



WHITE PAPER



Enterprise Wireless LAN and the Broadband Myth

The Challenge of 802.11 WLAN Coverage and Capacity

May 2006

Copyright © 2006, Extricom, Ltd. All rights reserved. Extricom, Interference-Free, TrueReuse, UltraThin, and the Extricom logo are trademarks of Extricom Ltd.

No part of this document may be reproduced in any form or by any electronic or mechanical means, including information storage and retrieval devices or systems, without prior written permission from Extricom, Ltd.

C O N F I D E N T I A L—Use, reproduction, or disclosure is subject to the restrictions in DFARS 252.227-7013 & 252.211-7015/FAR 52.227-14 & 52.227-19 for commercial computer software or technical data provided to the U.S. government with limited rights, as applicable.

Table of Contents

INTRODUCTION	3
THE “TRADITIONAL” ARCHITECTURE – CELL PLANNING TOPOLOGY	4
THE CAPACITY MYTH	5
RANGE LIMITATIONS	5
EDGE USERS.....	6
LIMITED NUMBER OF USERS.....	7
REUSE RANGE	7
INTERFERENCE RANGE	8
COVERAGE / CAPACITY	8
CONCLUSION.....	8
THE EXTRICOM SOLUTION	9
THE ARCHITECTURE	10
RANGE EXTENDED, HASSLE ELIMINATED	11
ALL USERS AT THE HIGHEST DATA RATE	11
“STACKED” CAPACITY	11
TRUE REUSE - OPTIMUM CHANNEL REUSE	12
SUMMARY	13
REFERENCES	14

Introduction

Businesses have an ever-increasing desire to “cut the cord” for voice, data, and video communications. This is fueling demand for pervasive wireless LAN (WLAN) networks. Enterprises are therefore poised to move from tactical, “hot-spot” use of Wi-Fi, to strategic and ubiquitous wireless services.

But this also raises the stakes. Aware of the implications of increased reliance on WLAN, forward-looking enterprises are re-examining long-held assumptions about WLAN benchmarks of performance, deployment, and utility. After all, what may have worked in the tactical, hot-spot model will not hold up when subjected to the ratcheting stress of large deployments running advanced applications that require seamless mobility and real-time, robust wireless connections.

The bottom line is that, relative to the traditional switched WLAN, performance and resiliency of the wireless network must increase, while at the same time the complexity, cost, and risk of deployment and maintenance must be reduced.

Traditional WLAN systems employ a cell-based topology. Unfortunately, the inherent traits of the 802.11 “Wi-Fi” specification cause this topology to be limited in its ability to deliver on the three core values of wireless: capacity (and broadband), coverage, and mobility. This document is focused on the real-world challenges to achieving the system capacity and overall bandwidth required by a ubiquitous enterprise WLAN deployment.

The analysis yields an inescapable conclusion: for strategic, enterprise-wide WLAN implementations to succeed, a rethinking of the system topology is essential. Extricom's patented Interference-Free™ WLAN architecture introduces a novel approach that eliminates the coverage and capacity limitations of traditional WLAN architectures, delivers zero-latency seamless mobility and communications resiliency, and dramatically simplifies the planning, deployment, and maintenance of the network. The result is an ideal solution for delivering a wire-like user experience for converged data, voice (VoWLAN), and video services.

One final note: while this document uses 802.11b networks as the working examples, the same arguments apply to an even greater extent to 802.11a and 802.11g WLANs because of their shorter transmission range at the higher data rate of 54 Mbps.

The “Traditional” Architecture – Cell Planning Topology

In a cell planning topology, the available radio channels are distributed among the WLAN access point (AP) as shown in Figure 1 below. Figure 1 shows the 802.11b/g case, in which there are only three non-overlapping channels available. Each AP (represented by a hexagon) is assigned a specific radio channel, and then the APs are distributed to form a cellular coverage pattern. The designer must take care to provide sufficient physical separation between any two APs that use the same channel, so as to minimize the interference between them. This is the topology that underpins traditional data-centric WLAN systems.




Figure 1 – Cell Planning Topology

Unfortunately, this topology is complex to deploy, optimize and maintain, and its real-world performance is inadequate to support pervasive WLAN deployments. At the heart of the problem is a fundamental inability of cell-planning topologies to cope with:

- The severe scarcity of channels, which makes co-channel interference inevitable. With so few channels, collision domain sharing and insufficient capacity are the other resulting challenges.
- Latency-plagued mobility, as a result of frequent handoffs between APs on different channels.
- Sub-optimal support of multiple services – In the cell planning topology, all users, all device types, and all traffic types share (i.e. contend for) each channel. This presents quality of service (QoS) challenges that are difficult to surmount, even with the new 802.11e standard.

This document focuses on the first item in the above list, namely the real obstacles to maximizing capacity and bandwidth in WLAN.

The Capacity Myth

“The Capacity Myth: it’s a common notion that if channels are reused, a traditional WLAN has virtually unlimited capacity.”

The common notion in cell-based WLAN is that if channels are reused, the WLAN has virtually unlimited capacity.

The theory goes like this: each Access Point (AP) in an 802.11b WLAN supports an 11 Mbps data rate; since there are three channels, cells can be laid out throughout the enterprise, with each cell surrounded by two cells on different channels. Co-channel cells are thus separated, and each cell can support users at maximum speed. The aggregate bandwidth of such a layout is, according to the common notion, virtually unlimited. Each cell provides the full data rate, multiplied by the number of cells in the cell layout. Figure 2 shows a typical cell layout.

In practice, a number of factors make the common notion completely unrealistic. Among the factors that contribute to much lower capacities in real-world settings are:

- Range Limitations
- Edge Users
- Limited Number of Users
- Reuse Range
- Interference Range
- Coverage / Capacity

Range Limitations

“Due to rate adaptation and ‘edge-users’, using large cells is devastating to the capacity of the deployment.”

An AP has a reach of 100 feet at 11 Mbps and 200 feet at 2 Mbps in a real-world enterprise with walls, doors, people, desks, chairs, filing cabinets, computers, and other RF interference. At 54 Mbps, the reach limitation is even more severe at 25 feet¹. To provide adequate coverage, APs are laid out in a tiled pattern of coverage cells, with neighboring cells using different channels (see Figure 2 below). The dark inner circles indicate the coverage area at 11 Mbps; the outer circle, the coverage area at 2 Mbps.

In real-world installations, cell size is always much larger than the area covered by the maximum data rate. This is necessary for adequate coverage. Furthermore, placing cells close together would be ineffective (see “Reuse Range”) and even counterproductive (see “Interference Range”). More cells increase the number of inter-AP handoffs to mobile users, causing delays and interruptions in service². Moreover, developing a cell plan for a large number of cells is very labor-intensive and expensive, especially since the cell plan must be constantly updated to cope with the dynamic nature of the radio-frequency environment.

On the other hand, using larger cells is devastating to the capacity of the installation. When a node (client or AP) receives another node poorly (because the client and AP are far away from one another), it uses rate adaptation, or “gearshift” to a lower data rate to improve reception. The area covered by lower data rates comprises $\frac{3}{4}$ of the cell. This means that most users will use the lower data rates. Not only does this decrease the aggregate bandwidth of the cell, but it also affects the experience of users connecting at high data rates (see “Edge Users”).

Result: aggregate bandwidth will be dramatically less than the theoretical bandwidth at the maximum data rate.




Figure 2: Cell Planning Using 3 Channels

Edge Users

“The edge user will always create a bottleneck that denies everyone a high-speed experience.”

Traditional cell-planning schemes invariably create edge users, that is, users that are too far from the nearest AP to connect at the higher data rates. Not only do edge users experience “gearshift” or rate adaptation, they also affect the experience of every other user in the Collision Domain.

If there is a single edge user using the AP who cannot achieve 11 Mbps access (because of distance or poor signal reception), all users will experience much slower overall performance. Even users who connect at the maximum data rate on a per-packet basis will experience much lower overall performance, because they must always wait for the edge user to relinquish the air time. Since the edge user’s transmission takes much longer than anyone else’s, the edge user will always create a bottleneck that denies everyone a high-speed experience³.

When both 802.11b and 802.11g technologies are used in the same environment, there is an inherent edge user problem: 11b supports slower data rates than 11g. Even though the problem is partly solved by the 802.11g standard itself, the protection mechanism greatly decreases capacity when the two technologies are used together⁴.

Result: Cell planning does not increase capacity. The slowest user determines the aggregated data rate.

Limited Number of Users

“As the number of users increases, the chances of collision increase rapidly, causing considerable deterioration of throughput.”

When the number of users in the enterprise WLAN becomes large (typically more than 60), the chances of collisions increase rapidly. Collision avoidance schemes make sure that a node first checks to confirm that no one within its listening range is transmitting in its channel. Before starting to transmit, a node listens on the channel for $50\mu\text{s}$ to ascertain that the channel is clear. The node then backs off for a random amount of time (random number $0-31 \times 20\mu\text{s}$; average wait = $310\mu\text{s}$) before transmitting. The random back-off decreases the likelihood that two nodes will start transmitting simultaneously. But the process of assuring that the channel is clear for transmission takes an average of $50\mu\text{s} + 310\mu\text{s} = 360\mu\text{s}$, which is a substantial time compared to the time it takes to transmit a data packet. If, after backing off, the transmission is still not received clearly, the random back-off time will be doubled (random number $0-63 \times 20\mu\text{s}$; average wait = $630\mu\text{s}$) before transmitting again. If that doesn't help, it is doubled again. This process can continue until random number = $0-1023 \times 20\mu\text{s}$; average wait of more than 10ms, which essentially means that communications has come to a standstill.

Result: *As the number of users increases, the chances of collision increase rapidly; double and even quadruple back-off can become frequent, causing considerable deterioration of throughput.*

Reuse Range

To increase capacity on the wireless network, frequencies are reused. Reusing a frequency requires careful cell planning. In practice, frequency reuse is seldom achieved, because unless cells are very far apart, they share collision domains. All nodes in the same collision domain share the same bandwidth. Placing two cells on the same channel within the same collision domain provides each with half of the bandwidth.

In accordance with the 802.11 standard, Clear Channel Assessment (CCA) must be used before a node transmits to ensure that the frequency is clear for transmission. The Physical Layer Convergence Protocol (PLCP) also limits the reuse range, but in practice, CCA will almost always be the determining factor of reuse range.

CCA measures the amount of energy in the channel without regard to packets, transmission speed, or even source of energy (whether it is an AP or a microwave oven). Since energy can be detected over hundreds of feet, CCA may place far too many cells in the same collision domain. In Figure 2, the blue cells and the red cells are very likely to be in the same collision domain. In other words, it is actually CCA reach that will determine the collision domain size, and not the cell size at the highest data rate.

Moreover, the CCA mechanism is notoriously unreliable, and may provide both false positives and false negatives. The CCA may indicate that the channel is clear when, in fact, it is not (false negative). This will result in co-channel interference. On the other hand, the CCA may indicate that the channel is not clear when, in fact, it could be reused.

Result: *Cells may be shared incorrectly; or worse, reuse attempted, resulting in interference, as described in “Interference Range”.*

“In practice, frequency reuse is seldom achieved because cells share collision domains.”

Interference Range

“Placing cells closer together leads to co-channel interference. Cell planning does not provide effective reuse in enterprise settings.”

Our laboratory tests show that the carrier-to-interference ratio for 802.11b CCK modulation must be at least 8dB, which means that the carrier must be 2.5 times closer than any co-channel interference to be able to read the data packet; that is, the client must be 2.5 times closer to one AP than to the next one on the same channel to be able to reuse the frequency channel without interference.

The cell must be much larger than the area that can be reached at the highest data rate. In order to prevent collisions, the next nearest AP must be at least $200 \text{ feet} \times 2.5 = 500 \text{ feet}$ from the client to prevent co-channel interference; therefore, the APs must be at least 700 feet apart. Practically speaking, to reuse a frequency three times, the domain must be approximately 1,200 feet across (reused cells will typically not be in a straight line). Figure 2 shows a typical cell plan that is 1,200 feet across. The red cells may be expected not to interfere with each other, allowing frequency reuse. Placing cells closer together leads to co-channel interference, which causes retransmission and rate adaptation, lowering capacity substantially.

Result: *Cell planning does not provide effective reuse in enterprise settings. In practice, it results in interference and reduced capacity.*

Coverage / Capacity

Examining Figure 2 above, the blue cells are too close to other cells to allow reuse. In other words, a space of 1,200 feet is required to allow the same channels to be reused a mere three times. Even then, reuse is not available except between cells on the outer edge of the site, and cells in the middle may hinder even that reuse.

Result: *A large layout sacrifices capacity for coverage. Even with a large layout, it may not be possible to reuse cells no matter how far apart they are, if interior cells interfere.*

Conclusion

“Users will seldom, if ever, be able to achieve much over a 2 Mbps experience on a traditional real-world enterprise WLAN.”

We have seen that range limitations, edge users, collision domains, effective reuse diameter, interference from intervening collision domains, and multiple users make 11 Mbps WLAN connections for 802.11b a theory rather than a reality. Users will seldom if ever be able to achieve much over a 2 Mbps experience on a traditional real-world enterprise WLAN.

The Extricom Solution

“Extricom allows APs to be spaced closely together providing highest-data-rate coverage throughout the enterprise.”

Extricom’s patented Interference-Free architecture is an innovative WLAN architecture that makes it possible to space APs close to each other, allowing quality reception/transmission and maximized transmission speeds (everyone can be close to an AP) without black holes or areas of poor or no coverage. This unique architecture completely avoids downlink contention so that performance is not affected and co-channel interference is completely avoided.

At the heart of Extricom’s innovation is a simple, yet powerful idea: eliminate the concept of cell-planning, and replace it with the “channel blanket” topology. The solution allows each radio channel to be used everywhere, on every access point, to create blankets of coverage. This approach is illustrated in Figure 3.

Within each channel blanket there is seamless mobility with no roaming latency, no co-channel interference, robust client connections that simply do not drop, and the ability to design for a guaranteed and predictable level of service for all users.

Also, using multi-radio, 802.11a/b/g UltraThin™ Access Points enables overlapping channel blankets to be powered by the same set of APs connected to a single wireless Switch.




Figure 3: Extricom Interference-Free Architecture

The Architecture

In the Extricom infrastructure, the central switch makes all of the decisions for packet delivery on the wireless network. The switch directs all of the traffic because the access points have no capabilities of their own... no software, no storage, no smarts, just radios. The clients don't even associate with the access point. Instead, the access point rapidly funnels the traffic back to the switch for processing.

As a result, a given 802.11a/b/g channel can be deployed at *every* access point, to deliver complete coverage, consistent capacity, and zero-latency roaming, without co-channel interference. And since all access points collaborate to provide service to the client, the user experiences a stable, wire-like connection that is highly resistant to the outside interference and RF signal variations.




Figure 4: Extricom Architecture

Thanks to this architecture, the system is dramatically simpler to plan, deploy, and maintain than traditional WLAN. That's because RF cell planning or specialized RF knowledge is not required.

In addition, the UltraThin APs enable true 'plug-and-play' deployment. Since it carries no software, the AP requires no configuration and is completely interchangeable with any other AP. Whenever changes are made to the system, there is no need to reconfigure, reboot, or otherwise maintain the AP. It is also never a point of possible security breach, since all security and configuration is performed centrally in the Wireless Switch. In terms of physical plant, no modifications are required either in Category 5 cabling or on the client side, which may use any off-the-shelf wireless Network Interface Card (NIC).

Range Extended, Hassle Eliminated

“With an Extricom WLAN, there is no trade-off between bandwidth and coverage; you will have both.”

The Interference-Free architecture allows many APs to be spaced as close together as required so that clients will always be within the maximum data rate range of an AP. High bandwidth coverage is complete and ubiquitous. If you discover an area that is not adequately covered, just add another AP. Since the Interference-Free architecture avoids co-channel interference and collisions, no RF engineering and cell-planning is required, so adding another AP will not interfere with an existing setup. Furthermore, since UltraThin APs contain no software, there is never anything to configure in the AP.

With an Extricom WLAN, there is no tradeoff between bandwidth and coverage: you will have both.

Result: Extricom provides highest-data-rate coverage throughout the enterprise, while eliminating RF site surveys, AP configuration, and “black holes”.

All Users at the Highest Data Rate

Since Extricom covers the entire enterprise with a blanket of closely spaced APs, there are no range limitations, and therefore no edge users. Every AP is close to another, allowing all APs to transmit at the maximum data rate. Even in mixed-mode environments (802.11b and 802.11g), Extricom allows clients with lower data rates to be placed on a different channel, thus eliminating the problem of the slower edge user. Consequently, no protection mechanisms are needed.

Result: Extricom eliminates the edge user problem. Since everybody connects at the maximum data rate, no one user slows down any other; all users experience maximum throughput.

“Stacked” Capacity

Since Extricom covers the entire enterprise with overlapping blankets on independent channels, it is possible to multiply the capacity at every point in the enterprise. In the case of the 2.4GHz band, this implies that local capacity can be tripled, by using all three channels on separate but overlapping blankets. While traditional cell planning uses three channels to provide coverage, Extricom uses them to provide capacity.

TrueReuse - Optimum Channel Reuse

“TrueReuse is able to overcome the limitations of the CCA mechanism, to provide frequency reuse with high spatial density.”

In addition to three channels everywhere in the enterprise providing full coverage at the maximum data rate, Extricom provides true frequency reuse with patented TrueReuse™ technology.

The idea behind TrueReuse is to take each channel blanket and multiply its aggregate capacity. While the blanket provides the coverage, mobility, and link resilience benefits already described, TrueReuse introduces a way to boost the aggregate capacity of that blanket.

TrueReuse essentially takes the concept of frequency reuse and actually delivers the benefit that had long been hoped-for in cell planning topologies. To increase the capacity of the channel, the Extricom switch uses its real-time knowledge of the entire system to decide when to permit multiple access points to simultaneously transmit on the same channel, to different clients, without causing co-channel interference. In essence, TrueReuse takes a single collision domain and dynamically splits it into sub-collision domains, thereby multiplying aggregate capacity.

How? The system dynamically measures the RF reception quality from each client on a packet-by-packet basis. These measurements are used to create a high granularity, real-time map of co-channel interference throughout the deployment. This map is then used to determine when simultaneous transmissions on the same channel will not cause co-channel interference. The switch uses its real-time information to overcome the limitations imposed by the CCA mechanism, and to provide efficient frequency reuse.

“With an Extricom WLAN, up to nine times the maximum aggregate data rate can be achieved in the 2.4 GHz band.”

The net result is a multiplication of the aggregate capacity of the channel blanket. The Extricom system will deliver three times to nine times more system bandwidth than a comparable cell-planning system. What’s more, this bandwidth multiplication is in addition to that which is achieved through the “stacking” effect of overlapping channel blankets.

The combined use of TrueReuse technology, and Extricom’s ubiquitous AP placement on the same channel, allows the Interference-Free architecture to greatly reduce the AP-to-AP distance necessary for effective reuse. By making decisions on a packet-by-packet basis, Extricom avoids co-channel interference and uses RF resources to the maximum.

Result: Extricom provides optimal channel reuse and very significantly increases capacity.

Summary

Broader adoption of Wi-Fi in the enterprise requires raising the bar on the performance and flexibility of WLAN, while concurrently cutting complexity, cost, and risk of ownership. The traditional approach to WLAN systems, based on a cell topology, is inherently limited. At the heart of the issue is the manner in which that infrastructure uses the few available radio channels, leading to a system that is suboptimal in terms of capacity, bandwidth, mobility, and resilience.

Extricom's patented Interference-Free WLAN architecture provides a true high-capacity wireless network. With this architecture, up to nine times the maximum aggregate data rate can be achieved. In the 5GHz band, this means reaching a maximum aggregate data rate in excess of 1 Gbps with 802.11a clients. Blanket coverage eliminates black holes, areas of poor reception, slow performance, range limitations, edge users, and limitations imposed by the Collision Domain range. TrueReuse replaces performance degradation caused by collisions with the ability to reuse the same frequency with high spatial density. The result is a wireless network that provides wire-like performance across the entire enterprise.

References

1. **802.11 Wireless LAN Performance White Paper**, Atheros Communications 2003
2. **Voice over WLAN**, Extricom 2004
3. **Performance Anomaly of 802.11b**, Martin Heusse, Franck Rousseau, Giles Berger-Sabbatel, Andrzej Duda, 2003
4. **IEEE 802.11g Network Behavior in a Mixed Environment**, Jim Zyren, Tim Godfrey, Menzo Wentink, Intersil Corporation 2003