White Paper

voice. Data. wireless. *Become one.*

MicroCell Deployments: Making a Bad Problem Worse for Pervasive Wireless LAN Deployments

Conventional Wireless LAN Wisdom

Wireless LANs are widely used in enterprises, homes and hot spots to provide connectivity to the Internet or network resources without wires. They are on their way to becoming ubiquitous, with almost everyone with a mobile laptop having the ability to connect to a Wi-Fi access point at home, on the road or at the office. Wireless LANs provide up to 54 Mbps of over-the-air bandwidth with 802.11g, reaching effective throughputs of ~24 Mbps in ideal conditions. 802.11g is the most popular wireless LAN standard currently being deployed, as it offers backwards compatibility with the previous generation of 11 Mbps 802.11b devices.

For larger hot spot or enterprise deployments, multiple access points are used to create pervasive, contiguous coverage. This allows users to move from one area to another without losing connectivity to the network.

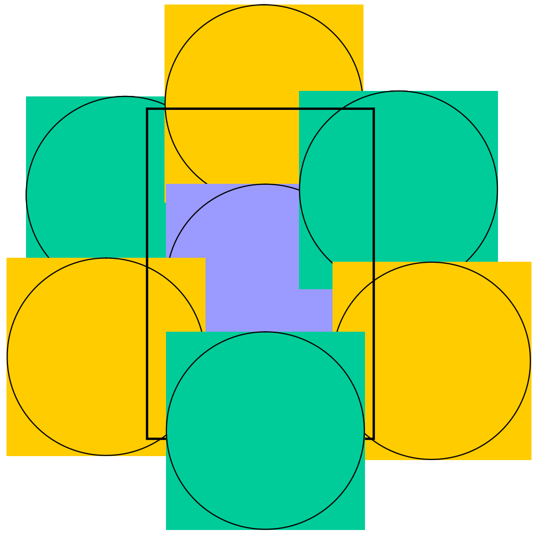

Figure 1: Multiple access points are placed contiguous to each other to allow for unbroken connectivity across a larger physical area.

As the client moves from one area to another, it recognizes that the signal it is receiving from the access point it is connected to is getting weaker. When the signal reaches a low enough threshold, the client will disconnect and search for a stronger signal from a new access point. This is termed handoff and is typically not noticeable by users of data applications.

The scenario described above sounds simplistic. Place access points such that you have contiguous coverage and mobile users will be able to roam and remain connected without any problem. Unfortunately, the situation is more complicated than that, and pervasive Wi-Fi deployments require enormous planning and effort in order to operate at a reasonable performance level.

The Challenges

Four basic facts about how Wi-Fi operates cause pervasive wireless LAN deployments to be difficult.

1) *Reducing data rate vs. cell size*
   With a Wi-Fi access point, the data rate reduces as you move farther away from it. Close to the access point, you may experience speeds of up to 54 Mbps (for 802.11g), but as you move farther away, speeds will drop to 48, 36, 24, 12, 9, 6, 2 and eventually 1 Mbps. Access point range can extend up to 300 feet indoors at the 1 Mbps rate. So, to achieve higher overall throughput, access points must be placed closer together.

2) *Limited non-overlapping channels*
   The 802.11b and 802.11g wireless LAN standards operate in the 2.4 GHz band. This band only offers three non-overlapping channels. Each channel is a separate ‘pipe' of bandwidth. All clients attached to the access point share that same bandwidth. Therefore, if you have 6 clients all attached to an 802.11g access point that achieves 24 Mbps effective throughput, at most they each have 4 Mbps. If you now place two more access points on different channels with contiguous coverage areas to the 1st one, and distribute the clients evenly among them, each 2 clients will share 24 Mbps of bandwidth, or have 12 Mbps each.

3) *Collisions and interference*
   As wireless clients are not physically connected together, they are unaware of when others are trying to transmit. As a result, it is likely that multiple clients in the same area will transmit at the same time. When this occurs, the clients are programmed to ‘back off.' That is, to wait some period of time before trying to transmit again. The waiting period increases with each collision; the idea being that the longer the interval, the higher the probability that the other client won't try to transmit at the same time. Back off causes the performance of the client to be reduced. If the client can normally transmit at 54 Mbps, but has to wait due to other collisions, then the effective transmission rate is reduced because of the time spent waiting.

4) *Handoff between access points*
   When a client roams from one access point to another, the time between disconnecting from the 1st access point and reconnecting to the 2nd access point is non-zero. This process involves the client searching for the next strongest signal by scanning the band. For some clients, this process can take up to several seconds.[1]

[1] "Its average handoff times ranged from about a half-second for one call, to just more than 1 second for the seven-calls-with-data scenario. While that kind of delay will be noticeable to callers, it was still by far the fastest roaming performance of any product." Network World, *Voice over Wireless LAN*, 1/10/05. A reference to the best handoff times they measured for the winning product.

Deploying for Density – Why MicroCells Exacerbate the Problem

With small Wi-Fi deployments, in a home, café or small office, most of the above mentioned problems are not noticeable. Either there are only one or two access points, or so few clients that the issues do not present themselves in a meaningful way.

All this changes with pervasive wireless LAN deployments across an enterprise campus. In this situation, you have tens to hundreds of access points and potentially thousands of clients. Universities and hospitals were some of the first enterprises to see the problems with a pervasive wireless LAN deployment, but as the desire for Wi-Fi connectivity grows, more and more enterprises will encounter the same challenges.

Microcells for Higher Density

With higher densities of clients, conventional wisdom says the enterprise will need to space access points closer together to achieve reasonable bandwidth per client. As an example, an 802.11b access point delivers an effective throughput of 6 Mbps. If for a given area, the density of clients requires more throughput, then shrinking the access point coverage area to put more cells in seems like a reasonable method to increase effective bandwidth. This technique is called *microcells*.

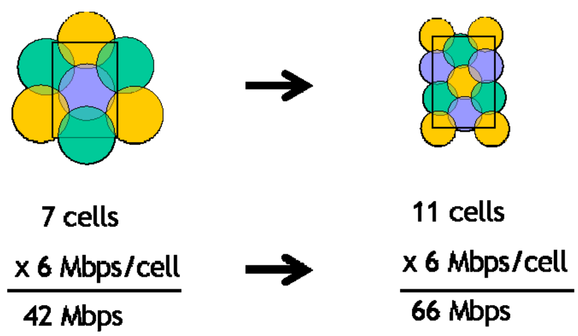

Figure 2: A microcell deployment increases the number of access points in a given area by reducing transmit power. It is widely recommended as a method to create more bandwidth for higher densities of clients.

Increased Interference and Collisions

Unfortunately, microcells increase interference and collisions. Cell size is reduced by decreasing transmit power. However, the access point signal will still travel beyond the 11 Mbps rate and as more access points are crowded into the same area, you end up with co-channel interference. That is, signals from different access points that are on the same channel will interfere with each other.

On the surface, decreasing the distance between access points seems like a simple answer to maintain high capacity for larger numbers of clients. However, access point RF propagation does not stop at the desired data rate. I.e. if a desired data rate of 24 Mbps is used to space access points to achieve a higher overall throughput, the access point signal will continue past the 24 Mbps radius through the 12, 6, 2 and 1 Mbps data rates (about 300 feet away).

This means that even careful planning to avoid adjacent APs having the same channel will not avoid the problem of co-channel interference. In fact, higher desired WLAN network bandwidths significantly aggravate this problem as more and more APs on the same channel interfere with each other. The result of this is a *decrease* in network capacity and performance. Because an 802.11 WLAN is a shared medium, the impact of co-channel interference is increased client collisions and corrupted packets (causing subsequent retransmissions) as the clients hear signals from the many APs and clients surrounding them. These increased client collisions and retransmissions cause more and more transmission delay effectively reducing throughput of the network. Today's generation of wireless LANs are unable to prevent this problem, consequently high density deployments are not able to maximize the performance of a 54 Mbps 802.11g network.

Hidden Nodes Increase Collisions

A further complication is termed 'hidden node.' Hidden node occurs when there are multiple clients within a single cell that are able to communicate with an access point, but cannot 'hear' each other. While the CSMA/CA protocol works well for those clients within range of each other, this will not always be the case. Clients may be separated suitably far from each other due to the Access Points range at low data rates, or separated by a physical obstacle.

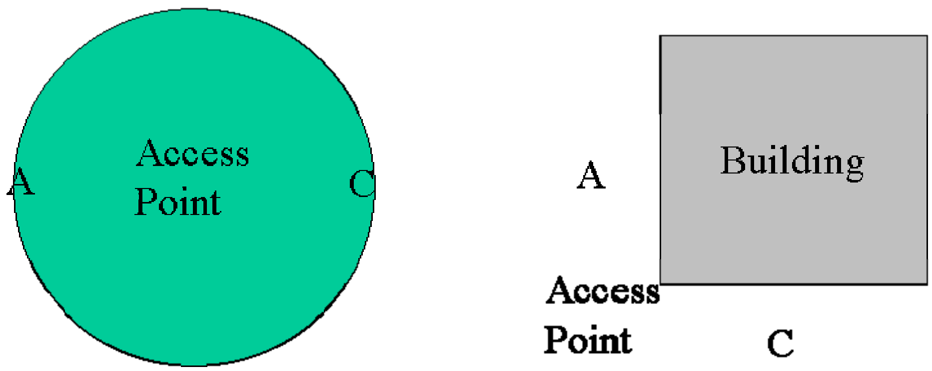

Figure 4. Increased collisions due to the hidden node problem reduces overall wireless LAN performance, and cannot be prevented by channel planning.

More Handoffs

An additional issue for latency sensitive applications such as voice over Wi-Fi is the significantly increased number of handoffs that a microcell deployment causes. And handoffs between access points cause delay which can destroy voice or video quality.

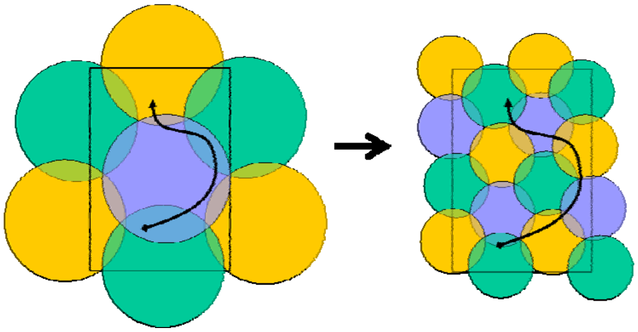

Figure 5: In the same physical space, the client on the left experiences one handoff vs. six handoffs in the same area using a microcell deployment. Multiple handoffs for voice clients can significantly degrade quality.

Increased Deployment Cost

Microcell deployments also significantly increase deployment cost in three ways.

1) Site survey - Due to the increased number of closely spaced access points, a much more careful site survey for the channel plan must be done to minimize co-channel interference.
2) Ethernet ports – Each Access Point requires an Ethernet port to bridge to the wired network. If there are no available ports, then new Ethernet switches must be installed.
3) Cable pulls – Vying with the site survey for expense, the cost of the cable pull can reach upwards of $500 per access point.

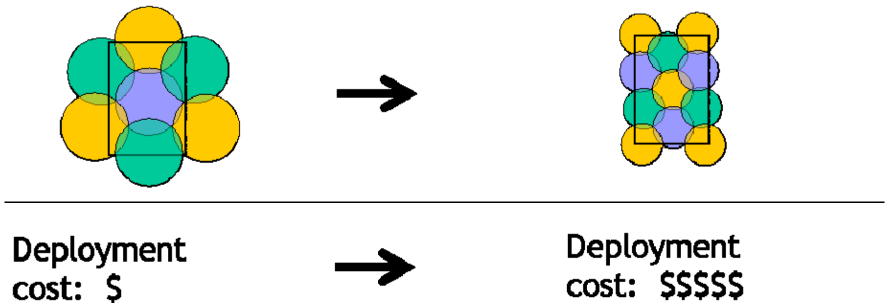

Figure 6: Microcells also significantly increase deployment cost and ongoing operational expense due to the significant increase in the number of access points.

The table below shows a typical deployment cost for a conventional wireless LAN deployment and one in the same area using a microcell topology to increase performance.

		Conventional Deployment	Microcell Deployment
# of APs		10	40
Cost per AP		$ 600	$ 600
Total AP Cost		$ 6,000	$ 24,000
WLAN Appliance/Switch Cost		$ 4,995	$ 22,995
Total WLAN Equipment Cost		$ 10,995	$ 46,995
Cable pull per AP		$ 500	$ 500
Layer 2/3 Switch cost per AP		$ 100	$ 100
Total Wired Infrastructure Cost		$ 6,000	$ 24,000
Site Survey (20 min per AP @ $100/hr)		$ 330	$ 1,320
Install cost per AP (20 min per AP @ $100/hr)		$ 330	$ 1,320
Total Installation Cost		$ 660	$ 2,640
Support cost per AP (20 min per AP per month @ $100/hr)		$ 3,960	$ 15,840
Total Support Cost		$ 3,960	$ 15,840
Total COO		$ 21,615	$ 89,475
% Cost Difference Using Microcells			314%

Table 1: A microcell deployment where the radius of the access point coverage is simply halved results in an increase in deployment costs of over 300% versus the conventional deployment.

Summary

Pervasive wireless LAN deployments require much larger, broader deployment of Wi-Fi access points than hot spots, cafes or isolated guest access in the enterprise. Microcell topologies that space access points closer together are recommended to ensure adequate throughput performance. However, four basic challenges including data rate reduction with range, limited non-overlapping channels, co-channel interference and collisions, and increased handoffs severely limit the effectiveness of microcell deployments. Costs for microcell deployments also increase dramatically not only to the additional equipment requirements, but the added time for complex site survey planning and ongoing maintenance. Clearly new solutions and deployment methodologies should be considered to make pervasive wireless LAN deployments feasible and economical.

