White Paper |

Enterprise

WLAN RF Architecture Primer:

Single-Channel and Adaptive Multi-Channel Models

Peter Thornycroft

ARUBA® networks

Single-channel WLAN architectures

Markets thrive on competition, where vendors are forced to innovate in order to differentiate, and customers pick winning products on the basis of performance and price. The WLAN market is currently blessed with a group of competitors differentiating anew with a broad range of technologies, including distributed switching architectures, 802.11n implementations and different models for of fixed-mobile convergence, to name a few.

In this paper we examine one of the questions central to the technology debate: the advantages and disadvantages of the single-channel RF architecture when compared to the adaptive, cellular-like alternative. The RF model is a significant aspect of WLAN architectures, and as we shall show, the initial selection of a single-channel model drives significant consequences, both good and bad. This article compares the two models in technical detail, allowing the reader to follow the science behind marketing claims, and develop an informed and objective opinion of the merits of vendors' solutions.

We first compare the three key challenges of WLAN architectures: handover and channel planning - which are somewhat improved in the single-channel architecture - and co-channel interference, which is greatly exacerbated. This mostly technical discussion, will demystify marketing claims that confuse performance effects due to the underlying architecture with unrelated features that make good sound-bites. We aim to clearly identify the strengths and weaknesses of each architecture.

Following this, the paper investigates the advantages and disadvantages of each model when building enterprise WLANs with respect to robustness, flexibility, scale, capacity, cost and performance. This section will introduce practical design considerations, i.e., the difficulties of building a functioning ecosystem when clients must be tuned to perform on a small number of deployed networks.

Our intent is to provide enough technical information for the reader to separate exaggerated marketing claims, driven by the fierce but healthy competitive environment, from the actual technical strengths and weaknesses of these architectures.

The adaptive model as a point of departure

We begin with a brief explanation of channel planning and handover in the adaptive RF model. The objective in this section is to explain the motivation for seeking a new and better way to accomplish mobility and planning, and thereby provide a backdrop for a deeper investigation of the single-channel model. The Wi-Fi networks we are considering are similar to the public cellular telephone networks in that they are formed of many radio transmitters, each covering a relatively small area, and provide coverage over large areas, numbering from a few to thousands of transmitters in a network.

Adaptive networks deploy access points using a different RF channel or frequency for each transmitter, in an arrangement similar to the cellular telephone networks. To build up an area of continuous coverage, access points are placed at intervals, with each providing coverage in its area, or cell, on a given RF channel. The use of different RF channels prevents interference in areas where cells overlap: if a client finds itself in an area where two cells using the same RF channel overlap, it will experience difficulty decoding the wanted signal in the presence of transmissions from the overlapping cell. This overlapping condition is avoided in the adaptive model by moving access points on the same channel physically as far apart as possible. Since

RF signals weaken over distance it is much easier to decode the wanted signal if the undesirable interfering signal is at a lower power level. One of the central design challenges in the adaptive architecture is designing coverage patterns to devoid of any same-channel overlap. The challenge arises because of scarcity of spectrum: there is usually a restricted number of available RF channels. For Wi-Fi there are only 3 non-overlapping 20 MHz channels available in the 2.4 GHz band, but more than twenty 20 MHz channels in the 5 GHz bands, depending on the technology used. This highlights an important distinction: the issues of channel planning are seriously challenging at 2.4 GHz, but not significant for networks operating in the 5 GHz band.

Adaptive RF model: channel re-use in non-adjacent cells

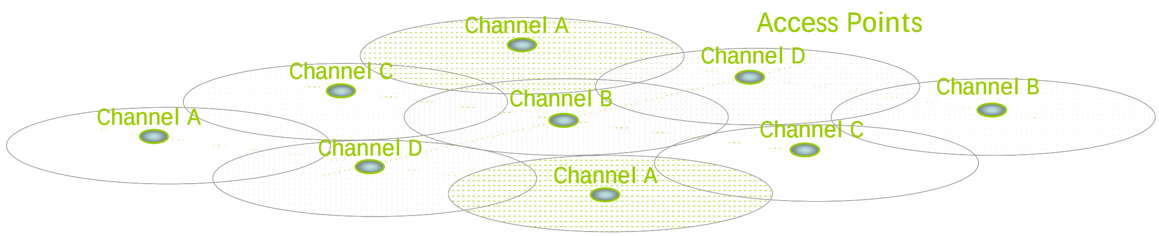

Since access point spacing is usually dictated by the required network capacity, and a limited choice of access point mounting locations is typically available, network planning engineers use transmit power to influence the size of each cell, and as much as possible identify the best RF channel re-use pattern across the network to avoid areas where same-channel cells overlap.

Adaptive RF model: channel assignment

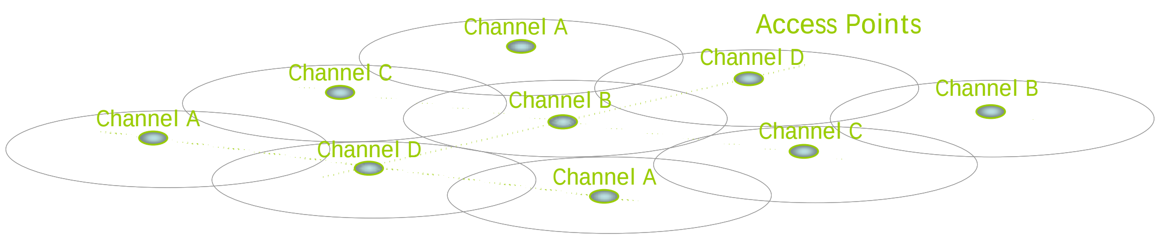

After continuous coverage, the second major requirement for enterprise Wi-Fi networks is mobility, that is, the process of handing over clients from access point to access point as they move from cell to cell. In the adaptive model, clients themselves choose to associate to a particular access point by selecting the appropriate RF channel and tuning out other access points transmitting on other RF channels. Transferring to a new cell is accomplished by the client switching its radio to work on the new access point's RF channel. The handover between access points is initiated by the client, and it is one of the more challenging functions of the adaptive architecture. First the client must decide it's time to handover, then select the target access point and switch to the new RF channel, and then re-authenticate at that access point. Each of these phases of handover is difficult to accomplish quickly, accurately and consistently, so standards bodies and vendors have devoted considerable attention to handover in recent years.

Channel planning and handover present significant complexity, so it is perhaps natural that engineers have explored ways to simplify them. The single-channel model is represented as a solution to both challenges.

Adaptive RF model: inter-AP handover

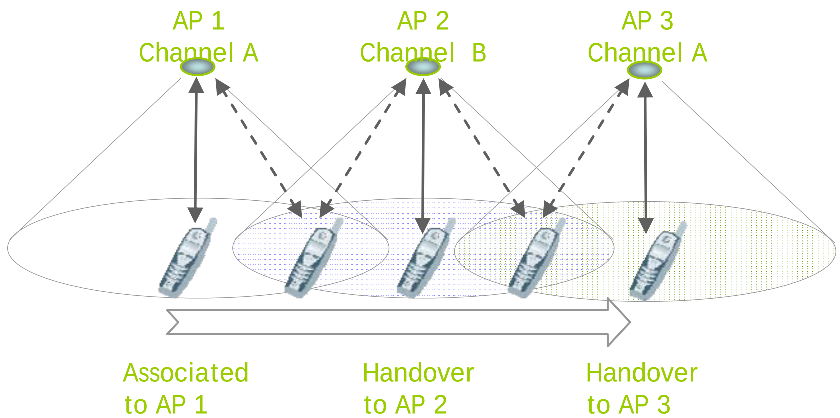

Note that the issues of channel planning and handover affect multi-access point networks, and are not applicable to single, isolated access points such as home Wi-Fi networks. The discussion in this paper is pertinent to enterprise Wi-Fi networks in which many access points are used to provide areas of continuous coverage.

The single-channel model: a solution to handover and RF planning

Single-channel networks use access points that are all tuned to the same RF channel or frequency. The simplest view of this model shows a number of access points with overlapping coverage forming a continuous region. Most implementations are designed so that the clients cannot distinguish between the access points providing coverage: instead the network decides which access point should transmit and receive data for a particular client. In other words the client is not involved in any handover decision. As clients move through a building, the network directs traffic to them via the nearest access point with available capacity.

Single-channel RF model: channel assignment

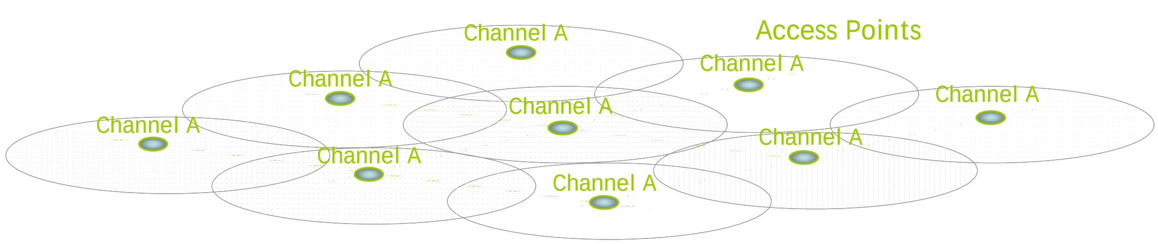

The main attraction of the single-channel RF architecture is that it eliminates the need for a handover mechanism involving the client. As we shall see below, while the core claim of simpler handover at the client is supported, moving the handover decision from the client to the network infrastructure comes with a new set of challenges.

A second attraction is that the traditional channel planning required for cellular networks should be unnecessary, because all access points are set to a common RF channel. The only decision to be made is to choose the best channel for the entire network. As we shall see, while RF channel planning is effectively eliminated, the architecture brings with it new challenges associated with planning and controlling cell overlap and selecting the RF channel used for the network.

A deeper investigation of handover

This section examines the simplification in network planning that arises from avoiding client-based handover. We conclude that while clients' behavior can indeed be simplified in this architecture, gains are offset by the added complexity and performance issues associated with implementing a single-channel RF architecture

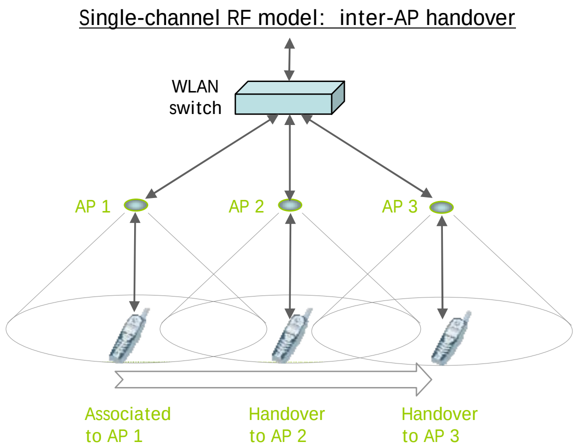

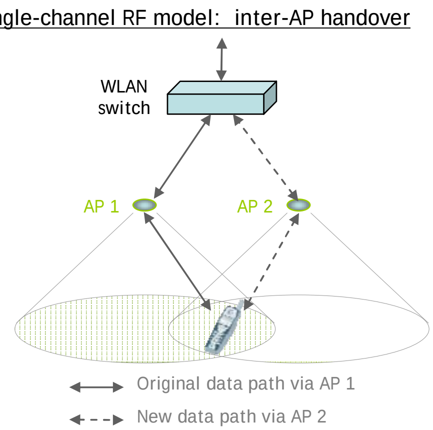

The diagrams above show handovers in a single-channel architecture. The key to these client-unaware handovers is that the access points are modified from the model defined in the IEEE 802.11 standards so that roaming clients can be tricked into believing they are always interacting with the same access point when they are in fact communicating with many different access points. This is accomplished by having all of the access points transmit the same BSSID or MAC address.

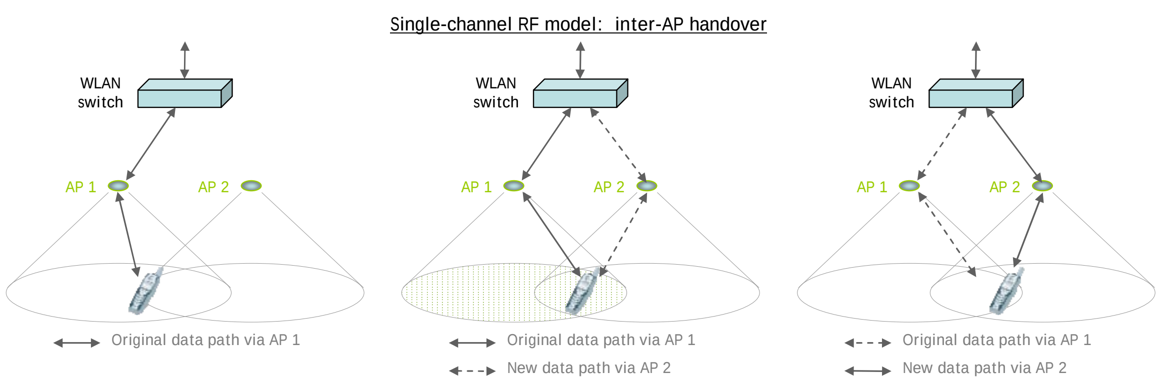

As the serving access point changes during a handover, the client must receive an uninterrupted flow of frames, meaning that header data and sequence numbers must increment smoothly as the serving access point changes.

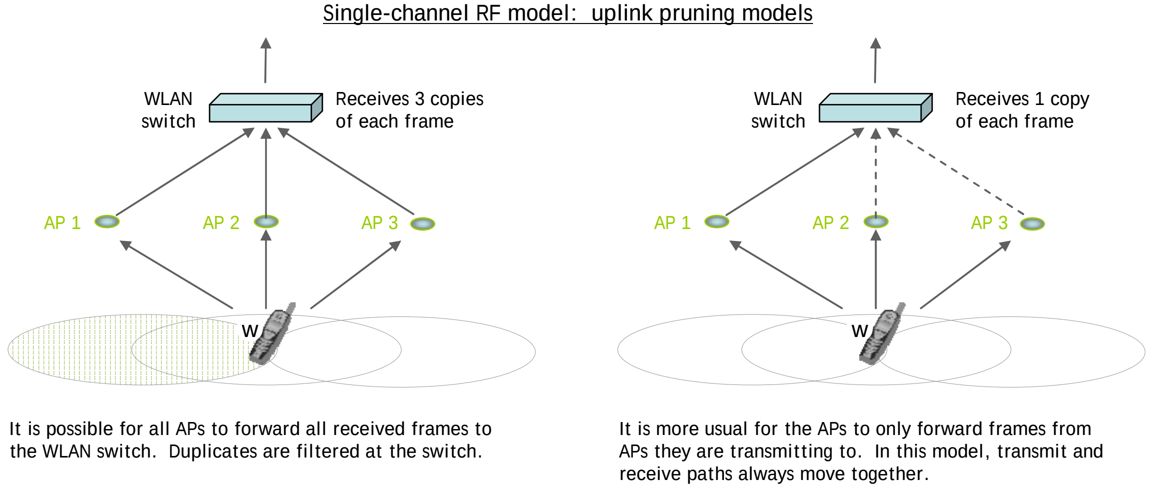

As shown in the diagram above, this is generally implemented by coordination between the WLAN switch and access points. As a client moves into an area covered by a new access point, the WLAN switch must sense the presence of the client and re-direct the flow of frames from the old access point to the new one. It is important that no gaps occur in the stream of frames during re-direction or data will be lost. Conversely, simultaneous transmission from different access points to the same client must be avoided, too – under the IEEE 802.11 standard, frames with duplicate sequence numbers indicate possible wireless intrusion attacks and result in data discards.

Managing the uplink direction is simpler. Each client transmission will be ‘heard’ by many neighboring access points, as they are all receiving on the client’s RF channel, so many duplicate copies of each transmission will be received across the network. The WLAN vendor must either prune the duplicate uplink frames at the access point, or forward all of them to the WLAN switch which can then apply a duplicate frame filter. The former requires close coordination between the WLAN switch and access points to ensure that one and only one access point is designated to receive for each client, the obvious technique would be to bind transmitting (downlink) and receiving (uplink) traffic for each client to a single access point, switching to a new access point for handover. The latter approach would provide better performance, as frames received in error at one access point would likely be successful at another, but has the disadvantage of significantly increasing overall traffic between access points and the WLAN switch, thereby also placing additional processing load on the switch.

We conclude that as the client moves from cell to cell, the WLAN switch must cleanly switch its data flow from access point to access point without any interruption or duplication of frames. It can certainly be argued that this is a lesser challenge than client-initiated handover in the adaptive model, but deploying a single-channel network requires considerable attention to switch processing capability and engineering coverage overlap between access points and across WLAN switches.

The optimum architecture for handover

While the adaptive model expects the client to make all decisions related to handover, the single-channel model moves decision making to the network side. One could argue that the former is a distributed approach, and hence more scalable. But there is certainly more information available on the network side, in terms of more complete knowledge of conditions at neighboring access points and even historical information about how clients reacted previously when in similar situations. This information is not available at the client.

Most experts agree that the optimum decision-making model involves a mix of information from the infrastructure and the client, and the IEEE 802.11 standards are moving in this direction. For example, 802.11k allows clients to report local RF conditions to the access point. There is also a provision for the access point to provide the client with a list of neighboring access points, information that would be more accurate and more timely than that available to the client from probes and scans of adjacent cells, as is the case today. 802.11v will also include a ‘directed handover’ feature whereby an access point can indicate to the client that it is time to handover to an adjacent cell. A number of similar functions will be part of the forthcoming Wi-Fi Alliance ‘voice-over-IP for enterprise’ certification, prompting phone vendors to design these capabilities into off-the-shelf Wi-Fi phones, and providing a migration path for infrastructure vendors using the adaptive architecture.

If the cellular industry is any guide, the decision-making model will evolve towards the adaptive model. The centralized handover control of first-generation cellular networks – which mimics the model used in single-channel WLAN architectures – proved unscalable. More capable and modern 2G and 3G architectures instead use a model based on shared network-client responsibility.

Co-channel Interference

Co-channel interference was touched on earlier in this paper when it was explained that overlapping cells on the same RF channel result in unwanted interference. To understand the implications of the single-channel model, we must look a little deeper. Co-channel interference is a phenomenon where transmissions from one cell spread to a nearby cell on the same RF channel, causing errors or dropped transmissions due to interference when they coincide with transmissions to or from devices in that cell.

Coverage with co-channel interference zones

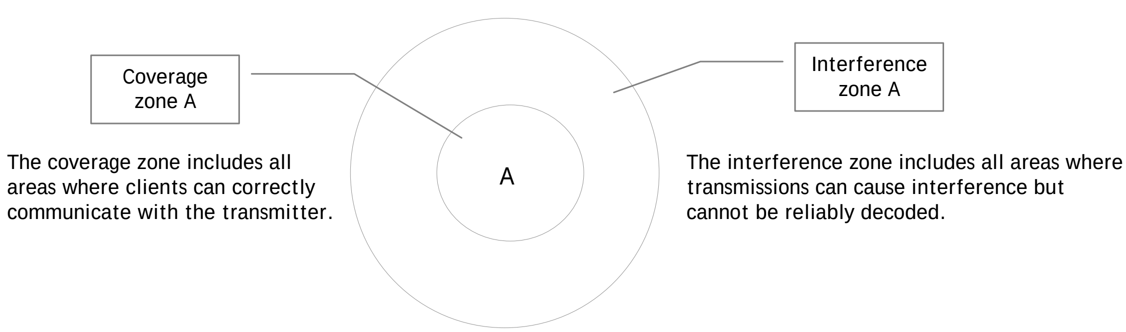

In practice, zones are not circular. They are 3-dimensional, and depend on RF conditions, especially indoors, and antenna patterns, among other factors.

Interference zones of clients and access points

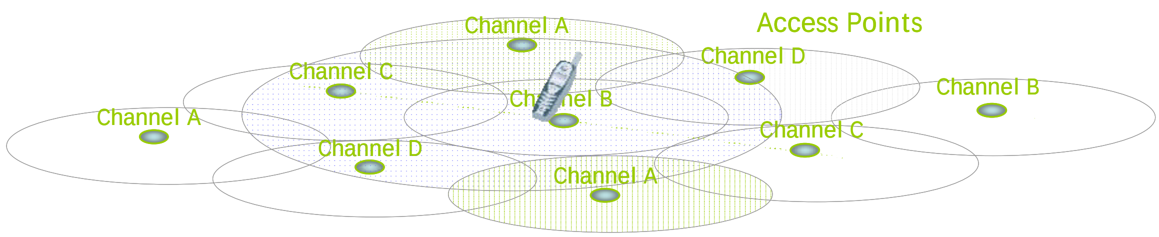

We define the 'interference zone' as the area around a client or access point where its transmissions will be powerful enough to affect devices' ability to decode other transmissions on the same RF channel. In practice the edge of the interference zone is not as clear-cut as shown in the diagram above: much depends on local conditions and the strength of the competing RF signal. Regardless of the simplified diagrams, the concept of the interference zone is central to any analysis of network performance issues.

In fact, client-originated is often the dominant form of co-channel interference in enterprise deployments today because clients may transmit at a higher power than the access points. This will become less important over time as clients implementing transmit power control (TPC) become more widespread (see the explanation of TPC below for details).

In the adaptive model, co-channel interference is limited because by design adjacent cells are nearly always on different channels. When fewer channels are available, access points on the same channel must be spaced closer, and co-channel interference becomes significant. While directional antennas and other techniques can be used to minimize the effect, co-channel interference limits transmission capacity and quality in many of today's Wi-Fi networks.

The mechanism whereby transmissions are inhibited in Wi-Fi is called carrier sense multiple access with collision avoidance for the media access control layer (CSMA/CA MAC). When a client or access point is about to transmit, it first checks the medium to see if any other transmissions are on the air. Even when such transmissions are detected from a distant cell on the same RF channel, if they can be successfully decoded they will inhibit the pending transmission. Similarly, if a device begins to transmit and senses another simultaneous frame on the air, it will cease transmission and enter a backoff mode. The effect is that transmissions in nearby cells on the same RF channel can prevent local devices from transmitting, even when the intended recipient of the transmission might not be in an area of interference.

Similarly, a device's receiver is tuned to its local cell's RF channel. When a transmission on the same channel from a neighboring cell is coincident in time, the signals will interfere. If the unwanted signal has comparable signal strength relative to the wanted frame, errors will cause a discard, and the lack of an acknowledgement will cause the sender to re-transmit the frame.

Mitigating co-channel interference

Two techniques are used to mitigate co-channel interference: spatial separation, using distance and multi-channel plans, and temporal separation using time-based coordination to avoid simultaneous transmissions.

Spatial separation is effective because the greater the distance between the devices causing and suffering interference, the lower the level of the unwanted received signal. Eventually the interfering signal is reduced to such a low level that it is no longer powerful enough to disrupt the wanted transmissions. Frequency reuse in a cellular model is simply a way to separate cells on the same channel by the greatest possible distance.

If spatial techniques fail, the second option is to control the transmission opportunities of the various clients and access points in the interference zone, ensuring that they do not transmit simultaneously. The Wi-Fi CSMA/CA MAC algorithm, enhanced for QoS support with WMM (wireless multimedia, based on the IEEE 802.11e standard) coordinates transmission opportunities within a cell. There is no mechanism in Wi-Fi, beyond CSMA/CA, to coordinate transmission opportunities across cells: cells operate independently and asynchronously. Spatial and temporal separation in CSMA/CA together effectively control co-channel interference in the adaptive architecture.

Since the single-channel architecture cannot use spatial separation, it must solve co-channel interference with a stronger temporal coordination mechanism. Such a mechanism was not contemplated in the IEEE 802.11 standards, and constitutes a proprietary extension that transgresses the spirit and intent, if not the letter of the IEEE 802.11 standard. Since clients are designed to meet the IEEE 802.11 standard, they don't necessarily accommodate such proprietary extensions. As we will see below, wayward client behavior is a not uncommon negative consequence of implementing a stronger temporal coordination mechanism.

Although we usually consider the situation where one access point's signal bleeds into another access point's cell, clients are also responsible for much co-channel interference. If a client is transmitting at a power level equivalent to the access point, a client at the edge of its cell may be closer to a neighbouring cell on the same channel, and thus more likely to cause co-channel interference. However, since access points usually transmit more often than clients, they may have a greater overall effect.

Co-channel interference (AP to client, AP to AP)

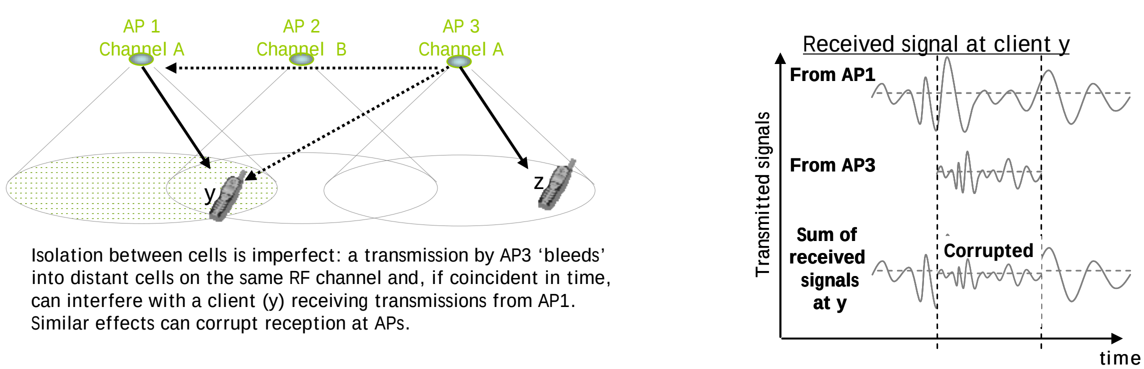

Isolation between cells is imperfect: a transmission by AP3 'bleeds' into distant cells on the same RF channel and, if coincident in time, can interfere with a client (y) receiving transmissions from AP1. Similar effects can corrupt reception at APs.

The diagram above shows how simultaneous transmission from APs that are well-spaced, but use the same RF channel, can result in interference at the client and loss of throughput in the network as corrupted frames must be re-transmitted.

More mechanisms for co-channel interference (client-originated, transmission inhibited)

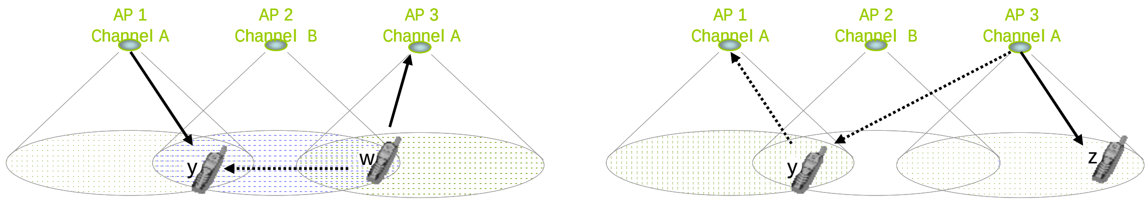

A similar effect, but client-originated. A transmission from w intended for its AP 3 bleeds over to y with sufficient power to corrupt its reception of a frame from AP1

Reception of a frame from a nearby cell on the same RF channel can also inhibit transmission, due to CSMA/CA. Here, client y wants to transmit to AP1, but must wait for AP 3's transmission to finish.

The diagram above shows how co-channel interference can be caused by clients as well as access points. A client's transmissions bleed through to another cell (on the same channel) and interfere with transmissions in that cell.

In these examples we have not distinguished between interference, hidden node and contention effects when transmissions collide since the effect is the same - retries reduce network throughput. If we investigate this phenomenon in more detail, it is clear that interference zones are graduated by signal strength. A client can by definition decode 802.11 frames from its access point. As the client moves farther away from the access point, it may be able to decode some transmissions encoded at low data rates (such as beacons and multicast) potentially even when it cannot join the cell because it cannot sustain the minimum data rate, e.g., the minimum rate is set to 12 Mbps where the multicast rate is 6 Mbps). This is the region where the client cannot connect but it can be inhibited from transmitting by decoding unwanted, co-channel transmissions. Transmissions may be too weak to decode, but strong enough to raise the noise floor and corrupt reception of wanted frames, resulting in a classic case of co-channel interference.

Co-channel interference in the single-channel model

Purveyors of single-channel architectures often claim that their WLANs eliminate co-channel interference, an extraordinary contention given that adjacent access points and clients are by definition on the same channel. One would expect that co-channel interference would be much more significant in the single-channel than the adaptive model.

The fact is that single-channel WLANs are indeed affected by co-channel interference in the same manner as the adaptive WLANs, but to a much more significant degree. Co-channel interference becomes one of the most significant challenges to performance in the single-channel architecture. As explained earlier, when a particular access point or client transmits, all others within its interference zone are inhibited as they sense the transmission, or forced to retransmit as the co-channel interference prevents them decoding the desired signal.

In the section below we will analyze various opportunities to mitigate co-channel interference based on the models above for client- and access point-originated interference, both within and across cells.

Coordinating transmissions within and across cells

Members of a Wi-Fi cell are all able to communicate with the access point, regardless of whether they can 'hear' one another. Leaving aside for the moment the actual mechanism used to coordinate transmissions across the multiple access points and clients in an interference zone, the most important consequence of solving co-channel interference in the time domain is the loss of network throughput. In the adaptive model, each cell is ideally isolated on its own channel. To avoid interference, the CSMA/CA MAC protocol ensures that only one transmission can be on the air (the 'medium' in Ethernet jargon) at one time – maximum throughput within the cell is achieved by enabling members of the cell to use every available millisecond of air time. In the adaptive model maximum throughput within the network is achieved by supporting simultaneous transmissions within adjacent cells wherein each transmission occurs on a different RF channel and thus will not interfere. In the single-channel model, this simultaneous-transmission scenario is not allowed.

Throughput of multi-cell networks, highlighting co-channel interference

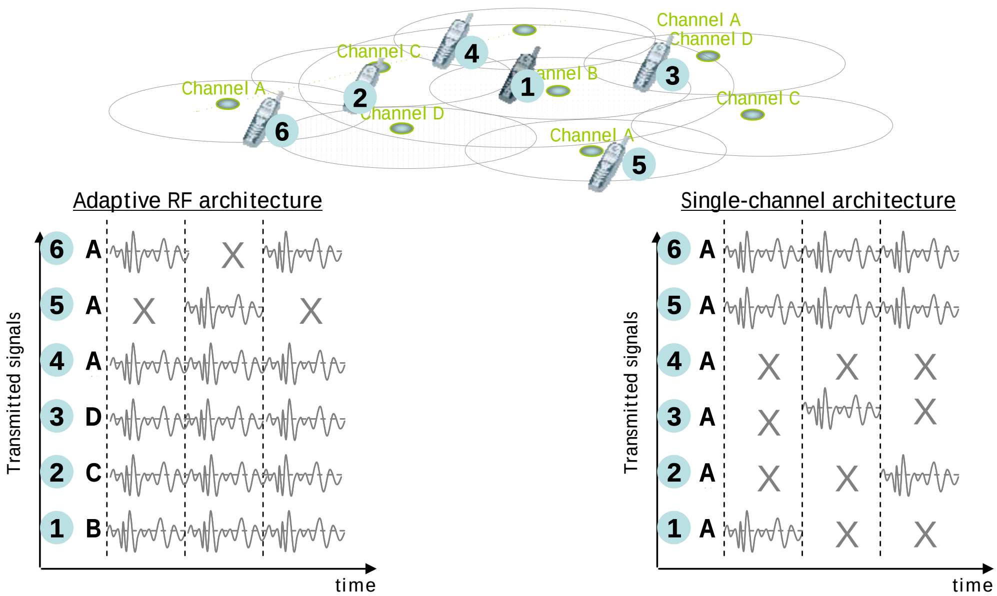

With 4-channel reuse, clients 5 and 6 co-interfere despite their geographical separation. Other clients are isolated by channel or distance separation

Single-channel model introduces considerable opportunities for co-channel interference, particularly client-client interference. Geographic separation allows some simultaneous transmission, for clients 5 and 6 in this example.

Coordination across access points is desirable for mitigating co-channel interference. Whereas adaptive architectures avoid adjacent access points on the same channel, single-channel networks by definition use the same channel for all access points. If adjacent access points transmit simultaneously, clients within their area of common, overlapping coverage will receive two coincident transmissions, with the result that neither will be decipherable. This can be mitigated only if the transmissions from one access point are coordinated in time with all adjacent (and some non-adjacent) access points or if distant access points cannot hear one another.

The diagram below shows the same effect in a different manner.

Time taken to transmit one frame to each of 8 clients, 2 clients to a cell

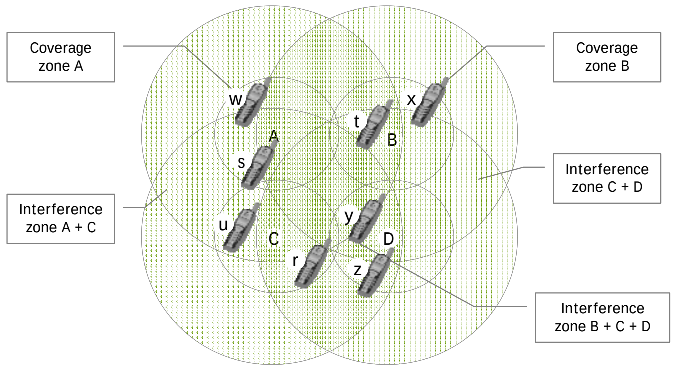

With a single-channel plan, 4 time periods are required

1	**A to w**	**B to x**	**C to r**	D to --
2	**A to s**	B to --	C to --	**D to z**
3	A to w, s	**B to t**	**C to u**, r	D to --
4	A to w, s	B to x	C to u	**D to y**

A 4-channel adaptive plan, 2 time periods

1	**A to w**	**B to x**	**C to r**	**D to y**
2	**A to s**	**B to t**	**C to u**	**D to z**

A 3-channel adaptive plan, 3 time periods

1	**A to w**	**B to x**	**C to r**	**D to z**
2	**A to s**	**B to t**	**C to u**	D to --
3	A to w or s	B to x or t	C to r or u	**D to z**

Assume there is one frame queued on the downlink for each client. How much simultaneity of transmission is possible? The diagram shows coverage and interference zones from each AP. The tables compare the time taken to transmit all frames, allowing for co-channel interference effects. A single-channel model is compared to the adaptive model using 3 and 4 RF channels. The tables show 'wanted' transmissions in bold. Other simultaneous transmission possibilities are noted, but not bold.

The diagram above compares the two architectures for the deployment model just described. The adaptive RF architecture allows simultaneous transmission by different clients whenever they are separated by RF channel, or spatially such that the signal strength is reduced below interference levels. Certain groups or pairs of clients and access points cannot transmit simultaneously, but generally there is good parallelism allowing for high throughput of data across the network.

The single-channel model forces equipment designers to seek alternative mechanisms to avoid co-channel interference in the time dimension, rather than just spatially, since wherever one looks in the network there are many access points and clients in any particular device's interference zone. The single-channel network must coordinate across multiple cells to avoid simultaneous transmissions. The diagram demonstrates that this results in a reduction in overall network throughput. This loss of capacity due to co-channel interference in multi-cell networks is the most significant drawback of the single-channel model.

Another, simpler example demonstrates this loss of capacity. Consider three adjacent cells forming a network. In the adaptive model, each operates on a different RF channel, so at any instant in time, three devices in the network can transmit simultaneously. In the single-channel model, only one device can transmit at any time. In the absence of secondary effects, the adaptive network supports three times the throughput.

Compensating for co-channel interference with the single-channel model: managing transmission opportunities

The analysis above shows that co-channel interference is a significant effect for single-channel networks, even when it is countered with an effective transmission synchronization mechanism. Although such mechanisms exist within cells as part of the MAC, they are much more difficult to realize across cells. For this reason we will now separate the problem into two parts. Since access points are closely controlled by the WLAN switch, we will first investigate ways in which access point transmissions can be coordinated in the single-channel architecture by controlling the downlink. Next, we will investigate ways in which vendors have sought to control client transmissions on the uplink, which is a much more difficult problem to tackle.

Controlling downlink transmissions

Downlink transmissions from access points towards clients can be coordinated when all access points are slaved to a central WLAN switch. The switch, which we have already seen is necessary to accomplish inter-access point handover, has a more difficult task when scheduling transmissions as it must work to tighter timing requirements: transmissions from access points must be scheduled to a few tens of microseconds for good accuracy. If timing is not sufficiently accurate, co-channel interference will not be improved. Few implementations of the single-channel model succeed in accurate coordination of access point transmissions across cells.

Another factor affecting transmission timing across cells is that Wi-Fi operates in unlicensed spectrum. This makes it more likely that nearby equipment or neighboring Wi-Fi networks will cause interference. Interference results in receiving errors, leading to immediate retransmissions. Even when a WLAN switch distributes downlink frames to an access point, the timing window is such that it cannot generate retransmissions because the round-trip path from access point to WLAN switch is too slow. This means the WLAN switch has even less control, as it cannot anticipate a priori which frames will need retransmission, or indeed at what rate they will be transmitted since the data rate can change frame-by-frame. The distributed decision making inherent in the adaptive model is the most effective mechanism for dealing with unpredictable RF conditions.

Once inter-access point coordination is abandoned, the remaining downlink operation is to handle transmissions within a cell, a familiar problem for adaptive and single-channel models alike. Since at least half of the traffic in the cell originates from the access point, controlling transmissions here is very effective in reducing contention time, and there are various techniques in the 802.11 standards to enhance this, such as allowing follow-on transmissions with short inter-frame intervals to avoid contending for the medium again. The downlink is also where QoS is primarily enforced, as WMM defines separate queues for each of four recognized classes of service. Even in the absence of client-based QoS, and before the advent of WMM-capable clients, downlink queuing provided an effective solution for multimedia transmissions.

Controlling uplink transmissions

Coordinating and controlling uplink transmissions, from clients is much more difficult. Indeed, despite some aggressive marketing from single-channel vendors that implies otherwise, it is not possible to control when a Wi-Fi client chooses to transmit because only the client makes this decision. Claims of time domain multiplexing (TDM)-like behavior are exaggerated beyond even a passing acquaintance with reality. While it is possible in Wi-Fi to inhibit clients from transmitting at certain times, and to apply layer 3+ techniques to control the overall flow of data, it is impossible for the infrastructure to specify the particular instant when a client should transmit. We will now review some of the control techniques used by single-channel vendors.

Single-channel architectures, as described above, must coordinate transmissions across the entire network in order to retain a modicum of overall capacity. The cleanest solution would be to use a scheduled or polled MAC algorithm. In such a scheme, the access point could learn when a client had traffic to send, and would then allocate that client a specific transmission opportunity while other members of the cell were inhibited. But standard 802.11 clients do not support this mode of operation, so the single-channel network architect must find a way to use - or misuse - an existing 802.11 mechanism to achieve a similar result.

The most obvious way to do this is by using the 802.11 network allocation vector (NAV). This indicator is contained in every 802.11 frame, and shows the expected duration of the frame. NAV is used by other stations to set their backoff timers: they know they will be unable to transmit at least until the NAV expires, as the medium will be occupied. By directing an access point to send frames with a misleadingly long NAV value, stations can be inhibited from transmitting for a period of time, effectively quieting the stations.

Note that the NAV prevents a station from transmitting within a given period, rather than instructing it exactly when to transmit (a condition that is impossible to achieve using standard 802.11 clients). Thus, single-channel architectures attempt to mitigate the co-channel interference effect by misusing a parameter in ways that were not intended by the authors of the IEEE 802.11 standard. As in all such situations, the practice can sometimes be successful, in other cases less so. It is out of concern for both standards compliance and unintended consequences that designers of the adaptive architecture have not adopted the non-standard practices, even though they are equally applicable to both architectures.

An additional technique involves manipulation at the TCP layer. This includes counters used for flow control between hosts: it is possible to throttle back a TCP/IP connection by changing these counters, or by buffering and delaying packets. This can be accomplished by having hosts use an unacknowledged-window protocol such that a delayed incoming packet will eventually result in delayed outgoing packets. This technique throttles the downlink traffic resulting in a concomitant reduction in the upstream data rate. Note that this is not granular enough to determine actual transmit times, but it can be useful in shaping overall traffic profiles. Such techniques are used by many networking vendors in a variety of situations, and are not unique to the single-channel architecture.

These ‘TDM-like’ scheduling techniques could be used in adaptive architectures, however most vendors have preferred not to introduce non-standard mechanisms that have little benefit in their architectures. However, they could certainly introduce such mechanisms if they anticipated that the improved performance outweighed the risks.

Modeling zones of interference

Since co-channel interference is a very significant impairment to the performance of a single-channel network, developing techniques to accurately model reachability and interference zones is highly desirable. Especially in larger networks, a client would be expected to ‘hear’ only neighboring devices, and clients on the other side of the building would be too far away to interfere with them. Estimating how far zones of interference extend is not easy, but the benefits are significant as several devices can probably transmit simultaneously without causing interference in such networks. Various techniques can be used for this, and a new IEEE 802.11 standard - 802.11k – is intended to allow a client to report activity it ‘hears’ on different channels. However, since 802.11-compliant clients cannot distinguish individual access points in a single-channel network, it may be difficult to use this standard technique in such situations.

Mapping co-channel interference to identify simultaneous transmission opportunities

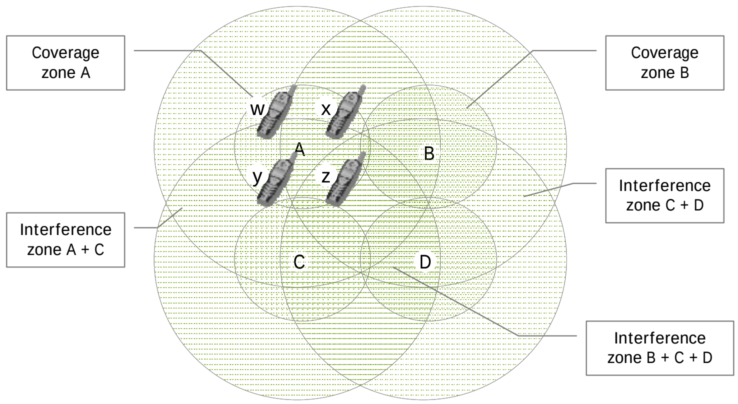

Simple interference map to coordinate transmissions from AP A

When AP A transmits	These cannot transmit simultaneously
A to w	None (e.g. all other APs can transmit)
A to x	B cannot transmit
A to y	B and C cannot transmit
A to z	B, C and D cannot transmit

In an ideal world, the WLAN switch would have perfect knowledge of the interference zones of all access points and all clients in the network, and would update its map in real-time. This would allow for the maximum possible simultaneous transmission and hence maximize network throughput. The computational challenges of maintaining an interference map for each device in the network are discussed later in this note, but we believe that most single-channel vendors have developed some form of reachability modeling to limit the effects of co-channel interference.

Transmit power control

Another technique to reduce co-channel interference involves controlling clients' transmit power using TPC. Network vendors already reduce access points' transmit power to reduce the coverage area while increasing overall capacity. However, most clients today transmit at a fixed (maximum) power level, usually far greater than their access point, exacerbating the client-originated co-channel interference effect above.

TPC allows an access point to dynamically control a client's transmit power level. Originally part of 802.11h and intended to avoid interference with other users of the 5 GHz bands, TPC has been adopted for general use in all Wi-Fi bands and will be incorporated in the forthcoming Wi-Fi Alliance 'voice-over-Wi-Fi for enterprise' certification. Rather than using a static setting for a fixed transmit power level, usually a maximum level, TPC allows the access point to sense when a client is transmitting at higher power than is necessary, and direct it to use a lower setting, almost frame-by-frame. The adoption of TPC in Wi-Fi clients will greatly reduce co-channel interference from clients, the dominant mechanism for transmission impairment in enterprise networks.

The effect of TPC will be applicable to both adaptive and single-channel architectures, effectively reducing the zone of interference caused by client devices, as they will tend to transmit at lower power levels.

Handover across WLAN switches

We showed earlier that in order to handover a client from one access point to the next, the WLAN switch managing the access points must have real-time information and make real-time decisions. This becomes extremely difficult when the access points are slaved to different WLAN switches. The speed of information exchange between switches has so far proved too severe a requirement for single-channel networks, so they do not offer a 'fast inter-switch handover' feature. Indeed, even chassis-based multi-switch clusters do not have this characteristic, and instead operate as functionally separate units. But in the adaptive architecture, inter-switch handover is a much simpler problem and is routinely accomplished through the use of proxy Mobile IP and related protocols.

Inter-WLAN switch handoff is a key differentiator between single-channel and adaptive implementations. Handover across multiple WLAN switches is routine with vendors using the adaptive architecture, but compromised and slow with a single-channel architecture. The lack of handover between layers is a serious limitation in a single-channel network using multiple RF layers spread across different channels. In such a situation voice calls will be interrupted or lost as users move between floors or around larger buildings.

Unintended consequences of the single-channel model

Since the 802.11 standards cover several thousand pages, there are often interpretations of what is and is not ‘compliant’ with the standards, and it is usually left to engineers to negotiate a solution when different interpretations result in problems in multi-vendor networks. However, the single-channel model is built on such a divergent underlying architecture that unintended consequences abound when clients designed for ‘standard’ 802.11 applications are used in a single-channel network. Frequently these problems arise from clients’ inability to distinguish individual access points, or access points’ inability to advertise their differences in beacons and elsewhere.

Voice clients routinely run background scans to find adjacent access points and their received signal strengths, so they have a candidate list when it is time to handover to a new access point. Usually the best choice will be the access point with the highest signal strength which, all other factors being equal, will also be the closest. In a single-channel network in which only one BSSID will be seen, the received signal strength will vary widely and swiftly based on the number of nearby single-channel access points. A standard voice client placed in such an environment can experience various symptoms of confusion, for instance when attempting to populate a list of handover candidate access points.

Another function that is difficult to implement in a single-channel network is client-based call admissions control. One standard approach calls for each access point to publish its current load as an element in the periodic beacon. Obviously if the network appears as one access point to the client, such a figure is meaningless. While the network can implement its own CAC mechanism, this can cause difficulties for standard clients developed and tested against standard access points.

Some third-party vendors’ location tracking and asset tag tracking systems have difficulty working where they cannot distinguish between different access points in the network. Systems where the network uses received signal strength at the access point to triangulate should still function, but client-centric models will not.

Many healthcare institutions are installing in-building RF distribution systems, often ‘distributed antenna systems’ (DAS) that repeat cellular network signals in hard-to-reach areas of the building. Where there is a desire to use this infrastructure for Wi-Fi, the single-channel architecture has particular difficulty in adapting. It is possible to use the two in combination, but many benefits are lost, because the only practical solution is a 1:1 ratio of access point units to remote antennas. This configuration negates the architectural advantages of DAS, splitting signals and then dynamically combining them to redistribute capacity.

While there would be some benefits to synchronizing access point beacons across the network, single-channel architectures have not yet implemented this and instead send transmit beacons isochronously.

However, timestamps for beacons are distributed from the WLAN switch. This allows a statistical condition in which two adjacent access points transmit beacons within a short interval, but network-side jitter results in the second transmission carrying an earlier timestamp. Some client implementations are not prepared for consecutive beacons, ostensibly from the same access point but with out-of-order timestamps, and NIC resets have been observed under such circumstances. The problems can potentially be fixed by the client designer, but would not be discovered in design and testing using ‘standard’ Wi-Fi environments and test scenarios.

When implementing the adaptive model, designers routinely switch access points between channels, to listen on the channels and build-up a picture of the RF environment across the band. This information is used when considering alternate channel plans, and for Intrusion Detection Systems (IDS), where unauthorized ‘rogue’ access points are detected and neutralized. It is possible to perform such functions in a single-channel architecture, but because the WLAN switch requires more information from the network, fewer opportunities exist for such time-slicing to alternate channels. Some security experts consider this weakens the IDS and rogue detection capabilities of such equipment.

Clients are built to work with the infrastructure model agreed upon by standards bodies, which means they exhibit some characteristics that would be unnecessary in a single-channel network. An example is the handover behavior of Wi-Fi phones. Even when the network only offers one channel, clients that have no special settings for such an environment will continuously scan as they seek an alternate channel on which to handover, should it be required. Since the client was not built to work in a single-channel network, when it hits a region of low signal strength it will seek any access point on any channel, which may have unintended consequences if, for instance, a network uses a single (but different) channel for each floor of a building. Such scanning also represents an unnecessary drain on battery life, but until clients are designed specifically for single-channel operation this is unavoidable, and there will be little improvement of battery life over an adaptive WLAN.

Multicast traffic

In Wi-Fi networks, multicast traffic is broadcast from the access point, but since it is potentially intended for all clients in the cell, it uses the lowest data rate configured, or the lowest in use, depending on the implementation. This low transmit data rate means multicast often considerably reduces the overall throughput of the cell.

Multicast in single-channel architectures

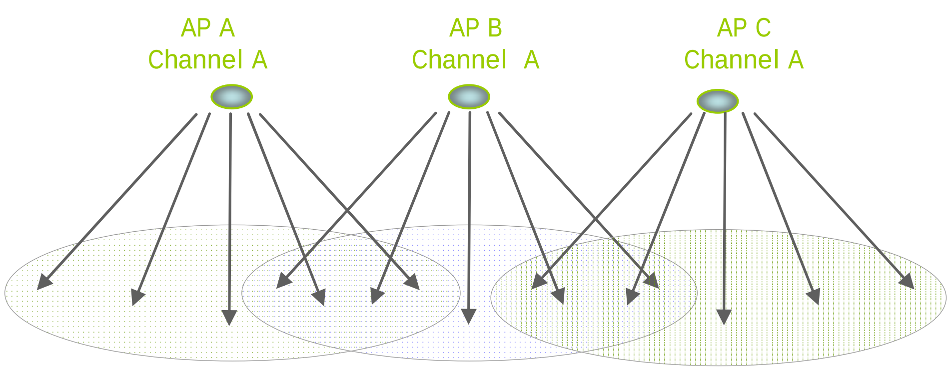

Multicast from the ‘single AP’ base will cause most clients to receive duplicate frames.

When applied to a single-channel architecture, multicast presents a problem. Multicast frames will likely need to be transmitted from every access point in order to reach every client, although a complex control model might be able to prune the tree somewhat. Multicast transmissions must be staggered across access points to avoid interference, but since standard 802.11 clients cannot distinguish different access points in a single-channel network, they are likely to receive several copies of the each multicast frame.

The upshot is that designers sometimes do not attempt to implement multicast in the single-channel architecture, and instead convert all multicast traffic to unicast. This is an adequate solution for some types of traffic, such as telephony push-to-talk groups, but when multicast is used for video transmission, many more frames must be generated compared with the adaptive architecture, adversely affecting network throughput.

Power save and sleep modes

The battery life of hand-held Wi-Fi devices has been improving but remains unsatisfactory for many users. Many techniques for reducing power consumption have been developed, including sleep modes in which the client informs the access point that it is going ‘off the air’ for a time, causing the network to buffer downlink traffic. Access points can normally maintain the sleep status of clients as local data that is used when it becomes necessary to advise the client of a handover to a new access point.

Since single-channel architectures do not inform the client when it transitions to another access point, so they must maintain a model of the client’s state centrally, in the WLAN switch. Additionally, in single-channel architectures the bit maps and other indications to clients that downlink traffic has arrived must be managed across access points rather than as local data. Managing these tasks represents yet another increase in the complexity and volume of real-time data that must be coordinated centrally by the WLAN switch. And it represents further opportunities for a mismatch between the expectations of clients developed and tested against a standards-compliant access point, and the unique infrastructure of single-channel architectures, e.g. state mismatches and clients stuck in sleep mode.

Silent clients

The silent client problem is in some ways a superset of the power save and sleep problems discussed above. Many handheld clients such as phones cease transmitting for minutes on end when not on a call, as they need to save battery life and every transmission consumes power reserves. A phone in the idle state periodically wakes up to listen to access points, both for incoming traffic (advertised in the beacon when the client is in sleep mode), and to assure itself that the access point is still in range. When the phone senses it is losing the access point’s signal in an adaptive WLAN, it enters handover mode to seek out and authenticate with a new access point. In this way, the WLAN switch can track its movement from cell to cell.

The sleeping or silent client forces the network to transmit in all cells (or 'ping' periodically)

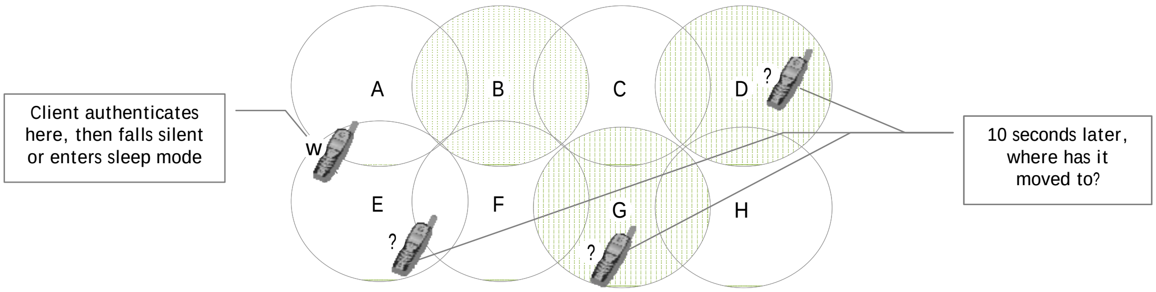

However, in a single-channel architecture the phone will not recognize when it is moving from one cell to another, and if the client is not transmitting, neither will the network. Thus a phone may appear in one part of the network, then after a quiet period of a few minutes, reappear on the other side of the building. Two solutions are available, neither of them palatable. The network can assume that any client not heard from for a few milliseconds could be in any cell, but that would create more overhead and more co-channel interference by requiring broadcast, multicast and traffic map transmissions from every access point in the network for all such clients. Alternatively, the network can send ping-like transmissions to such clients, forcing an acknowledgement and revealing their location. The latter is more often adopted in practice, but it creates unintended consequences such as excessive battery drain on clients.

Special considerations for 802.11n

The new high-speed standard for Wi-Fi, known as 802.11n, introduces many very powerful techniques to improve both throughput and the reliability of communications. Most of these techniques are based on the use of multiple antennas at either or both the access point and client. Multiple antennas are used for multiple input, multiple output (MIMO) and exploit the spatial dimension at the RF layer. One consequence of MIMO is that RF conditions now depend not only on the position of the transmitting and receiving devices, but also on how the antenna signals are driven or combined. Range can be extended using spatial techniques, allowing an access point to cover more area, or to reach a client in a location where previously there was a low RF signal.

There is as yet only limited field experience of the effects of 802.11n on network performance regardless of the RF model used. One aspect that will be a challenge for the single-channel architecture, however, is in the calculation of handover points in overlapping coverage. 802.11a, b and g, all use single antennas, and as a client moves away from a serving access point the received signal strength decays in a relatively smooth manner, while the signal strength of adjacent access points increases smoothly, allowing handover decisions as shown earlier. However, due to spatial processing, with 802.11n we expect the signal strength of the 'old' access point to remain high for much longer. Similarly, the adjacent access point will receive a very weak signal from the device until it starts to target it as a MIMO client, at which point the signal strength will dramatically improve. This means that the inter-cell overlap zone will increase and exhibit non-linear effects. As we saw earlier, this will make handover decisions more difficult for the single-channel architecture.

A second effect of 802.11n will be to increase the number of access points potentially receiving from each client, further increasing the processing burden on the WLAN switch.

Silicon for the single-channel architecture

As we have shown, the single-channel architecture requires significantly tighter coordination between access points and the WLAN switch, as the switch must maintain more information about the state of each client, and coordinate across access points to a much greater extent than in the adaptive model. This has consequences for Wi-Fi infrastructure vendors. Access points are built around Wi-Fi integrated circuits that implement most of the packet-handling and supervisory functions. Since silicon vendors design for the volume market and they implement standard Wi-Fi reference models, they test and verify operation against standards-compliant clients. This can be a problem for single-channel infrastructure vendors, as their reference model diverges from the known standard.

Single-channel networks require more information to be passed from the access point to the WLAN switch, which translates either into extra functionality on the access point silicon or custom firmware images that force the integrated circuit to operate in a non-standard manner. The consequence is that, while single-channel vendors are seldom late to market in releasing new hardware such as 802.11n access points, the hooks to support true single-channel operation often lag by many months. Early product versions often don't support inter-cell handover, for instance, and adding this feature may require upgrades and retesting of an already installed system. The fast pace of innovation in enterprise Wi-Fi becomes even more difficult to maintain for single-channel vendors, as they must add and regression test their proprietary extensions to support single-channel operation.

Building scalable single-channel WLAN switches

The discussion above demonstrated that the single-channel model allows a simplification of client behaviour, as each client only 'sees' a single access point. However, many of the decisions such as handover are moved to the infrastructure side of the network, and must be dealt with in the WLAN switch. It is worth listing the information such a switch must accommodate to support these additional tasks:

i. An interference map should be drawn up, in order to mitigate the effects of co-channel interference. Ideally, this should dynamically track interference zones for each client and access point in the network, although a simpler model might be sufficient based on access point interference zones, with each client mapped to its current serving access point. With such a map, it is possible to identify simultaneous transmission opportunities, where two devices (e.g., access points) can transmit at the same time without causing interference at the desired client;

ii. A 'serving access point' map identifies each client and binds it to the current cell;

iii. The WLAN switch must maintain real-time information about the signal strengths of each client as received by every access point so it can make handover decisions;

iv. Many aspects of clients' state must be maintained in the WLAN switch, in particular the sleep mode, multicast groups and any frames buffered for sleeping clients.

None of the items above are needed in the adaptive architecture, except perhaps (ii), as the clients take care of the functions in a distributed, if uncoordinated manner.

Given the above requirements, a single-channel WLAN switch requires more processing power for a given population of clients and access points, than the equivalent in an adaptive architecture. Since there is a direct correlation between processing power and product cost, the amortized cost per access point will be higher.

Single-channel RF model: inter-cell overlap region too small

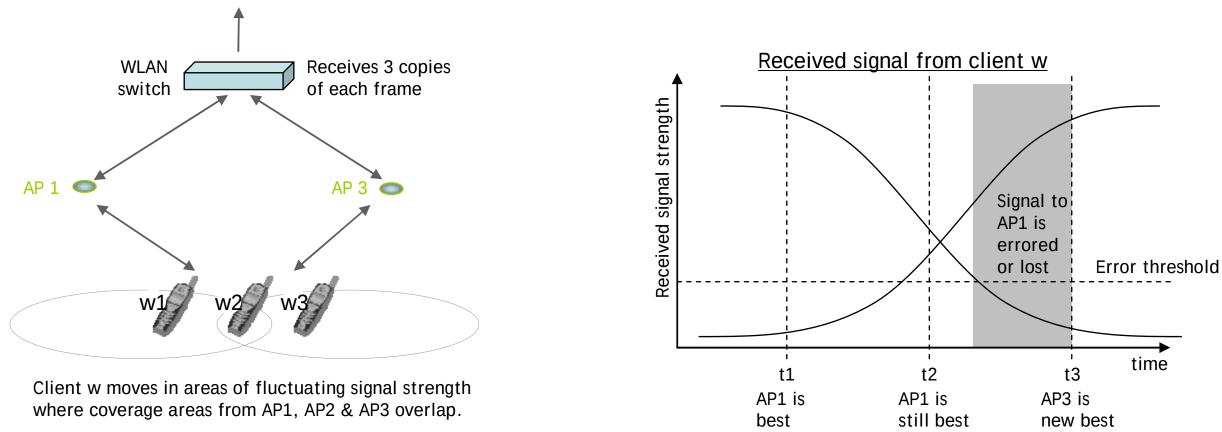

The WLAN switch calculates whether to transmit & receive frames via AP1, AP2 or AP3, based on estimates of the received signal strength etc. from all APs in range. Because the WLAN switch is simultaneously evaluating handover for all clients in the network, it cannot make instant decisions, but can only calculate periodically at times t1, t2, t3. In this example, the calculation at t2 indicates AP1 is the best choice, but the signal is already degrading, and is lost before the handover decision is revisited at time t3.

Additionally, the real-time response requirements are much more stringent, as calculations, particularly (iii) concerning split-second handover decisions, must be very fast. Indeed, the WLAN switch cannot be calculating the ‘best’ access point for each client continuously – it must time-slice such decisions, which means they are effectively scheduled and each client is assessed every few tens of milliseconds at best. Since a client’s situation can change considerably within such a window, a single-channel network must be carefully planned to provide a smooth, relatively wide overlap area between adjacent cells, ensuring a client does not lose the old cell’s signal before the handover algorithm can catch up with it.

RF deployment architectures in the enterprise WLAN environment: the freeway analogy

The enterprise WLAN market has a history of more than a decade, and there is a well-established performance envelope defining the technical requirements for this equipment. We know that sources of RF interference will crop up regularly but unpredictably, and that many enterprise networks will be overlapped by adjacent WLANs run by other, legitimate organizations. We also understand the requirements for scale, client density, cost and performance targets for clients on the WLAN. The rest of this document presents practical uses of the adaptive and single-channel architectures.

We begin with a metaphor – an adaptive WLAN is a freeway with a variable number of lanes which can be added or deleted as the need for capacity grows and shrinks, respectively. The challenge confronting a network engineer is how to sense the capacity requirements and correctly adjust the number of lanes without wasting airtime.

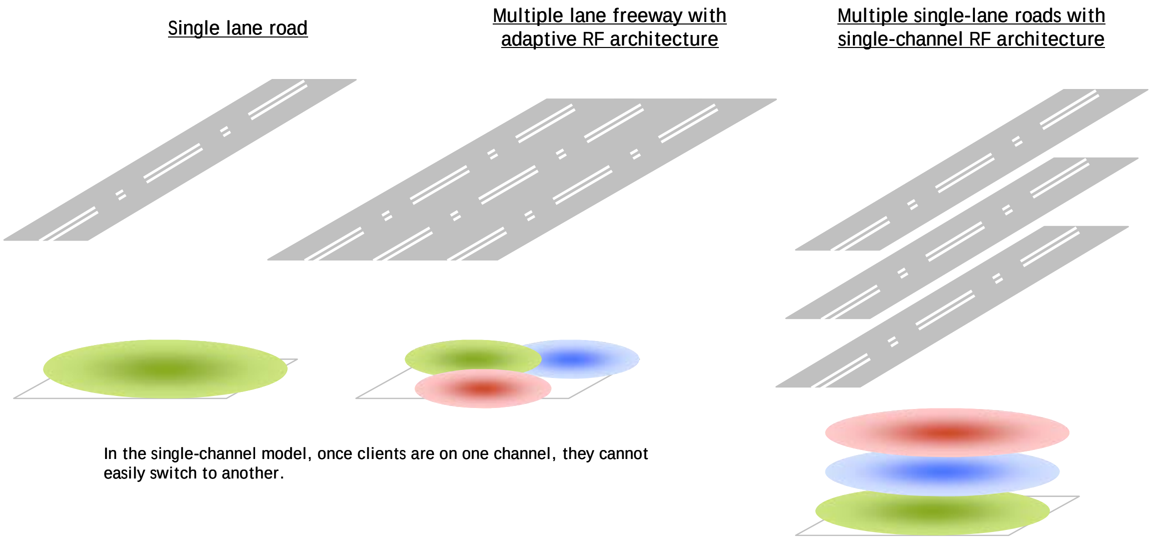

In the single-channel model clients can only handover using the channel on which they first connected, analogous to using only one lane of the freeway. Two consequences result. First, it is not possible to load-balance across channels: if one channel (or road lane) is congested, users cannot be easily directed to other roads, even if they are under utilized. Second, when a channel is introduced into the network to relieve congestion or provide capacity at a local bottleneck, coverage for that channel must be extended across the entire network, which can entail considerable cost. It is no coincidence that single-channel vendors initially advocated multi-radio access points with four or more radios, even though the cost penalty hampered early market acceptance.

The following paragraphs discuss the consequences of these two consequences, as well as other inherent limitations with the single-channel architecture.

Robustness & flexibility of RF channel planning

While we often think of buildings as having clean, isolated RF environments, the reality is quite different. Radio waves travel both short and long distances on unpredictable paths. Further, the RF spectrum used by Wi-Fi is unlicensed - only the maximum transmit power is regulated - so anyone can set up a transmitter at any time, at any location. And in-band interference crops up from a variety of sources: microwave ovens, lighting systems, cordless phones, Bluetooth radios, and nearby Wi-Fi networks operating in the same bands.

Adaptive multi-channel WLANs recognized and solved these problems many years ago. Access points continuously monitor the RF environment on their own and other channels, feeding information to software that repeatedly calculates the optimum RF plan, adapting as the environment changes.

The existence of nearby interfering devices means that the RF plan must be flexible for a network to function reliably. If access points are set up with fixed channel assignments, inevitably over time some areas will become impaired due to non-Wi-Fi interference (generally within the network) or overlap with other (legal) Wi-Fi networks towards the edge. The result is that channel assignments for the affected access points must be modified.

This real-world requirement highlights a basic weakness of the single-channel model. It is impossible to modify one access point's RF channel in isolation: the whole network must change. Every time a change is made in such a network, all access points will switch their channel, often rebooting in the process, resulting in a network-wide outage in the order of minutes before it settles. And to work effectively the newly selected channel must be free, and remain free from interference everywhere in the enterprise.

By contrast, the adaptive model allows access point channel assignments to change individually. When new sources of interference are detected, a new plan can be implemented in a controlled way, access point by access point, and with minimum disruption to clients. Often it is necessary to adjust only a handful of access points to accommodate a new source of interference.

For example, consider an enterprise office situated in a large building with other tenants. Many of these neighbors will operate Wi-Fi networks that could transmit on the same RF channel used by adjacent enterprise access points. In the adaptive model, a quick, automatic recalculation of the RF plan will change the channel assignments of the affected access points with minimal disruption. However, a single-channel network faces significant disruption when the channel in use has to be changed. Indeed, it is quite possible that in aggregate the neighbors would use all three channels available at 2.4 GHz at different edges of the enterprise network, causing major complications for the single-channel model. In field tests it is not uncommon to detect dozens of nearby networks in a large, multi-tenant building deployment.

Another important example of a channel-switching stimulus is radar avoidance. Most national regulators only allow use of the 5 GHz band, or parts thereof, only if dynamic frequency selection (DFS) is enabled. This feature, from 802.11h, periodically senses the presence of other users of the 5 GHz band, primarily military and weather radars. Upon sensing such a user, the access point must quickly vacate the channel to avoid interfering with the other system. Although we expect such radar systems to be relatively static in their channel assignment, regulators are forcing ever more sensitive detection algorithms on Wi-Fi devices, causing frequent events where the detection mechanism is falsely triggered. When such an event, real or false, causes a channel switch, a single-channel network will suffer a hugely disruptive event, as all clients are disconnected and forced to find the new RF channel. An adaptive Wi-Fi network, however, will only need to adjust the small number of cells using the affected channel, typically with minimal disruption.

Positioning access points

As we noted earlier, installers are often constrained with respect to the location at which access points can be mounted in enterprise Wi-Fi networks. In the adaptive model, exact planning and installation locations are not required, because the RF channel and transmit power are automatically recalibrated once the network is operational. Single-channel networks are not able to change the RF channel of individual cells, so transmit power is the key variable for network tuning. But some single-channel vendors are unable to reliably change even this variable: instead power is fixed at the maximum legal transmit power. This limitation removes the ability to reduce the size of cell coverage, an important tool for increasing the capacity of the network on a per user basis.

Balancing overlapping coverage with co-channel interference

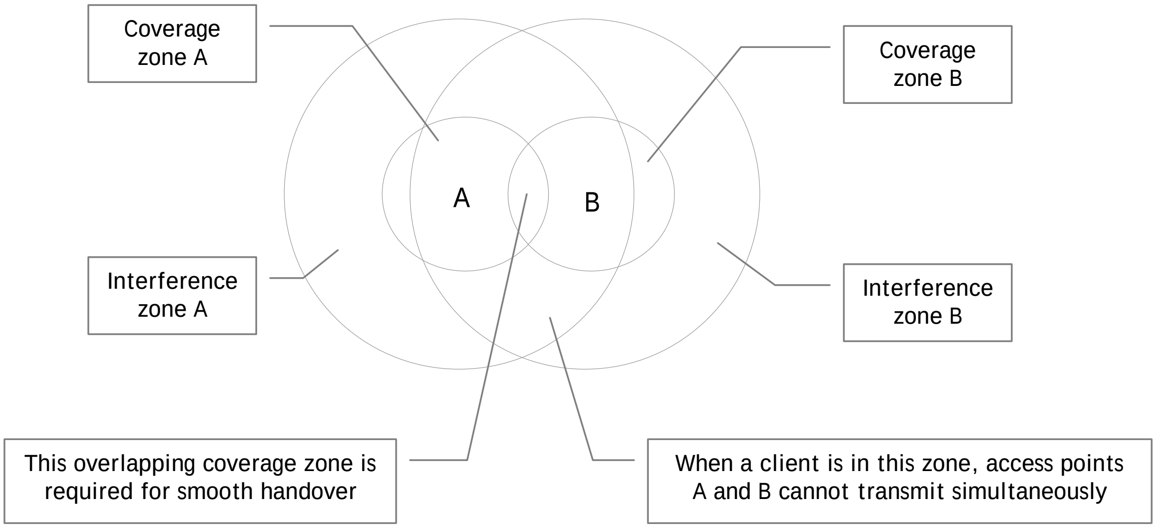

One reason for this inflexible approach in the single-channel model is that the inter-cell coverage overlap area is a critical parameter for good handovers. We showed earlier that, because the WLAN switch must make handover decisions for each client in the network, it is forced to time-slice, as it cannot simultaneously and continuously process signal strength figures from all access points. This means that each client is visited for a handover snapshot only periodically. If the time period between snapshots is close to the time in which a phone client can roam across the inter-cell overlap region, the WLAN switch may fail to handover the client before it loses coverage, leading to gaps in the speech signal.

For this reason, the inter-cell handover region must be relatively large, but as access points move closer together, co-channel interference becomes worse and network capacity suffers. Given a high, fixed access point transmit power, access point placement is the only variable left to the network designer in controlling this key variable, and the exact placement of access points and antennas can be critical in some vendors' implementations of the single-channel model. In some situations, a minor change in the RF characteristics of the environment, such as moving furniture, can make the difference between good and poor handovers.

Fairness

All IEEE 802.11 protocols are backwards-compatible and most Wi-Fi networks support a mixed population of 802.11a/b/g and possibly 802.11n clients which must all share common access points. Because newer clients are capable of higher speeds, some vendors have suggested that newer clients be given priority over older clients. As a numerical example, consider a cell with an 802.11b client connecting at 5.5 Mbps and an 802.11g client at 36 Mbps. If both clients are allowed equal opportunities to seize the medium, they will achieve throughput proportional to their, data rates, so the newer client should be able to send 6x the traffic of the older one. From another point of view, allowing the 802.11b client to use 50% of the airtime will reduce the capacity of the cell – it is time on the medium multiplied by data rate that determines throughput, and the introduction of a slow client in a cell can seriously reduce capacity.

This effect may not be desirable, and some vendors advocate artificially restricting older clients' use of the medium to avoid it. The reason single-channel vendors give it prominence may be that some form of control over client transmit opportunities is much more significant in that architecture, as we showed earlier. However, there may be situations where the reverse is desired – all clients need to be allowed equal throughput on an equal footing, regardless of their capabilities. The broader solution is to provide a policy engine and enforcement mechanism for traffic shaping, which can then be tuned by network managers, and all enterprise Wi-Fi vendors, regardless of internal architecture, should support such a feature.

Scale & capacity

Real-world enterprise WLANs cover large buildings with many floors and diverse user populations. This gives rise to a number of requirements for flexible RF deployment to support more capacity, denser client population, and so on. The adaptive model accommodates this requirement by adding more access points locally, and shrinking the cell size.

Adaptive RF model

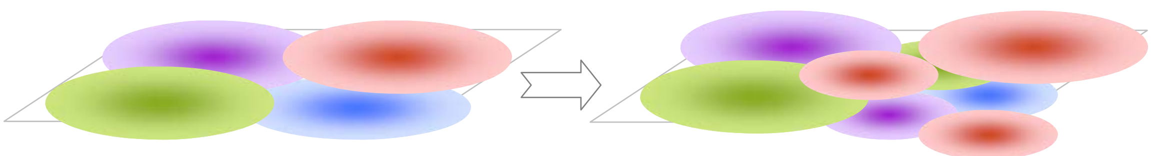

The adaptive RF model allows capacity to be added as and where required, by adding new access points and reducing the cell size.

Wi-Fi works in multiple frequency bands, and enterprises routinely use two bands in parallel at 2.4 and 5GHz. This essentially doubles the capacity of the WLAN. While some clients operate on only one band, most PCs are dual-band capable.

Multi-band operation is a simple extension of the adaptive RF concept, where handover from one access point to the next entails a change of channel: there is no difference if the new access point is in a different band. However, when applied to a single-channel architecture there is a significant increase in complexity.

If a new channel is added anywhere in the network, it becomes necessary to build out coverage on the new channel all across the network.

Single-channel RF model

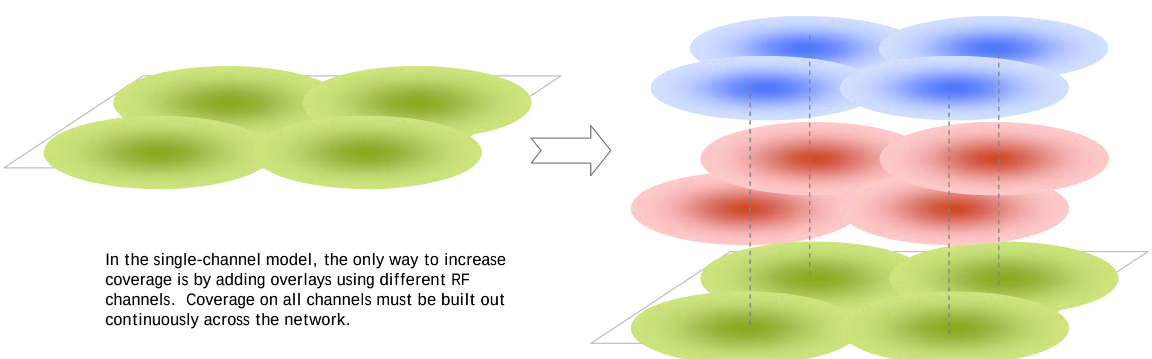

In the single-channel model, the only way to increase coverage is by adding overlays using different RF channels. Coverage on all channels must be built out continuously across the network.

Single-channel vendors faced with such a problem might protest that clients can indeed switch channels, even in such an architecture. However, if such a handover is used, it will need to be of the same quality as any other inter-AP handover – meaning all access points will need to support the requisite number of channels. Also, single-channel vendors seek to modify clients to make them ‘stickier’, that is, less prone to switch access points, because this allows them to perform better in a single-channel network. Introducing a second parallel network on a different RF channel means that such modifications must be removed, and it becomes difficult to control clients: how do they know whether to stay on the previous channel when the RF signal degrades, or to switch to the new one?

Scale is another requirement for modern WLAN networks, as many networks now cover millions of square meters using hundreds of access points. The adaptive architecture is well-proven in these large networks. Single-channel vendors prefer in practice to use a different channel on each floor of a building because it makes the network more manageable when interference and neighboring networks are present (see above). Again, this requires clients to switch channels as they move about the building, introducing the ‘classic’ Wi-Fi inter-AP handover function.

The amount of real-time calculation required of the WLAN switch also limits the scale of single-channel networks. Taking an ‘air-traffic control’ analogy, the single-channel model requires the central WLAN switch to steer every plane in the air. As discussed earlier, this mandates constant reporting of the position of each plane, and time-bounded decisions on when to change direction or speed: if the controller fails to re-direct a plane in time, it may lose its bearings and go adrift. This challenge increases with the number of planes in the sky or clients in the network. Conversely, the adaptive architecture allows each client to make its own decisions, based on observations of its surroundings and general guidance. Distributed decision-making clearly leads to a much more scalable architecture: whereas WLAN switches using adaptive architecture can control thousands of access points, single-channel WLAN switches scale only to hundreds of access points, and as shown above, do not support fast inter-switch handoff.

Client density

High client density is claimed to be a unique advantage of the single-channel architecture. A small area with many clients can be covered using a number of parallel single-channel networks, leading to superior client density over an adaptive network. Or so the argument goes.

In reality, each single-channel overlay is limited by co-channel interference, as explained earlier in this paper. In a small area in which every device has an interference zone covering the whole network, it is impossible for more than one device to transmit simultaneously on each channel, regardless of the number of access points. So each single-channel overlay reduces to a single access point or cell equivalent. If three channels are available, the network capacity cannot exceed three simultaneous transmissions, the same result as for an adaptive architecture.

Single-channel vendors sometimes claim to control client transmit schedules and deliver ‘TDM-like’ behavior. As discussed earlier, this is impossible to achieve for data traffic. We have already shown that even schemes that inhibit transmission by all clients in a cell are counter to the intent of standards and have unintended consequences. Wherever the majority of the traffic is data traffic, such claims are wide of the mark.

There is a grain of truth in the ‘TDM-like’ behavior when applied to voice clients, as voice frames are already periodic, transmitting at intervals of 20 or 30 msec. Many devices can also synchronize uplink transmissions to the downlink: when they receive a frame, they immediately transmit in return if there is a frame to send. This allows an access point (or WLAN switch) to buffer downlink frames until the desired transmit slot, thereby influencing uplink transmissions from the client. This is a clever and legal technique, but it can be equally applied to both the adaptive and single-channel models. In fact, large voice networks are largely self-organizing, as voice clients tend to synchronize transmissions in a sequential manner as a by-product of the CSMA/CA MAC algorithm.

The intent of the client control schemes above is to mitigate co-channel interference, which would otherwise devastate network capacity compared to the adaptive model. But another effect they cannot counter is the reduction in overall capacity of the network, compared to the adaptive model in which multiple RF channels are used as one pool of bandwidth. To return to the freeway analogy, once a client has joined a particular lane of traffic, it must stay there. If one lane is congested in the single-channel model, it is difficult to switch clients to adjacent lanes, even if those lanes are unused. In Wi-Fi terms, while the adaptive architecture will allow smooth handover between cells for load balancing, the single-channel architecture will not be able to handover clients from access point to access point. To overcome this issue in areas with high density clients, lecture halls in particular, single channel networks are typically overlaid with multiple parallel channels and the clients are distributed among the access points. This scenario is only valid when clients remain fixed – it cannot be used when clients must roam because it requires a duplication of the multiple single-channel overlay throughout the facility.

The usual approach when serving high client densities with the adaptive architecture is to use multiple access points, set the network to use the maximum available number of channels, and allow automated RF planning tools to optimize the channel plan. Such systems have allowed a network of 30 access points to serve client densities of 700+ people in an area of 2000 sq meters with good performance and without any special planning, a figure exceeding any claims of single-channel vendors. An additional benefit of this architecture is that clients can roam freely in and out of the high density areas – without requiring an expensive overlay in the rest of the facility – by simply migrating to new channels as the 802.11 standard intended.

Cost

Constructing a WLAN using single-channel architecture becomes expensive as soon as multi-channel coverage is required. Instead of switching radios in each access point to the desired band as and where required by the client population, it is necessary to provide building-wide coverage on each band. The consequence is that single-channel vendors have long advocated access points sporting three, four or more radios. The use of many-radio access points does indeed result in high-capacity, continuous coverage, but at a price premium to networks built on the adaptive model. Consequently, vendors often discuss the multi-radio, overlapping coverage design for single-channel networks, only to considerably simplify the network design – and capabilities - when the excessive cost becomes apparent.

Conclusion

In this paper we have explained the essential differences between a single-channel WLAN model and an adaptive model. Single-channel proponents have to date focused on the one undoubted benefit of the architecture - it removes the need for the client to initiate and execute an inter-access point handover. In theory, a handset or laptop computer is not aware that a handover is in process. Unfortunately, the overall result of moving to a single-channel model is an unfavorable tradeoff, as the deleterious side-effects of the non-standard single-channel architecture outweigh this single advantage.

The situation is further complicated because over time, single-channel vendors have attempted to make virtues out of a number of features that they have in fact been forced to introduce in order for their architecture to function at all. The attempt to coordinate transmission instants, both across the network and within cells, was initially an imperative to mitigate serious co-channel interference effects. It has somehow morphed into claims of 'TDM-like' behavior that do not stand up to rigorous analysis.

Even claims of stress-free handover are exaggerated because single-channel vendors still must work with 'standard' devices that were not designed to have handover remotely managed. In practice many off-the-shelf Wi-Fi clients – Intel Centrino, for instance – can become unstable in the presence of multiple access points with the same BSSID. Additionally, the WLAN infrastructure must implement a handover mechanism which adds complexity in the network infrastructure and, in particular, in the real-time data the WLAN switch must process. The tighter linkage between access point and switch makes inter-switch handover extremely challenging to implement, and can show instability where access point-to-WLAN switch communication is imperfect, for instance with remote access points, limited bandwidth, or errors on the link.

Also, despite the promise of the single-channel architecture, customers have found that it is important to closely control the overlapping coverage between access points, and to run different floors of a building on different channels to avoid unwanted side-effects. Therefore, while partially solving the handover problem, the single-channel architecture introduces a number of significant adverse consequences in other areas.

Co-channel interference is a more significant problem, because the interference zone for each client covers many more devices. Single-channel proponents have attempted to mitigate this by using transmit scheduling. While control and scheduling of access point transmissions is achievable, and indeed is practiced by all Wi-Fi vendors, attempts to schedule standard 802.11 clients have not proven as successful. Such 'TDM-like' schemes have not been revealed in engineering detail, leading most industry observers to conclude that they violate 802.11 standards in different ways. Regardless, these schemes have not been demonstrably effective in negating the adverse effects of co-channel interference in single-channel networks.

Single-channel architectures cannot adapt gracefully when interference and neighbouring networks make the current channel unusable in localized areas of the network. If there are multiple nearby networks, the pool of available interference-free channels may be drained, forcing repeated and disruptive RF channel changes across the network. Adding channels to mitigate this problem requires building out coverage on that new channel across the whole coverage area of the network. This adds unnecessary cost, and is akin to overbuilding a highway – it simply doesn't make the best use of airtime from the available channels.

Single-channel architectures require network-wide build-out of access points supporting the non-standard mode of operation. While the infrastructure can be constructed, the clients that have to run on it were designed and tested for an adaptive RF environment. That explains why many clients that work well in adaptive RF environments become unstable or unreliable in single-channel deployments – single-channel technology is proprietary and not the design target for 802.11 clients.

Finally, one of the alluring claims of the single-channel architecture is its alleged simplicity of deployment. Place access points anywhere, and the networks will configure themselves. The reality is that single-channel networks are susceptible to more and more diverse deployment considerations than adaptive RF networks – co-channel interference, adjacent network interference, roaming into and out of high density areas, channel overlays. Adaptive RF architectures accommodate these issues because they were anticipated in the 802.11 standard, and standard mechanisms are in place to address them. In addition, adaptive RF networks have evolved to the point where most installation set-up, including comprehensive channel planning surveys required in the past, are now managed automatically – without manual intervention. Any improvements in simplicity of deployments have thus long since been negated.

Taken together adaptive RF networks offer the claimed benefits of single-channel networks, but without the deployment concerns, and have the benefit of being based on the open standards for which all Wi-Fi clients were designed. The advantages are real, not virtual, and the scope and size of adaptive RF deployments – the largest in the world across all industries and markets – are a testament to their performance, scalability, and value.

About Aruba Networks, Inc.

People move. Networks must follow. Aruba securely delivers networks to users, wherever they work or roam. Our unified mobility solutions include Wi-Fi networks, identity-based security, remote access and cellular services, and centralized network management to enable the Follow-Me Enterprise that moves in lock-step with users:

- Follow-Me Connectivity: 802.11a/b/g/n Wi-Fi networks ensure that users are always within reach of mission-critical information;
- Follow-Me Security: Identity-based security assigns access policies to users, enforcing those policies whenever and wherever a network is accessed;
- Follow-Me Applications: Remote access solutions and cellular network integration ensure uninterrupted access to applications as users move.

The cost, convenience, and security benefits of our unified mobility solutions are fundamentally changing how and where we work. Listed on the NASDAQ and Russell 2000® Index, Aruba is based in Sunnyvale, California, and has operations throughout the Americas, Europe, Middle East, and Asia Pacific regions. To learn more, visit Aruba at http://www.arubanetworks.com.

WP_RFARCH_US_080228

1322 Crossman Ave. Sunnyvale, CA 94089-1113
Tel. +1.408.227.4500 | Fax. +1.408.227.4550 | info@arubanetworks.com
http://www.arubanetworks.com

WP_RFARCH_080228_US