WHITE PAPER

802.11n and Voice: High Throughput Does Not Guarantee Good Voice Quality

Author: Joe Epstein, Sr. Director of Technology
Date: May 2010

TABLE OF CONTENTS

INTRODUCTION

Wireless infrastructures are increasingly essential to support growth in use of mobile devices. Users expect anywhere, anytime access – and organizations increasingly must be able to deliver. For enterprises, 802.11n is clearly the right choice for designing a new wireless network: it enables higher performance than wired Fast Ethernet and has extended range, allowing applications to run wirelessly that were originally designed for the once-fatter wired pipes.

But what about such a purely-wireless application as voice over Wi-Fi? Voice was designed not to have the bulkiest throughput requirements, but to need a predictable, high-quality, always-on service. Mixing it with the raw speeds of 802.11n and the results can be quite surprising: **high throughput does not guarantee good voice quality**. In fact, combining the two can simply ruin call quality, greatly reducing the maximum call density whenever voice and data are used together on the same network.

The only way out is for the WLAN technology to be designed explicitly around optimizing the *application* experience, not just the bits and bytes of connectivity. Strong wireless innovations are needed to simultaneously get the best voice and data performance, providing the predictability and stability not seen on traditional wireless LANs.

WHAT VOICE REQUIRES

The reason why a network designed for high-performance data can both suffer from and destroy the quality of voice is that data and voice have different requirements. Voice, as it is sent over IP-based networks, is composed of short but frequent transmissions, carrying small packets of audio from one party of the call to the other. Overall, voice throughput demands are very low—a typical voice call requires as low as a tenth of one percent of the throughput of a web download. But the requirement to precisely, and regularly, send traffic at fixed intervals—such as 50 times a second—means that the voice traffic is placing a higher demand on the network in terms of its need for regularly free *air time*.

Let's use an analogy. At a busy intersection, a crowd of people crossing with the green light will only disrupt traffic for a short amount of time. However, jaywalkers crossing whenever they want—even if there is only one person who is crossing at a time—can bring traffic to a grinding stop for far longer stretches.

Because a voice call will not even work unless these frequent, but tiny, packets get through, 802.11n gives voice the highest priority over the air to allow voice to transmit ahead of data. However, this triggers another problem. A flood of voice calls at once will jam up the air, as too many high-priority voice packets compete for being transmitted at exactly the same time. In the context of the jaywalking analogy, this is like adding a single turnstile to the pedestrian crossing, watching both the car (data) and pedestrian (voice) traffic slow. In fact, with 802.11n, the cars can go much faster, threatening the pedestrians even more.

Finally, voice requires predictable RF service. Web traffic is flexible—*elastic* is the term for it—and can survive without causing substantial user disruption even when the available bandwidth and quality varies. This translates to wireless networks, where simple web access may simply paper over unstable or fluctuating networks. But voice is inflexible—*inelastic*—and thus cannot tolerate fluctuating networks. The wireless must be the same from minute to minute to keep the voice quality from fluctuating with the network; it must be the same from day to day if the maximum number of voice calls needs to be equally predictable across those days.

WHY MICROCELL DOES IT WRONG

The legacy way of building wireless networks—even 802.11n networks—is called microcell. Microcell wireless infrastructures are only designed for consumer and casual use: increasingly unable to handle rapid growth in devices, latency-sensitive applications and performance requirements. Microcell refers to the dual notions of deploying more access points than needed, at low transmit powers, low ranges, and different channels, and allowing the power levels and channels of the access points to vary frequently over time.

The reduced range and frequent changes of microcell were introduced almost a decade ago, to improve the quality of the network compared to the single-access-point technologies of the 1990s. However, the technique is only a partial solution from an older time when checking email was the leading application, and the limitations of the legacy technique leads to a number of difficult challenges for voice.

The first problem microcell introduces to voice is from the reduced range. Microcell architectures reduce the range of access points to reduce the interference caused within their architecture by the coverage overlap of adjacent access points. However, this reduction means that each WLAN phone spends more of its time in regions at the edge of each access point's coverage, where signal strengths are low and quality is poor, rather than in regions of high signal strength.

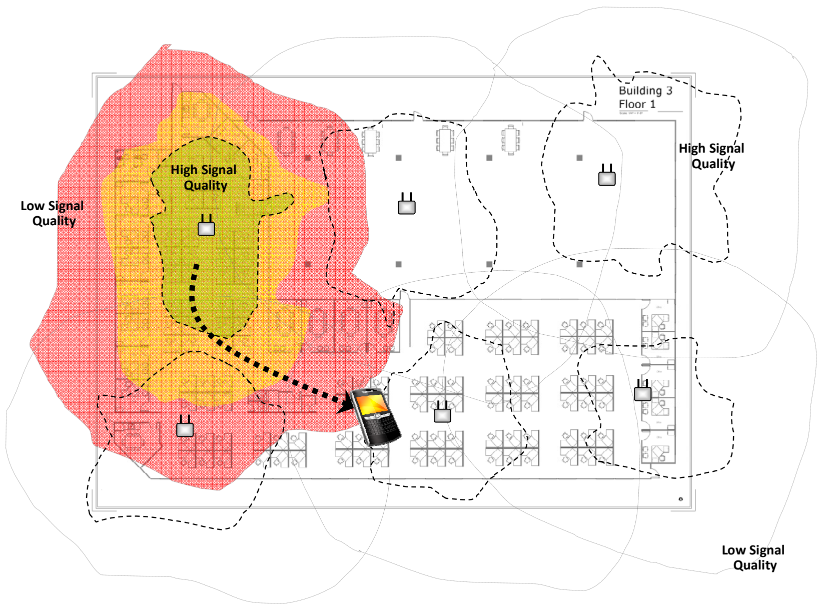

Figure 1 – Microcell causes phones to see signal quality drop severely for every access point, leading to network instability, poor quality, disruptive handoffs, and a sharp increase in the complexity in understanding and diagnosing the wireless network.

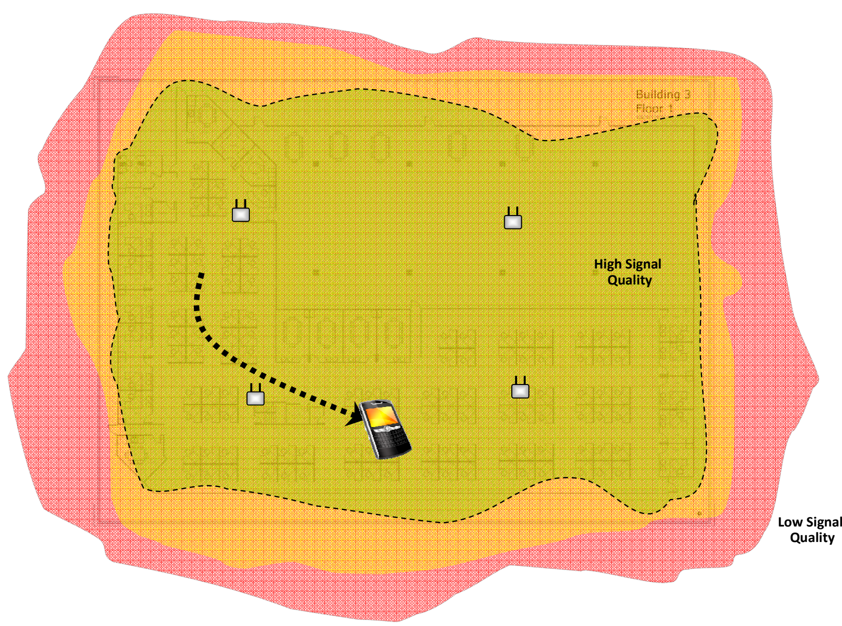

Figure 2 – Meru with WLAN Virtualization produces a network with stable, consistently high quality coverage throughout the deployment, ideal for predictable voice quality with no penalties for handoffs. Phones never experience poor signal quality, because the network always hands them off to the highest-quality access point.

The single-channel network depicted in this figure has roughly the same wireless capacity as the microcell network in Figure 1, yet with fewer access points needed.

The second problem microcell introduces to voice is from the use of multiple channels to produce only one 300Mbps layer of wireless networking (or two such layers for dual-band access points). With a microcell architecture, if one access point is on a given channel, no access point in range of that one can use the same channel. Therefore, they all must choose different channels—even though all of the access points belong to the same network. Because microcell technology has had a run of nearly a decade, many people are familiar with this technique and may not find it surprising. However, for wireless technologies, this approach is highly unnatural—and unnecessary—for producing a uniform, high-quality wireless network. Compare this RF problem to broadcast radio. If the microcell method for channel allocation were used for car radios, then a person driving would only be able to find one FM station at a time—that means only one type of music for the entire trip. Odder still, although the content stays the same, the actual number of the channel constantly changes while the car travels. This peculiar paradigm became a part of wireless LANs only because enterprise-sized networks needed some stop-gap measure to prevent the interference, and this stop-gap nature is why microcell architectures are legacy.

The variation of channel from place to place directly leads to the problem of voice handoffs. Because the channel changes, a phone moving through the network is not only forced to experience a low-quality region at the outer edges of each cell it moves through, but upon hitting that low-quality region, it has to scan through a multitude of other channels to find which one is providing better service. This happens at each access point cell boundary, which can be every couple of dozen feet in a microcell network. The low-quality fringes and the time taken away to scan will directly impact the quality of the voice call.

The third problem microcell introduces to voice is from the frequent changes to access point ranges and channel selections over time. Microcell architectures perform this frequent auto-configuration as an attempt to deal with the heavy burden of contention and co-channel interference generated by the clients that move though the network, and as clients move, their network is thus forced to move as well (channels and power). But there is a stark difference between taking contention and interference for granted and trying to balance two evils than directly managing the contention and interference to reduce or eliminate those evils entirely. In fact, changing power levels and channels is actually the wrong thing to do for the stability of the network. The changes the microcell networks make are unpredictable by nature and cannot be improved to result in predictable changes to the RF environment. The frequent variation means that the network is never the same way twice—unpredictability in, unpredictability out, so to speak—and that unpredictability carries directly over to voice quality and capacity. And although the changes a microcell network makes in configuration may make an improvement to some metric or another, they are rarely the changes that improve the quality of voice, as voice quality (predicted or measured) is not an input into the algorithms used to make the changes.

The fourth, and perhaps most important problem, microcell introduces is that it leaves the handoff decision to the client. It is impossible to create a client that hands off predictably well in most environments without per-installation tuning. Microcell, however, places control of the handoff decision in the handoff decision in the client, and thus is at its mercy. Combine on one network multiple revisions of voice clients, or clients from multiple vendors, and predictability is very difficult to achieve.

In short, voice needs a stable, consistent, network to carry its inflexibly regular traffic. The very nature and design of microcell networks, on the other hand, are based on unstable, inconsistent coverage, varying moment-to-moment. The consequences of running voice on such a network are impacts to both CAPEX, for more access points at lower power levels, and OPEX, where additional resources have to go in for trying to manage network quality and service levels for a constantly-fluctuating network.

Nevertheless, 802.11n wireless is the right technology to use for a wireless network. How can this paradox be resolved?

WHAT MERU DOES RIGHT

Meru takes the approach of recognizing that the stability and predictability of a high-performance, high-capacity wireless network is paramount to being able to run applications such as voice over wireless. Meru's WLAN virtualization architecture is able to introduce wireline-like predictability by moving ahead from the notion of ever-changing, overly-dynamic microcell to a system that strives to maintain the same coverage and performance day-in and day-out, limiting changes to the minimum necessary to preserve the quality of the network while allowing for optimal behavior in shared WLAN air spaces.

Meru's WLAN architecture is based on two principles: pooling the wireless resources into a predictable foundation using *Virtual Cell™*, and partitioning those now-homogenized resources to each client using *Virtual Port™*. Together, the WLAN virtualization architecture brings the benefits of reduced CAPEX and OPEX directly to wireless—for both voice networks and data networks—surpassing the benefits promised but not often delivered for "rightsized" WLAN.

Virtual Cell consolidates the disparate cells of wireless coverage from each access point into one layer of coverage, turning up the power levels that microcell turns down, and removing the low-quality cellular fringes by setting all access points in the Virtual Cell to the same channel. Together, networks can be planned out as simply as one places light bulbs, just ensuring that there is adequate coverage (illumination) everywhere without worrying about from which bulb the light shining on one area comes. Rather, the focus is coverage, with channel layering for extra capacity, where needed. No more unpredictable channel and power fluctuations, as self noise, or interference from other access points and their clients, is handled using advanced algorithms that avoid having to reconfigure the network to adapt to the client, but rather ensure that the clients are transparently made to connect directly to the optimal access point. The higher power levels of a Virtual Cell network, as well as the elimination of co-channel interference concerns and the related site-survey planning process, leads to a sharp reduction in CAPEX for installing a Meru WLAN over a microcell WLAN—up to 30% fewer access points, and the elimination of costly pre-installation RF planning and prediction services.

Virtual Port focuses the pool of high-quality wireless established within the Virtual Cell to each user, customizing the network experience by establishing the notion of a wireless "switch port". Each client is given its own, unique virtual WLAN, where only it and the network are members. Within each Virtual Port, the client is given its fair share of service—no more, no less—no matter what the other clients are doing and how they use their network. This provides per-device isolation in much the same manner as Ethernet switch ports do, when compared to the preceding and obsolete hubs, where every port accessed the same, shared resource rather than one unique resource per port. Furthermore, because each client remains associated to its same Virtual Port—without having to reconnect—no matter how many access points it hands off between, the network takes control over the handoff decision, providing predictability and uniformity to the chaos of microcell, where each client makes its own decision.

Meru's WLAN Virtualization technology allows for zero-millisecond, deterministic handoffs between access points within the Virtual Cell, providing precisely the sort of continuation of quality that voice calls require. Devices hand off in a WLAN Virtualization network only from the point of view of the network, as the device sees no break in its connection and no establishment of a new connection. Instead, the device gets to take its connection with it, no matter where it moves on the network. No between-access-point drops, no missed connections.

WLAN Virtualization is able to solve the key problem of data interfering with voice in ways microcell networks cannot address. For Meru, because each client gets its own Virtual Port, the network is able to ensure that the voice traffic doesn't just get prioritized over data (and thus jammed into a bottlenecking turnstile), but rather is metered out with the optimal amount of voice capacity. This ensures that voice quality is maintained under dense voice conditions, while at the same time eliminating the problem of voice coexistence, thus removing the obstacles for higher data performance. Meru's WLAN Virtualization has the tools to optimize both for voice and data at the same time.

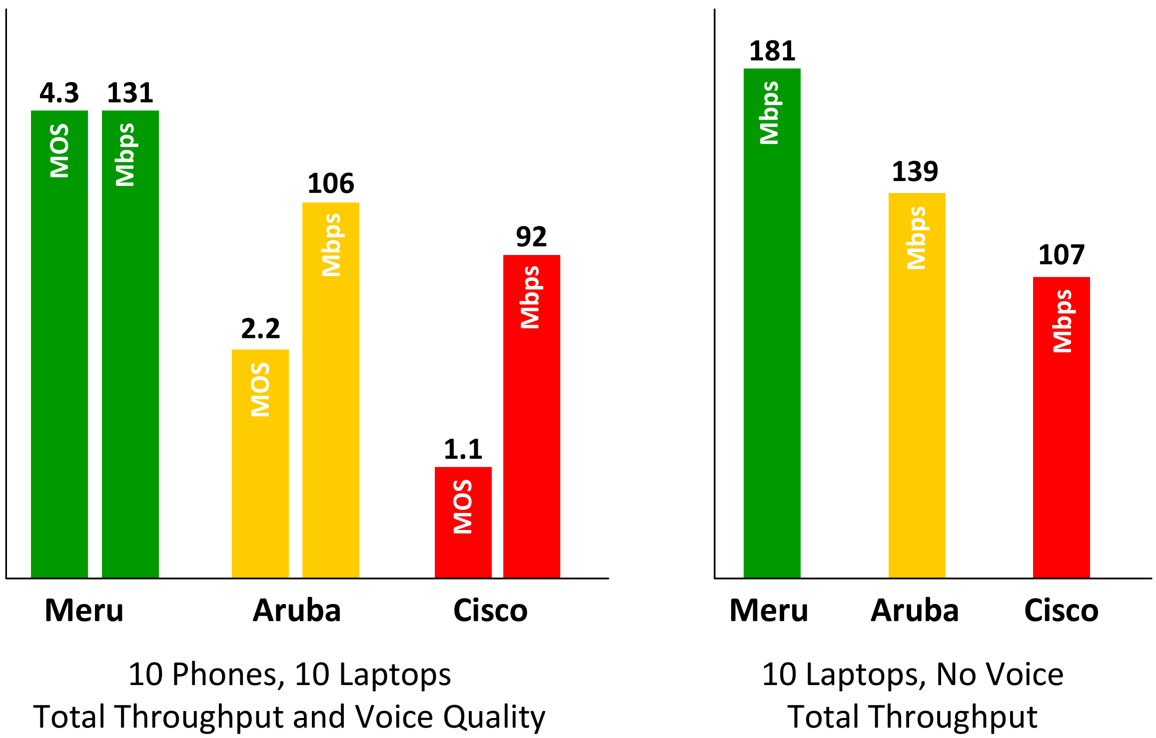

Figure 3 – Proof points for 802.11n voice, from a comparison by *Novarum*[1]

CONCLUSION

Networking technologies have evolved, and just as the raw link speed technologies have improved, the architectures have had to improve to keep up. For wireless, the introduction of

[1] "The Limits of WLAN Capacity" http://www.novarum.com/novarum_blog/2009/03/the-limits-of-wlan-capacity.html

previous wireline applications as mission-critical components of the wireless network creates a number of challenges that are only intensified by the introduction of new link standards—especially 802.11n. Thankfully, architectures evolve in wireless as well, and when it comes to the challenge of voice—of building a predictable, stable, toll-grade quality network that must also serve the needs for the existing and future IT data use—switching away from legacy microcell-based wireless to the more flexible, stable, and diagnosable WLAN Virtualization architecture can address these problems, all with a lower CAPEX and OPEX than previous-generation solutions can provide.