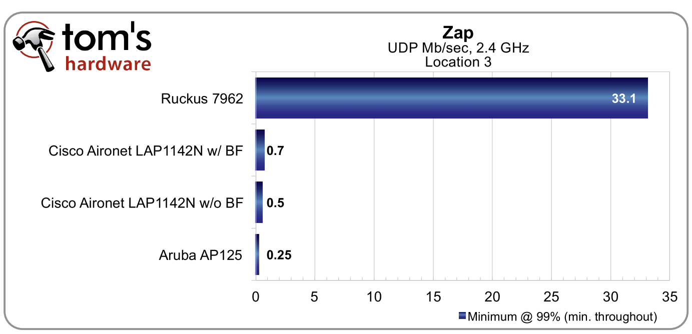

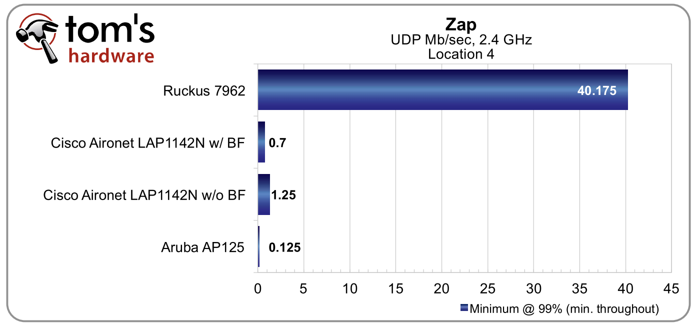

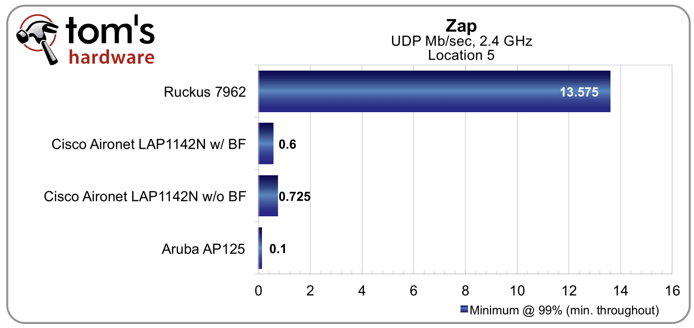

When we switch over to the 5 GHz band, there's more oddness. Ruckus exhibits none of the location 1 black holes found with Aruba and Cisco. Why? We're not sure. This is much worse from Aruba and Cisco than I expected. So I grabbed the trusty Wi-Spy Spectrum Analyzer to scan location 1 for interference.

Wi-Spy results showed nothing out of the ordinary. In fact, the conclusion that both of the Aruba and Cisco products are unusable for video streaming is inescapable, although not unexpected. Actually, these results are typical for 802.11n products, which is why you haven't seen many carriers or manufacturers marketing WiFi for video. Yes, there are a few exceptions, but the proof is in the pudding, not the PR. It's simply been impossible to sustain a quality experience. Once more, recall my shock at the beginning of this piece.

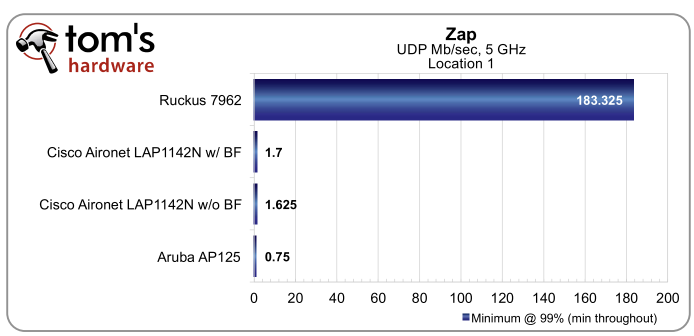

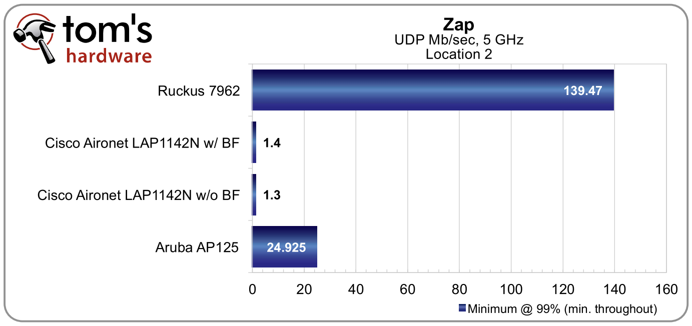

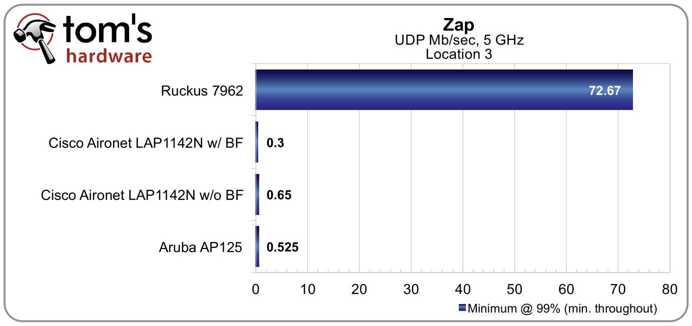

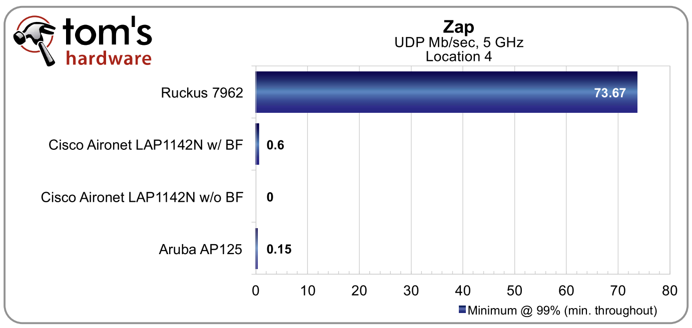

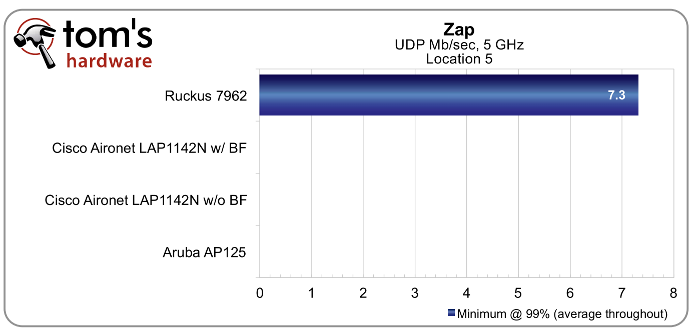

Chariot at 2.4 GHz

By the time I got to Chariot, I was already 1.5 days into a 2-day testing window. In four hours, my test space and many thousands of dollars of on-loan equipment were going to turn into pumpkins. So, as much as I'd like to tell you there was some technical reason for why I only tested the 7962 against the 1142 with beamforming enabled, the truth is that I ran out of time and had to make a choice. I decided the few hours I had left for Chariot testing would be a battle of the beamforms—on-chip versus on-antenna. Fight!

Quite the difference between TCP and UDP numbers, no? Also, remember that 2.4 GHz tests at close range can show half the TCP throughput speed of 5 GHz. This is why Netgear put its HD/Video wireless kit on the 5 GHz band. No one should be too surprised when Ruckus only manages an average throughput of 67 Mbps at location 1, although this is double what Cisco pulls in. Numbers decline incrementally at locations 2 and 4, with Ruckus showing more of a drop-off. The contenders reverse roles at location 3, with Cisco showing the greater loss over distance.

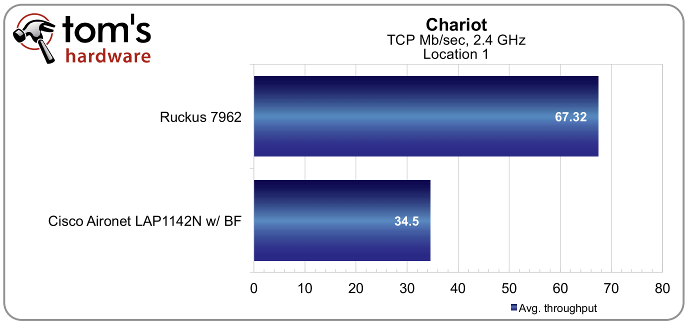

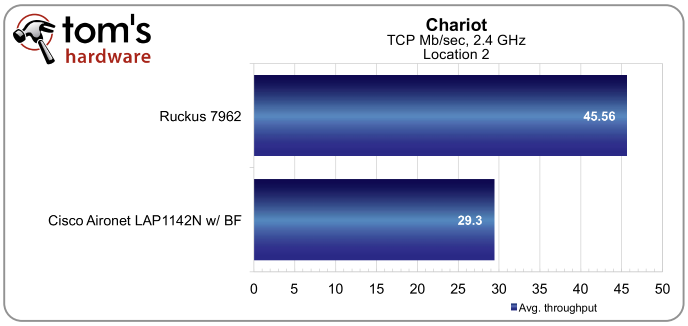

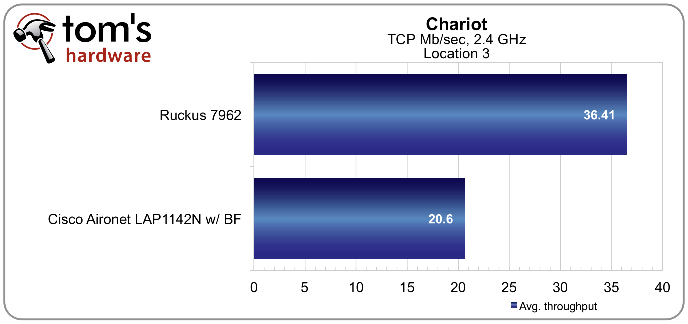

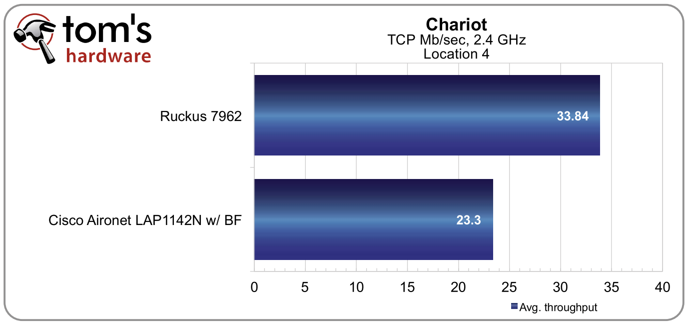

Down in location 5, both access points are able to hold a connection. The numbers we see here show Cisco hitting about half of the throughput we saw with Zap on 50% tests while Ruckus only falls off by about one-third.

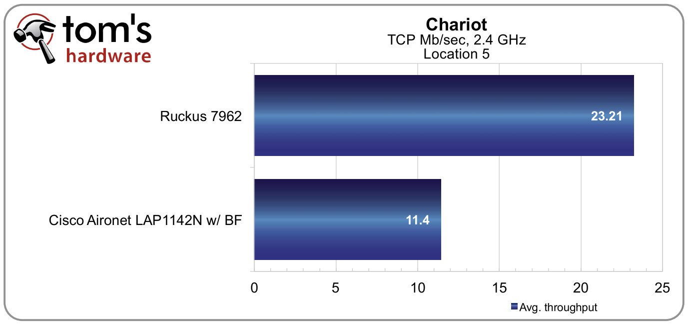

Chariot at 5 GHz

There's no point in kicking a dead horse. Across the first four locations, BeamFlex beats Cisco's on-chip beamforming by anywhere from roughly 1.5x to 3x. Location 5 is the exception because the 1142 failed to connect.

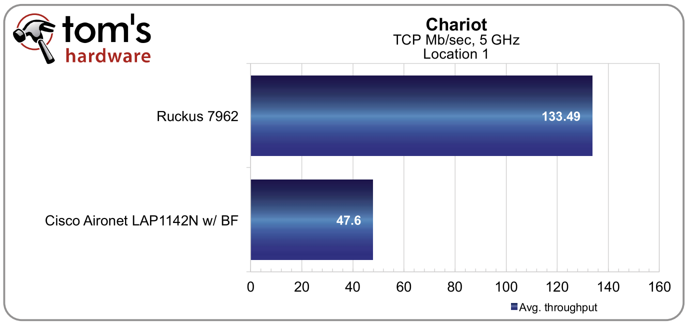

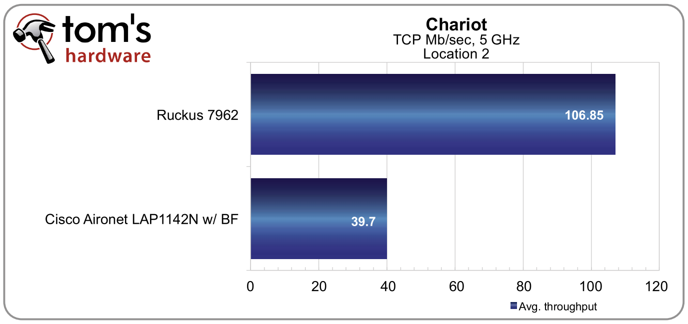

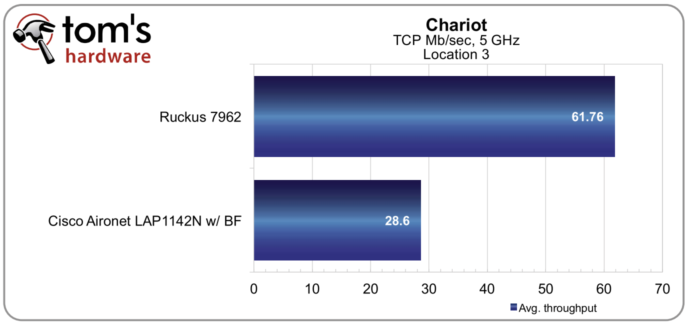

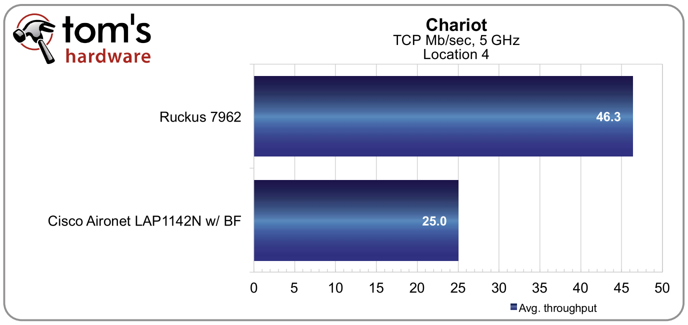

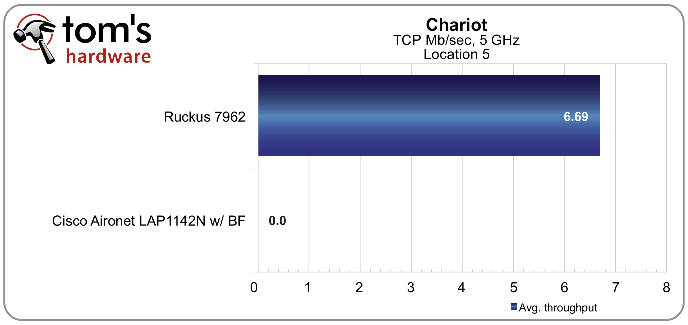

It might also be helpful to view all five locations in line graph format. As you can see, the relative advantage enjoyed by the 7962 is greater at the closer distances.

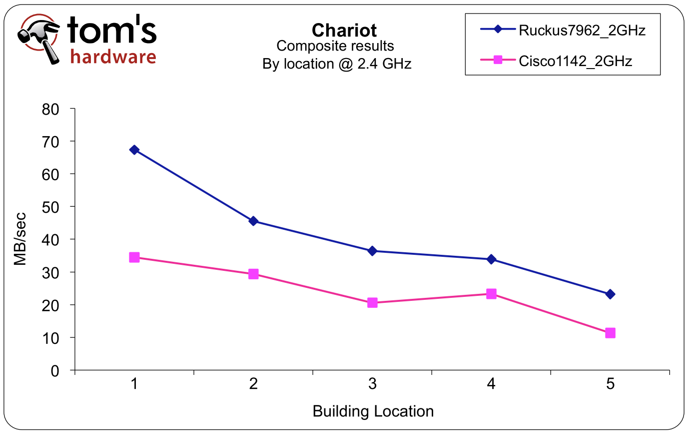

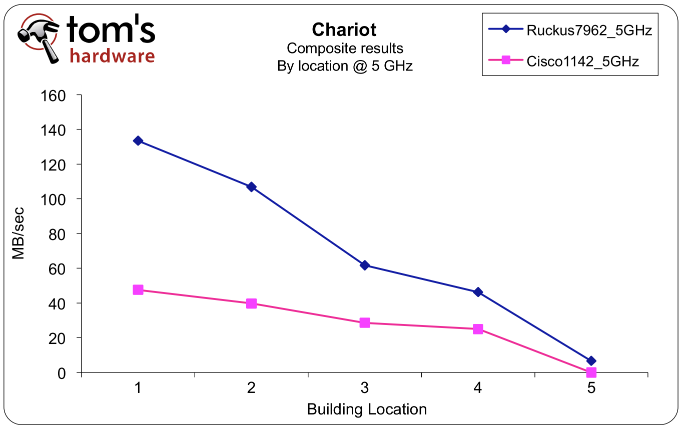

Conclusion

I don't need to sell you on how the world is moving to being video-centric. If you've stuck with us this far, it's probably because you see the writing on the wall and know that you're either going to have to string Ethernet cabling through your walls or else find some sort of wireless solution for your present and future video needs. Moreover, CAT5e or CAT6 only get you so far. The day when most of us have more video-oriented devices in our pockets than on our desks and walls can't be far off. Woe to the hapless consumer who doesn't have a wireless backbone up to the task of distributing streams wherever needed.

I don't want you to walk away from this article with the single message: "Ruckus rocks!" That's not the point. What we've seen here is that on-chip beamforming, at least in the way that Cisco has implemented it on the 1142, barely has any effect. No wonder the feature arrives disabled today. However, beamforming in principal can have a tremendous effect. Ruckus clearly shows that all 802.11n up to present has merely been a preface. This is the next level, and so far there's only one company standing on it.

My hope is that this article will raise some eyebrows and spur the industry to advance. With the 1142, Cisco largely relied on existing designs. The level of innovation was minimal, and it shows. We need more companies like Ruckus willing to invest two or three years into rethinking the problem and taking wireless communications forward in multiples of performance, not single-digit gains. Yes, there will be interoperability wars. Yes, the pricing will be double what you pay for non-beamforming equivalents. But in return, we'll be able to do and enjoy things with our wireless LANs that simply can't be done today. ■