ZEBRA

DESIGNING A HIGH DENSITY WLAN

Executive summary

Wireless LANs (WLANs) have traditionally been implemented for low bandwidth applications such as web browsing and point-of-sale (POS). But with the evolution of Internet based applications and mobile devices, millions of users are nearly always online and expect to be connected wherever they are — at work and home, on the road, in hotel, cafes, universities and more.

As a result, the deployment architecture of Wi-Fi networks is shifting from large cell coverage areas to application-specific coverage areas. Since Wi-Fi is a half-duplex shared medium, network administrators are finding it increasingly challenging to plan, deploy and maintain high density wireless LANs (WLANs) in areas such as university auditoriums, theme parks and airports. This paper outlines the key challenges and provides a methodology that can be utilized to address the dilemmas network administrators face when designing high density networks (Figure 1).

Figure 1: Design methodology

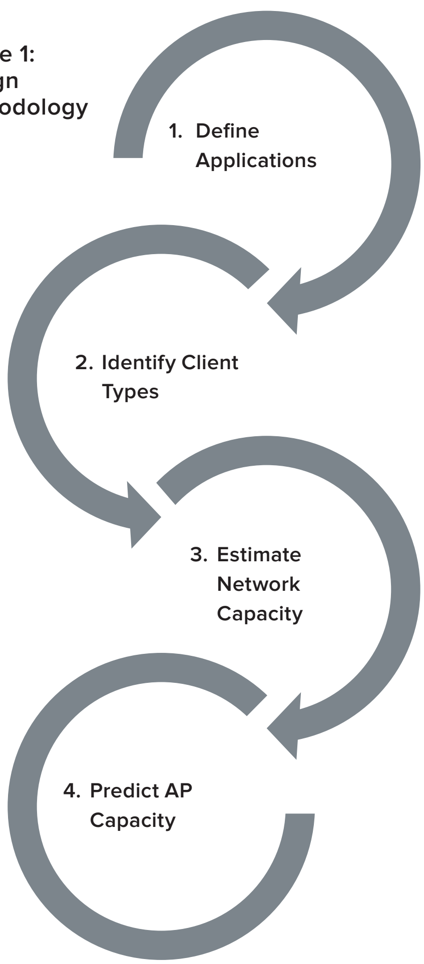

1 DESIGN CONSIDERATIONS: DEFINING THE OVERALL REQUIREMENTS

While the packet data rates for the various Wi-Fi technologies are well known — 802.11a/b/g/n/ac — these numbers are easily misunderstood as actual throughput per client. For example, 802.11n access points (APs) are advertised as 300 Mbps — however, due to numerous factors, this does not translate to 300 Mbps for every client device. Understanding these key factors and how they affect network performance and overall capacity is an important step in developing a successful design for a high-performance high density WLAN

The first step in the development of a successful high density network design starts with defining specific requirements, factoring in the number of devices, types of devices and the specific types of applications that will be running on those devices. Following is an example:

THE WLAN NETWORK SHOULD BE CAPABLE OF SUPPORTING 150 CLIENT DEVICES LIKE TABLETS AND SMART PHONES STREAMING 1 MBPS YOUTUBE STYLE VIDEOS

2 DEFINING THE APPLICATIONS APPLICATION THROUGHPUT = 1 MBPS YOUTUBE STYLE VIDEOS

The next step is to define all the applications that need to be supported on a WLAN and their related bandwidth requirements, since both will impact AP capacity and AP placement. It is important to understand that Service Level Agreements (SLAs) for watching educational videos is very different from supporting voice calls. High throughput intensive applications such as video will reduce the amount of clients supported on a single radio, leading to requirements for a higher number of APs. While figuring out how applications affect capacity is challenging, it is a crucial step in designing a high performance high density WLAN, which is why our sample requirements definition above includes estimated application bandwidth.

Figure 2: Application throughput

APPLICATION	BANDWIDTH
Netflix HD Quality[1]	5 Mbps
Netflix DVD Quality[1]	3 Mbps
FaceTime[2]	1 Mbps
YouTube[3]	500 Kbps – 1 Mbps
Web Browsing	500 Kbps – 1 Mbps
Skype (HD)[4]	1.5 Mbps
Google Hangout[5]	1 Mbps
Google Play[6]	320 Kbps
Facebook video calls[7]	500 Kbps

AN APPLICATION'S BANDWIDTH REQUIREMENTS PLAY AN IMPORTANT ROLE IN DETERMINING THE AP CAPACITY. IF THIS PARAMETER IS UNKNOWN, IT IS STRONGLY RECOMMENDED TO USE NETWORK SNIFFING TOOLS TO ESTIMATE HOW MUCH BANDWIDTH THE APPLICATION ACTUALLY REQUIRES.

The following table lists some common bandwidth requirements to help with this step.

Source:

https://support.ne_lix.com/en/node/306
http://support.apple.com/kb/HT4534
https://support.google.com/youtube/answer/78358?hl=en
https://support.skype.com/en/faq/FA1417/how-much-bandwidth-doesskype-need
https://support.google.com/plus/answer/1216376?hl=en
https://support.google.com/googleplay/answer/1391343?hl=en
https://www.facebook.com/help/325947034156919/

3 IDENTIFYING CLIENT TYPES CLIENT DEVICES = TABLETS AND SMARTPHONES

There has been an explosion in the various kinds of client devices, from smartphones and tablets to laptops. Most devices shipped today are 802.11n compliant, however they don't support the maximum 802.11n data rate of 300 or 450 Mbps. To conserve the battery life of mobile devices, manufacturers limit the devices to lower data rates, for example, 65 Mbps. This has an impact on the overall network throughput and needs to be factored into the design of a high density network.

With 802.11a/g networks, all clients support identical data rates. However 802.11ac/n introduces varying modes of operation: single stream to multiple streams; 20MHz, 40 MHz, 80MHz or 160MHz operation; multiple radio chains and so on. While two 802.11n certified clients will interoperate, they may not necessarily have the same speeds and feeds.

For example, with a 300 Mbps capable AP with a 3 x 3 : 2 MIMO (such as Zebra's AP7131), the effective TCP throughput can be close to 190 Mbps (1500 byte frames). This assumes that the client is also capable of 300 Mbps speed. With the same AP, if you associate a smartphone or tablet that is only capable of 65 Mbps data rates, the effective TCP throughput will be around 30-40 Mbps (1500 byte frames).

Why is there a difference? This is analogous to a highway with a speed limit of 300 mph. Now if one car travels at 65 mph, the efficiency of the highway is significantly reduced. Adding more and more cars traveling at 65 mph will lead to a traffic jam. The same is true for wireless networks. As the number of lower data rate clients increases, the overall available bandwidth shrinks, since these devices spend more time on air.

4 ESTIMATING NETWORK THROUGHPUT

The total network capacity is a function of technologies (a/b/g/n/ac) and applications. 802.11n devices support varying speeds and generate overhead in presence of legacy clients that can lead to throughput degradation. Understanding these complexities will aid in a better design.

Zebra conducted tests to demonstrate how throughput is affected by increasing the number of clients as well as mixed mode environments. The tests were conducted using the industry's leading test tool, IxChariot®.

By simulating real-world scenarios, IxChariot enables the prediction of network application performance under realistic load conditions prior to deployment, allowing enterprises to implement new applications and devices without putting service quality at risk.

Figure 3: Throughput test results

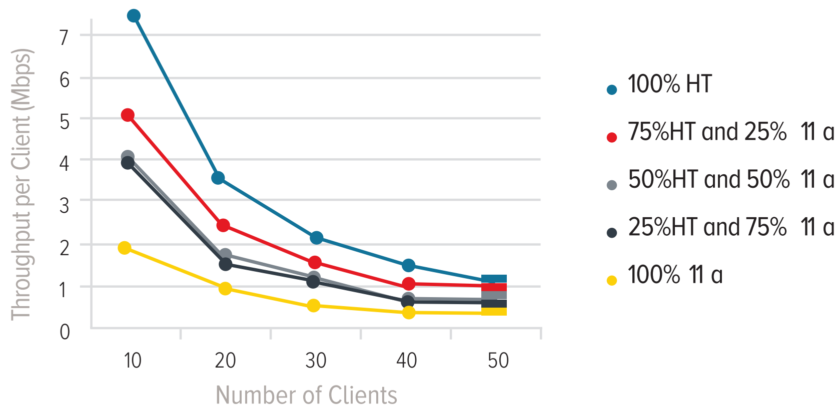

The test was conducted using:

- Fifty (50) 802.11n and 802.11a clients
- A single AP radio – 20Mhz Channel
- The IxChariot test tool (HighThroughput.scr)
- Security: WPA2-PSK

The impact of increasing the number of clients

Let's consider the 100 percent 11n scenario in Figure 3. With 10 clients connected to the AP, the total network bandwidth is equally shared by each client. As Figure 3 illustrates, each client is allocated over 7 Mbps, for a total of over 70 Mbps (7 Mbps * 10 clients) of effective bandwidth. As the number of clients increases, the collision domain increases, leading to degradation in throughput — analogous to a traffic jam on a highway.

With 50 clients in the network, the throughput per client decreases over 80 percent, from approximately 7 Mbps to 1 Mbps. The effective overall network throughput decreases just under 30 percent, from 70 Mbps to 50 Mbps (1 Mbps * 50 clients).

These numbers illustrate the importance of factoring in the effects of collisions when designing a high density network. A Collision Factor of 0.8 can be used to help predict AP capacity, which will be addressed in Section 5.

Mixed mode test

Figure 3 also illustrates the impact of various ratios of mixed modes. We start with 10 802.11n clients and slowly replace some of the devices with 802.11a clients. The impact is dramatic — the throughput per client drops from 7 Mbps to 4 Mbps (50% 11n and 50% 11a) per client. Why? Mixed mode deployments are synonymous to cars and bicycles sharing the same road. As more bicycles enter the road, traffic jams begin to build up — analogous to throughput degradation in the wireless LAN. A Mixed Mode Factor of 0.6 can be used in AP prediction as demonstrated in section 5.

THE EFFECTIVE NETWORK THROUGHPUT IS DRASTICALLY AFFECTED BY MANY FACTORS, SUCH AS APPLICATION TYPE, CLIENT TYPE, COLLISION DOMAIN, AND MIXED MODE RATIO, JUST TO NAME A FEW. TAKING THESE FACTORS IN TO ACCOUNT ENABLES THE INTELLIGENT DESIGN OF A HIGH DENSITY NETWORK THAT WILL PROVIDE SUPERIOR PERFORMANCE ON ALL CLIENTS AND ALL APPLICATIONS.

Once the type of clients (such as smartphones and tablets) and applications (such as YouTube streaming videos) have been defined, Total Network Throughput must be calculated. Using Chariot and a 65 Mbps client, the effective throughput can be calculated as follows:

Aggregate Network Throughput = 40 Mbps

$$\frac{\text{Aggregate Network Throughput (40 Mbps)}}{\text{Application Throughput (1 Mpbs)}} = \text{40 Clients per radio}$$

Next, the Collision Factor needs to be taken in to account:

Client Count = Base Client Count * Collision Factor

= 40 * 0.8 = 32 Clients per radio

When the Collision factor is calculated, the number of clients per radio is reduced from 40 to 32.

Now the following formula can be used to determine how many APs are required to support the total number of clients — which is 150 in our example.

$$\text{Total Number of Radios} = \frac{150}{32} \sim \text{5 Radios}$$

5 PREDICTING ACCESS POINT CAPACITY

Based on parameters defined in the previous section, a rough estimate can be made to predict AP Capacity. The first step is getting a basic client count per AP using the following calculation:

Thus, three dual radio APs are required to support 150 clients streaming YouTube videos. If there were a mix of legacy clients, the Mixed Mode Factor would be utilized as well. The result would be a further decrease in the number of clients supported per radio, increasing the number of access points required.

There are other factors that affect the AP Capacity and should be taken in to consideration as well, including channel availability and channel interference, discussed in the following sections.

Channel Planning

The availability of a channel depends on the regulatory domain where the AP will be deployed. The first step in channel planning is to obtain the list of allowed channels. In the 2.4 GHz band, there are limited number of channels that can be used. In the US, channels 1 through 11 are available.

Based on the channel overlap depicted in Figure 4, only three channels can be used; 1, 6 and 11. Thus if there are more than three APs, depending on the spacing between the APs, re-using channels can be detrimental, leading to degradation of network performance. In an auditorium-like environment, the AP placement can be very dense and channel re-use may not be possible. Also, using 802.11n channel bonding is not recommended.

Figure 4: 2.4 GHz spectrum

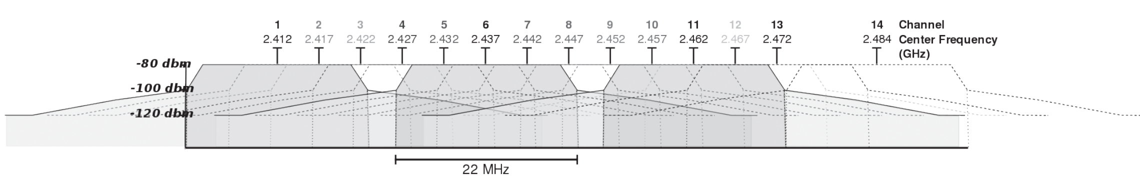

http://en.wikipedia.org/wiki/File:2.4_GHz_Wi-Fi_channels_%28802.11g_WLAN%29.svg

Figure 5: 5 GHz available channel list

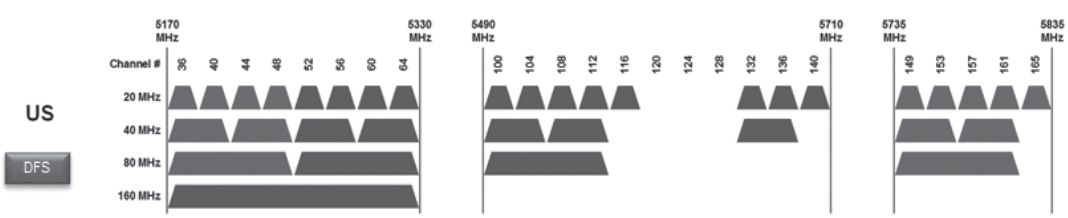

The 5 GHz band in the U.S. has more non-overlapping channels. For example, the Zebra AP 7131N access point has 21 available channels.

This provides more room to collate APs, without the need for channel re-use. The list of available channels varies based on the model number of the AP and country of operation. Please refer to the AP technical documentation to verify the available channel list.

One thing to note when deploying in the 5 GHz band is the Dynamic Frequency Selection (DFS) requirement on certain channels. DFS requires APs to detect radar signals and change channels. Out of the 20 channels in the U.S., only 9 channels are non-DFS (36-48 and 149-165). Figure 5 shows a list of allowed channels in the U.S.

The big question is whether to use 20 MHz or 40 MHz channels. If the client density is high and the clients are smartphones, tablets and other single stream clients, then using 40 MHz does not provide any gain. In fact, it can lead to interference as the number of channels reduces and frequency re-use may become necessary. With 802.11ac on the horizon, 80 MHz and 160 MHz are additional choices — although these will not be available until the second generation of 802.11ac. Like 802.11n, using channel bonding may not be beneficial if there are a substantial number of legacy clients.

Co-Channel and Adjacent Channel Interference

In a high density deployment, APs can be placed very close to each other (such as in an auditorium). If frequency needs to be reused, two APs on the same channel can interfere with each other. The severity of the degradation will be dependent upon AP power and the distance between the two APs. Thus when frequency reuse is implemented, the APs should be placed as far apart as possible and configured with lowest possible power to help minimize interference and performance degradation.

PREDICTING AP CAPACITY REQUIRES ADMINISTRATORS TO DEFINE APPLICATION REQUIREMENTS, SUCH AS CLIENT TYPES. OTHER FACTORS CAN BE UNCOVERED WITH A SITE SURVEY, SUCH AS ALLOWABLE CHANNELS AND INTERFERENCE, WHICH CAN DRAMATICALLY IMPACT THE AP COUNT.

On the other hand, adjacent channel interference occurs when APs are configured on adjacent channels. Since the side lobes of the channels overlap, there can be significant interference if the APs are not far enough apart. As seen in Figure 6, when two APs are placed 10 feet apart on channel 1 and 6 respectively at 19dm output power, there is significant intereference. The side lobe of the AP on channel 6 extends all the way to channel 11. Hence it is important to space APs appropriately and configure the power levels to reduce interference. A thorough site survey can ensure such interferences are minimized.

Figure 6: Adjacent channel interference

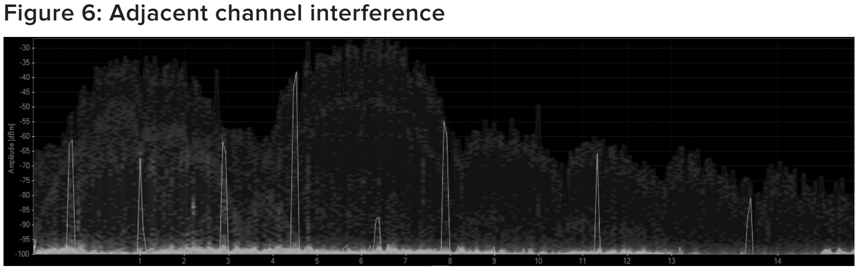

6 802.11AC: 5TH GENERATION WI-FI, BUILT FOR HIGH DENSITY WLANS

The first generations of 802.11ac products have been released. These infrastructure products support up to 1.3 Gbps data rates if three spatial streams and an 80 Mhz channel width are enabled. While most first generation clients, especially smartphones and tablets, will not support three spatial streams, the increased data rate will help boost network performance.

802.11ac is built on 802.11n, and uses the same concepts to boost throughput — such as channel width, short guard interval and aggregation. The big difference comes at the physical (PHY) layer through improved modulation techniques (256- QAM). The new modulation technique packs more symbols (data) into the available bandwidth and hence increases the data rates. However, the higher data rates are not sustainable over long distances — the data rate drops significantly as clients move away from the AP. This can benefit high density deployments if it is feasible to create small cells with high data rates.

Another important 802.11ac feature is Multi User MIMO. This feature enables the AP to communicate with multiple clients simultaneously. It also requires clients to implement a feedback mechanism. This feature will only be available in second-generation chipsets.

Conclusion

An effective density design depends on a lot of factors and is complex. A successful network design is dependent upon:

- The number and types of clients, as well as SLAs
- Data rates
- Application types and required bandwidth
- Mixed mode client ratios
- The impact of co-channel and adjacent channel interefence.

By factoring these elements into the network design, it is possible to create a network design that can support the densest deployments.

802.11AC DOES INCREASE THE DATA RATES TO 1.3 GBPS IN THE FIRST GENERATION CHIPSETS, HELPING TO INCREASE NETWORK CAPACITY FOR HIGH DENSITY DEPLOYMENTS. BUT LIKE 11N, THIS COMES AT PRICE — CHANNELS ARE SACRIFICED.

FOR MORE INFORMATION ON DESIGNING HIGH-CAPACITY WLANS, PLEASE VISIT WWW.ZEBRA.COM/WLAN OR ACCESS OUR GLOBAL CONTACTS DIRECTORY AT WWW.ZEBRA.COM/CONTACT