
VALIDATED REFERENCE DESIGN



VERY HIGH-DENSITY 802.11ac NETWORKS

Theory Guide

Version 1.0

Chuck Lukaszewski, CWNE #112
Liang Li



Copyright

© Aruba Networks, Inc. All rights reserved. Aruba Networks®, Aruba Networks™ (stylized), People Move Networks Must Follow®, Mobile Edge Architecture®, RFProtect®, Green Island®, ClientMatch®, Aruba Central®, Aruba Mobility Management System™, ETips™, Virtual Intranet Access™, Aruba Instant™, ArubaOS™, xSec™, ServiceEdge™, Aruba ClearPass Access Management System™, AirMesh™, AirWave™, Aruba@Work™, Cloud WiFi™, Aruba Cloud™, Adaptive Radio Management™, Mobility-Defined Networks™, Meridian™ and ArubaCare™ are trademarks of Aruba Networks, Inc. registered in the United States and foreign countries. Aruba Networks, Inc. reserves the right to change, modify, transfer or otherwise revise this publication and the product specifications without notice. While Aruba Networks uses commercially reasonable efforts to ensure the accuracy of the specifications contained in this document, Aruba Networks will assume no responsibility for any errors or inaccuracies that may appear in this document.

Open Source Code

Certain Aruba products include Open Source software code developed by third parties, including software code subject to the GNU General Public License (“GPL”), GNU Lesser General Public License (“LGPL”), or other Open Source Licenses. The Open Source code used can be found at this site:

http://www.arubanetworks.com/open_source

Legal Notice

ARUBA DISCLAIMS ANY AND ALL OTHER REPRESENTATIONS AND WARRANTIES, WEATHER EXPRESS, IMPLIED, OR STATUTORY, INCLUDING WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, NONINFRINGEMENT, ACCURACY AND QUET ENJOYMENT. IN NO EVENT SHALL THE AGGREGATE LIABILITY OF ARUBA EXCEED THE AMOUNTS ACUTALLY PAID TO ARUBA UNDER ANY APPLICABLE WRITTEN AGREEMENT OR FOR ARUBA PRODUCTS OR SERVICES PURSHASD DIRECTLY FROM ARUBA, WHICHEVER IS LESS.

Warning and Disclaimer

This guide is designed to provide information about wireless networking, which includes Aruba Network products. Though Aruba uses commercially reasonable efforts to ensure the accuracy of the specifications contained in this document, this guide and the information in it is provided on an “as is” basis. Aruba assumes no liability or responsibility for any errors or omissions.

ARUBA DISCLAIMS ANY AND ALL OTHER REPRESENTATIONS AND WARRANTIES, WHETHER EXPRESSED, IMPLIED, OR STATUTORY, INCLUDING WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, NONINFRINGEMENT, ACCURACY, AND QUIET ENJOYMENT. IN NO EVENT SHALL THE AGGREGATE LIABILITY OF ARUBA EXCEED THE AMOUNTS ACTUALLY PAID TO ARUBA UNDER ANY APPLICABLE WRITTEN AGREEMENT OR FOR ARUBA PRODUCTS OR SERVICES PURCHASED DIRECTLY FROM ARUBA, WHICHEVER IS LESS.

Aruba Networks reserves the right to change, modify, transfer, or otherwise revise this publication and the product specifications without notice.



1344 CROSSMAN AVENUE | SUNNYVALE, CALIFORNIA 94089
1.866.55.ARUBA | T: 1.408.227.4500 | arubanetworks.com

Table of Contents

Chapter T-1:	Introduction	5
Chapter T-2:	What Is “The Channel?”	6
	Becoming Aware of Different Meanings of “Channel”	6
	Definition of the Channel Entity	7
	Collision Domain Properties	8
	More Comprehensive Collision Domain Model	9
	Putting the Model to Use	11
	Understanding Rate Efficiency	11
	Understanding Payload Domain vs. Collision Domain	12
	Understanding Time Efficiency and Utilization	12
	There Is No Spoon	13
	Multiple Collision Domains	13
	Overlapping Collision Domains	14
	Collision Domains of Stations	14
	Take the Red Pill	15
Chapter T-3:	Understanding Airtime	16
	What is Airtime?	16
	Airtime Structure	17
	Data Rates for 802.11 Data MPDUs	20
	Data Rates for 802.11 Control Frames	24
	Effective TXOP Data Rate	25
	Building a Frame Time Calculator	25
	Performing What-If Analysis	26
	Effects of Arbitration	27
	Average Frame Size Measurements In Live Environments	30
	Aruba Administration Building	30
	Football Stadium	31
	What is the Relationship Between Airtime and Bandwidth?	33
	Why is Wired Bandwidth Fixed but Wireless Bandwidth Varies?	33
	Summary	34
	Bibliography	35
Chapter T-4:	How Wi-Fi Channels Work Under High Load	36
	Channel Capacity Is Inversely Proportional to Client Count	36
	Defining the Contention Premium	38

	Explaining the Contention Premium	38
	Collisions and Retries Are Not the Cause	39
	Downward Rate Adaptation Is Not the Cause	40
	Ruling Out TCP Windowing	41
	Control Frame Growth Is the Critical Factor	42
	Average Frame Size Decreases with Load	43
	Causes of Control Frame Growth	44
	MIMO Works!	45
	Per-Client Throughput	46
Chapter T-5:	Understanding RF Collision Domains	48
	How the 802.11 Clear Channel Assessment Works	48
	How Co-Channel Interference Reduces WLAN Performance	49
	How Adjacent-Channel Interference Reduces WLAN Performance	50
	ACI Interference Example	51
	Measuring the ACI Impairment	52
	Interference Radius of Energy Detect and Preamble Detect	54
	A Real World Example of 802.11 Radio Power	55
	Containing CCI By Trimming Low Data Rates Is a Myth	56
	Minimum Requirements to Achieve Spatial Reuse	57
	Controlling ACI	57
Appendix T-A:	Aruba Very High-Density Testbed	58
	Testbed Justification	58
	Testbed Design	58
	Topology	59
	Channels	60
	SSID Configuration	60
	Automation	60
	What is a Client Scaling Test?	61
	Why No High Throughput or Legacy Clients This Time?	61
	Comparing with Other Published Results	61

Chapter T-1: Introduction

Welcome to the Theory guide of the Aruba Very High-Density (VHD) Validated Reference Design (VRD). The Planning guide explained what a VHD network is, presented a structured methodology for dimensioning an end-to-end system, explained how to choose APs and antennas, and introduced the three basic radio coverage strategies that can be used. The previous guide (Engineering and Configuration) covered capacity planning, configuration, channel planning, and security architecture. That guide is intended for wireless engineers responsible for deploying 802.11 networks.

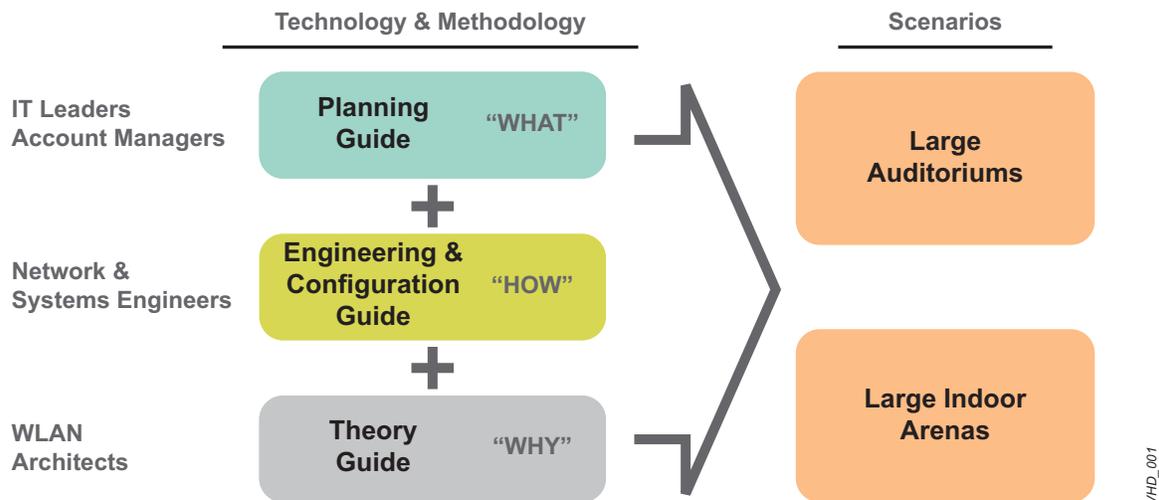


Figure T1-1 Organization of the Very High Density VRD

This guide is the most technical of the series. It is aimed at architect-level technical staff of our customers and partners, or those holding expert-level technical certifications in the wireless and networking fields. After reading the four chapters of this volume, you should be able to:

- Understand and visualize what an 802.11 channel is
- Understand, explain, and measure actual airtime consumed by 802.11 transmissions
- Understand, explain, and forecast the behavior of a VHD 802.11 channel in a range of operating load conditions
- Understand, explain, and compensate for 802.11 collision domain interference radius in your designs

Whereas the first two volumes were focused on explaining the “what” and “how” of VHD networks, this Theory guide addresses the topic of “why.” After you have fully comprehended the material in this document, you should be able to understand and explain each of the engineering and configuration recommendations made in the previous guides.

All readers should also read the appropriate Scenario document for their particular high-density use case.

Chapter T-2: What Is “The Channel?”

“You're here because you know something. What you know you can't explain, but you feel it. You've felt it your entire life, that there's something wrong with the world. You don't know what it is, but it's there, like a splinter in your mind.”

Morpheus to Neo, The Matrix (1999)

As a wireless architect, you have been explaining radio systems to others for a long time. You have drawn many circle diagrams with APs in the center to explain radio cells. You use the phrase “the channel” without thinking about it – and yet you have always known that those circles and that phrase are leaving out something important, something vital. What that might be you can't explain, and no textbook or vendor guide you have ever read has helped.

For the majority of conventional deployments this missing element doesn't seem to matter, so it's easy to ignore. But something about the performance of high-density environments you have worked on reminds you that more is happening under the surface than meets the eye.

This splinter in your mind is the role that time itself plays an 802.11 channel.

Time is an even more scarce resource than spectrum. There is never enough spectrum to be sure. But time cannot be rewound, and inefficiently used airtime is wasted capacity that can never be recovered. Wasted airtime can be the difference between success or failure in VHD design, or at least between an average performance and a great performance.

Your experience as a wireless architect has taught you to “see” radio. You can look at any environment and instantly know where to place radios and how the resulting antenna patterns will propagate. But the radio coverage is merely a flat one-dimensional view. You must also learn to “see” time to achieve a true multidimensional picture of an 802.11 system. With this enhanced vision, you will build faster, more robust WLANs. More importantly, you will be able to make entirely new arguments when third parties want to take your VHD system in a direction that you know will be harmful for all concerned.

This Theory guide covers a range of topics that are essential to take your architectural knowledge to a new level. But these topics ultimately boil down to airtime, and in particular, the effect of airtime conflicts between radio cells on the same channel or center frequency.

Becoming Aware of Different Meanings of “Channel”

The word “channel” appears 708 times throughout these VRDs. Sometimes “channel” is used in the context of a particular slice of the frequency spectrum that is allocated for Wi-Fi® use, such as channel 6 or 149. Every network engineer is familiar with this usage and instinctively understands it.

When referring to blocks of spectrum, the book uses phrases such as “9 channels,” “21 channels,” or “the 5-GHz band.” “Channel bonding” falls in the same category. When we discuss the regulatory rules that apply to specific spectrum, we use the phrases “DFS channel” and “non-DFS channel.” Again, it is fairly clear that these references are to a particular frequency range somewhere between 2 GHz and 6 GHz.

However, equally often in the guides of this VRD, “channel” is been used to describe a definite entity with specific properties and performance characteristics. Some examples include:

- ...that is because the **capacity of the channel** actually decreases as the number of clients increases...
- ...the baseline assumption for any high-density network is that **the channel is very congested**...
- ...the term “**average channel throughput**” in formula (1) is meant to capture all of these effects for a given environment...
- ...in a conventional deployment, when a new interference source is detected that **degrades channel quality**...

What exactly is this entity called “the channel”? Clearly we are not referring to spectrum, at least not in the direct sense. Is the channel entity a real, physical thing or an abstract concept? Where are its boundaries? Are they fixed or fluid? Why does a channel have any properties at all beyond its bandwidth? How are these properties to be measured?

The purpose of this entire volume is to help you develop an intuitive understanding of the channel entity and the answers to these questions. Almost every aspect of the theory behind VHD WLAN performance ultimately boils down to this construct called “the channel.” So we begin by carefully defining exactly what is meant when the term is used in this way.

Definition of the Channel Entity

Simply put, the channel entity is an 802.11 collision domain.

What is a collision domain? As always, the details are critical to understand:

- A collision domain is an independent block of capacity in an 802.11 system.
- A collision domain is a physical area in which 802.11 devices that attempt to send on the same channel can decode one another’s frame preambles.
- A collision domain is also a moment in time. Two nearby stations on the same channel do not collide if they send at different times.
- Finally, collision domains are dynamic regions that are constantly moving in space and time based on which devices are transmitting.

The concept of a collision domain is specific to the 802.11 MAC layer. All radio systems can interfere with one another if two transmitters attempt to send at the same time on the same frequency. However, 802.11-based technologies are unique because they apply carrier-sense multiple-access with collision avoidance (CSMA/CA). As you probably know, the collision avoidance mechanism uses a virtual carrier sensing mechanism as well as a physical energy detection mechanism. What you may not be aware of is the role that frame preambles play in the virtual carrier sense, and therefore the true shape of the collision domain in both space and time.

Do not use the word “cell” as a synonym for collision domain. Cells are typically engineered areas where the SINR or RSSI exceeds a specific target value. The so-called “cell edge” is the radial distance from an AP at which this value is hit. However, the collision domain extends until the SINR goes below the preamble detection (PD) threshold. The area of the cell is far smaller than the area of the collision domain.

By convention, collision domains are normally drawn as a circle around an AP that contains a number of clients, like [Figure T2-1](#), which we introduced in [Chapter EC-2: Estimating System Throughput](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide:

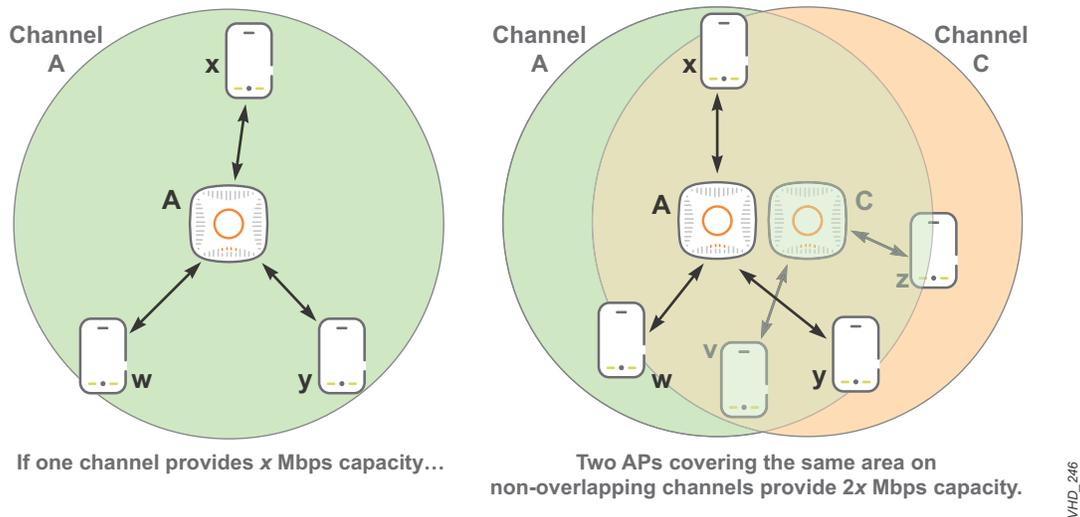


Figure T2-1 *Simplified Collision Domains*

Of course, such diagrams are vastly oversimplified and ignore many complexities of real radio cells. Though this diagram is adequate for most discussions about Wi-Fi, it is completely inadequate for our theory conversation. In particular, this figure shows only one dimension: distance from the AP. In this guide, we are interested primarily in two much more important factors: airtime and data rate. Therefore, we need a much richer model of a collision domain.

Collision Domain Properties

To construct a more complete view of an 802.11 collision domain, we start by defining three critical properties:

- Time
- Data rate
- Range

Time is the linear flow of time inside the physical area that is covered by the collision domain. Truly independent collision domains on the same radio channel also have independent time flows. The sending and receiving station pairs in each domain can transmit at the exact same nanosecond without being blocked by the other pair. Therefore, their airtime is independent.

Data rate is the speed of a particular transmission during a specific time slot. Data rate depends directly on the signal-to-interference-plus-noise ratio (SINR) that is measured at the receiver. In 802.11ac, data rate is expressed as a modulation and coding scheme (MCS) value from 0 to 9.

Range is the physical distance between a sending and receiving station. It is also expressed in terms of SINR. In this case, the “edge” of a collision domain is the SINR needed to decode the Signal field (L-SIG) of the Legacy Preamble, which must be sent using Binary Phase Shift Keying (BPSK) modulation. Range is also called the PD distance. BPSK requires an SINR of 4 dB. Preambles that fall below this value become noise.

Therefore, the physical edge of any collision domain is always determined as the distance at which the SINR is equal to 4 dB. The distance is less if there are impairments like walls, structures, or human bodies.



Collision domains have many other properties, including channel width and channel model. We ignore these properties for now.

More Comprehensive Collision Domain Model

Figure T2-2 is a new diagram of an 802.11 collision domain using these properties.

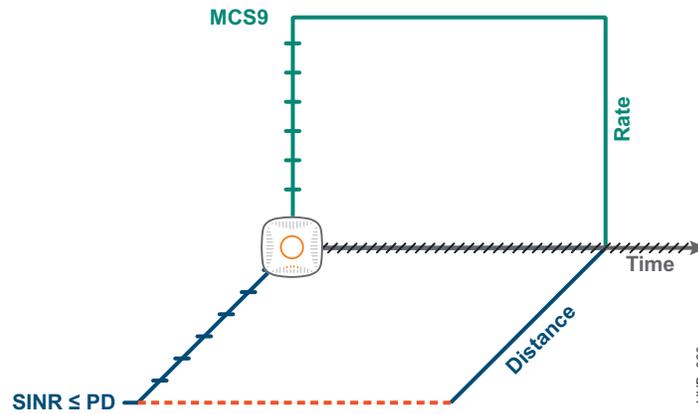


Figure T2-2 Multidimensional Model of Collision Domain

In Figure T2-2, these properties are laid out on three different axes. Assume that the AP is at the intersection. Time moves from left to right, and continues in perpetuity. Data rate is expressed on the vertical axis from MCS0 to MCS9. (We are ignoring spatial streams for now.) Finally, distance is shown on the back-to-front axis. Distance is expressed in SINR, and it stops where the SINR drops below the PD minimum.

The relationship between SINR and MCS value is well understood. This relationship can be calculated from the data sheet of any AP vendor, and it typically has an exponential shape to it due to the r^2 nature of radio signal decay. In Figure T2-3, the diagram is redrawn to show the data rate ceiling across the range of distance to the AP.

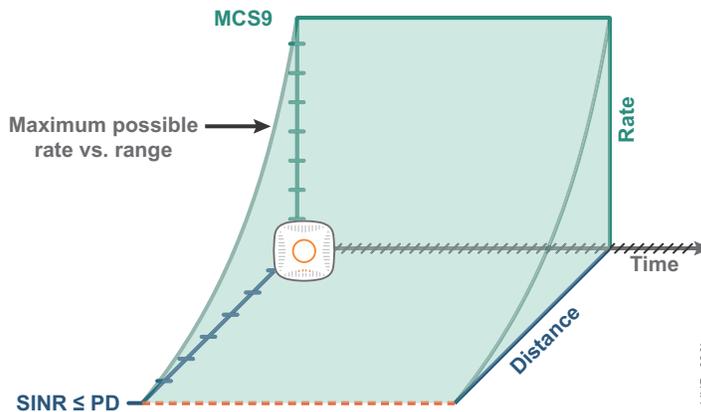


Figure T2-3 Collision Domain Model Showing Maximum Data Rate

Now that you have the basic idea, we add details to the model as shown in [Figure T2-4](#). If we assume that the AP is at the intersection of the three axes, then we are only showing half of the coverage of the AP. So the distance and data rate ceiling must be drawn in the other direction. (Of course the cell radiates in all directions, but in this approach only two are shown.) The distance between the two cell edges is the collision domain, which aligns with the area-based definition given earlier.

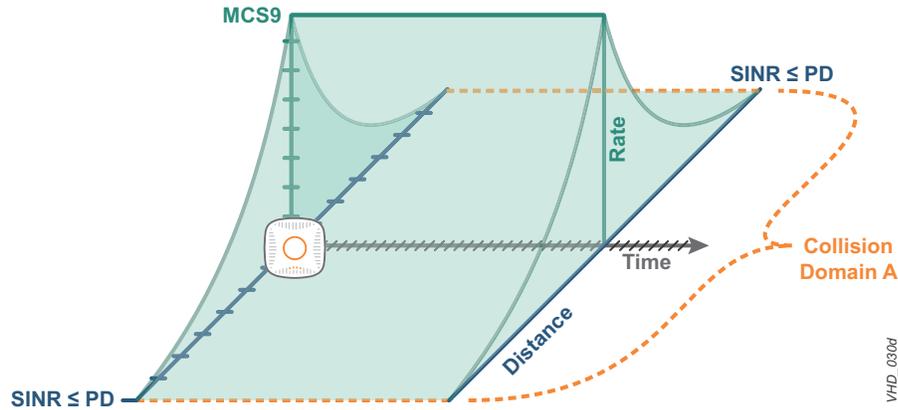


Figure T2-4 Adding Omnidirectional Coverage to Collision Domain Model

Finally, we must add clients to the “cell”. How shall the clients be placed now that time is part of the model? The answer is to show the position of clients on the distance axis as they gain control of the channel over time. [Figure T2-5](#) adds these complexities.

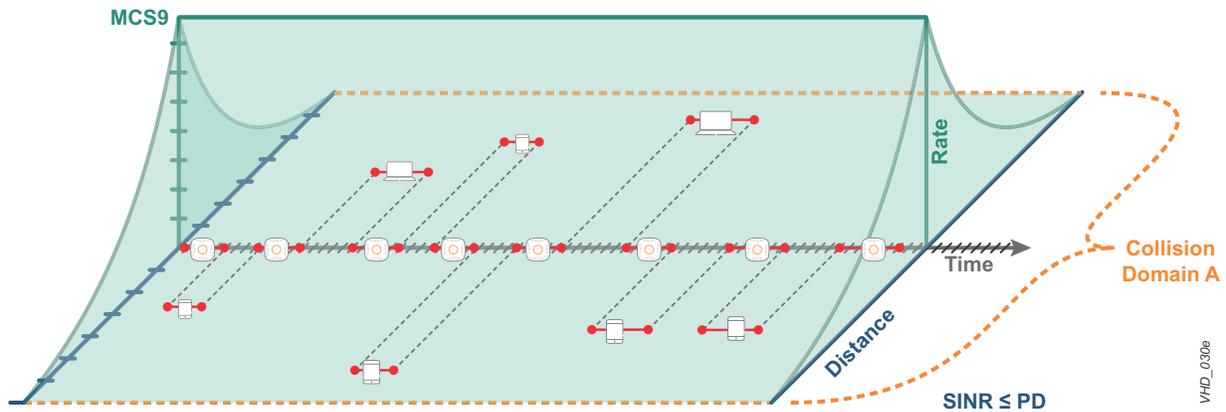


Figure T2-5 Complete Collision Domain Model with Clients

To keep things reasonably simple for now, we are intentionally ignoring the fact that the collision domain actually is dynamic. It is constantly shifting in space and time based on which devices are transmitting.

At this point, you may think that the circle model is much easier, despite its many simplifications. But if you want to understand what is actually going on in VHD environments, we must find a way to add airtime and data rate to the picture.

Putting the Model to Use

This multidimensional model of a collision domain is a tool that can be used for a variety of practical purposes.

Understanding Rate Efficiency

As a general rule, every transmission in a VHD collision domain should use the maximum possible rate for all three 802.11 frame types: data, control, and management.

Figure T2-6 is a 2D slice of the model, which focuses on the data rate and distance axes. The vertical axis are the 802.11 legacy and MCS rates grouped by the modulation they share in common. The horizontal axis shows distance from the AP.

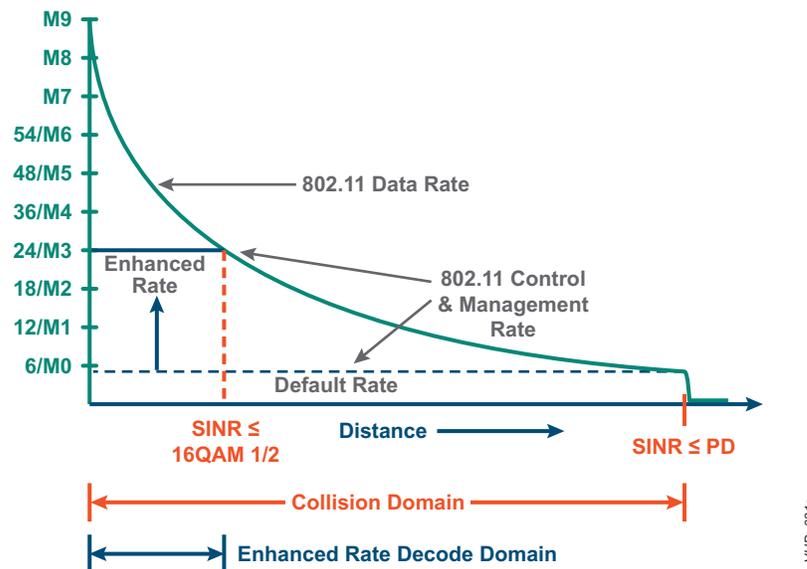


Figure T2-6 Using the Collision Domain Model to Understand Data Rate Efficiency

This chart shows several important points:

- The average PHY data rate that is used for data frames between any station (STA) and an AP should follow the rate curve shown in green. If the rate seen on air is less than expected, this indicates an operational problem or an issue with the system design.
- The average PHY data rate used for control frames should be pushed as high as it will reliably go. Do not accept the default values in VHD environments. Figure T2-6 shows a dotted blue line for the default 6 Mbps setting and a solid blue line for a 24 Mbps setting (16-QAM modulation with $\frac{1}{2}$ coding).
- The average PHY data rate that is used for management frames should likewise be pushed much higher than the defaults for the same reasons.

As you think about your SSID rate configurations, you always want to push the rate used as high as possible toward the allowable limit on the curve. [Chapter EC-3: Airtime Management](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide discussed this issue in depth across many different types of 802.11 transmissions.

Understanding Payload Domain vs. Collision Domain

Figure T2-6 also corrects a common misunderstanding among WLAN engineers and architects.

The conventional wisdom is that as the data rate for control and management frames is increased, the cell size “shrinks”. A higher SINR is required to decode the faster rate, so that payload is not decodable beyond a specific point.

This same thinking is behind the common practice of “trimming out” low OFDM data rates.

However, the chart clearly shows that if the payload rate is changed, that change **does not** alter the interference range of the legacy preamble detection. Those preambles must be sent using BPSK and they cannot be changed. So the collision domain size is **unaffected**. Distant STAs that decode the preamble still mark the channel as busy for the full duration of the frame even if the payload cannot be recovered.

By the way, trimming out low control and data rates does have many practical benefits, which are discussed at length in [Chapter EC-3: Airtime Management](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide. But trimming those rates does not change the size of the collision domain.

Understanding Time Efficiency and Utilization

Rate efficiency directly affects airtime efficiency and channel utilization. To visualize this concept, [Figure T2-7](#) takes a different 2D slice of our model and focuses on the rate and time axes.

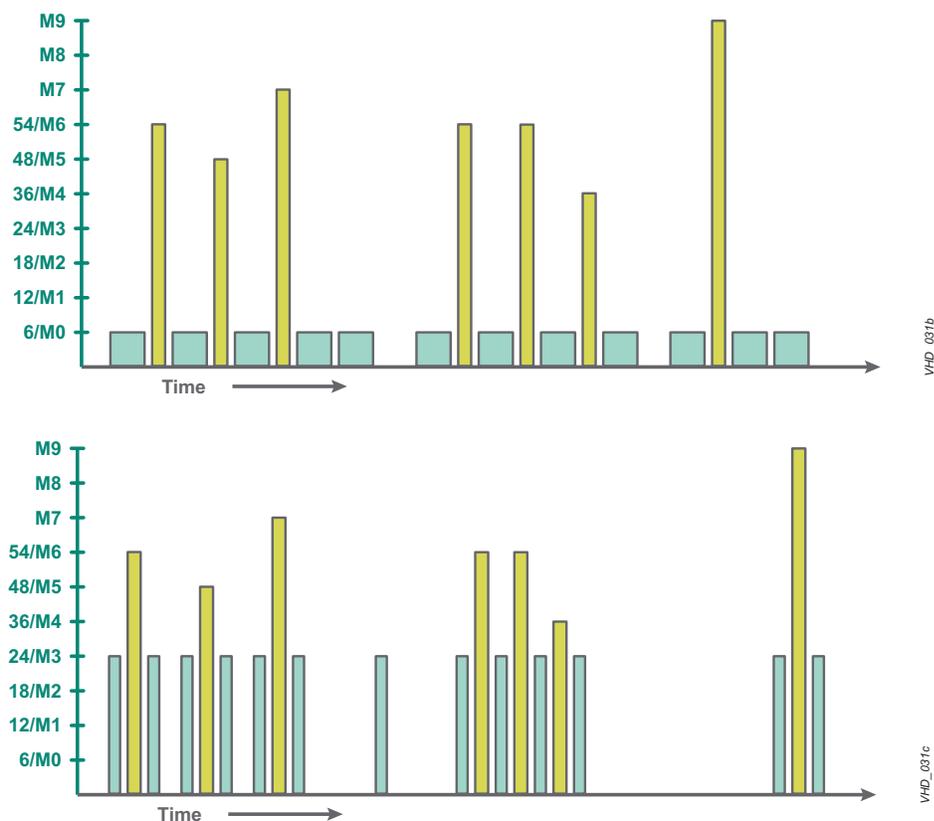


Figure T2-7 Creating Capacity By Using Higher Rates to Increase Airtime Efficiency

This highly oversimplified view is meant to show the relative time consumed by the same sequence of data packets with two different control frame rates. On the top is a default 6 Mbps rate, and on the bottom, the same sequence is shown, but using a 24 Mbps rate.

Chapter T-3: Understanding Airtime explains data rates in detail. For now, you need only know that a control frame of a given size will take 4 times longer to send at the default rate than at a 24 Mbps rate. This concept is manifested in the timeline view, which effectively shows channel utilization. When you raise the control rate, each station gets off the air faster. Capacity is created by increasing the idle time during which the channel is free for other users.

There Is No Spoon

It must be stressed that this multidimensional collision domain model is just that. Like the circular cell drawings, this model does not exist in a physical sense, although every 802.11 radio cell does operate according to these principles. And like all models, it intentionally simplifies a more complex reality. The purpose of the model is to give you a mental framework to begin to understand the interdependency of time, data rate, and range, and to begin to “see” the time dimension.

Multiple Collision Domains

We defined an 802.11 collision domain as the physical area in which two 802.11 stations can decode one another’s legacy preambles. A secondary definition is that time flows independently in each collision domain on the same channel. Our multidimensional model can be extended to show this by adding a second collision domain next to the first.

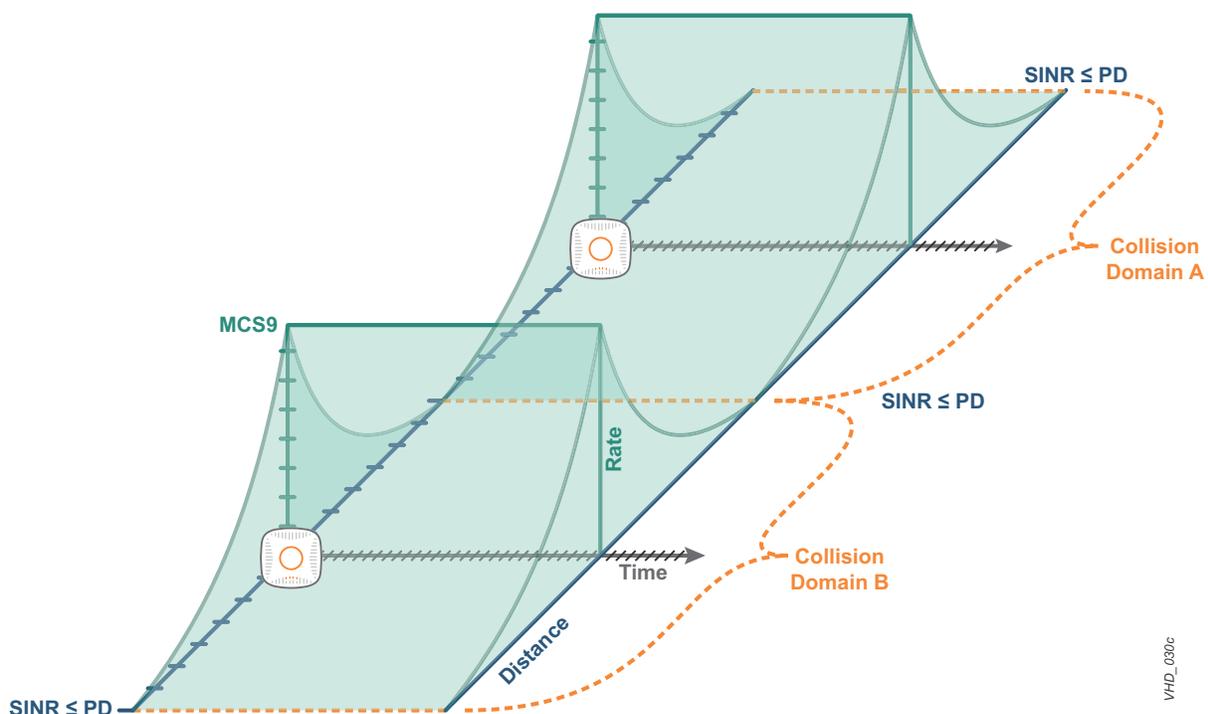


Figure T2-8 Two Non-Overlapping 802.11 Collision Domains

Imagine a second AP on the same channel at the new axis intersection, with its own PLCP PD interference radius. In this example, the cells have been spaced with the unrealistic assumption that this point is exactly halfway between the APs. The APs would have to be several hundred meters apart in free space for their PD distances not to overlap at all, or somewhat less if there are walls or other structural materials in the way. In this case, each AP has a truly independent collision domain. It can be said that the “channels”

are independent from a capacity perspective (even though they are on the same exact center frequency). This situation of course is the ideal of every dense network design.

Overlapping Collision Domains

However, if these APs overlap on the distance axis to any extent, then they are no longer independent in time in that region. [Figure T2-9](#) shows the far more common case of an enterprise deployment with a -65 dBm cell edge and approximately 20 m (65 ft) AP-to-AP spacing. This would be the case in a three-channel plan in 2.4-GHz.

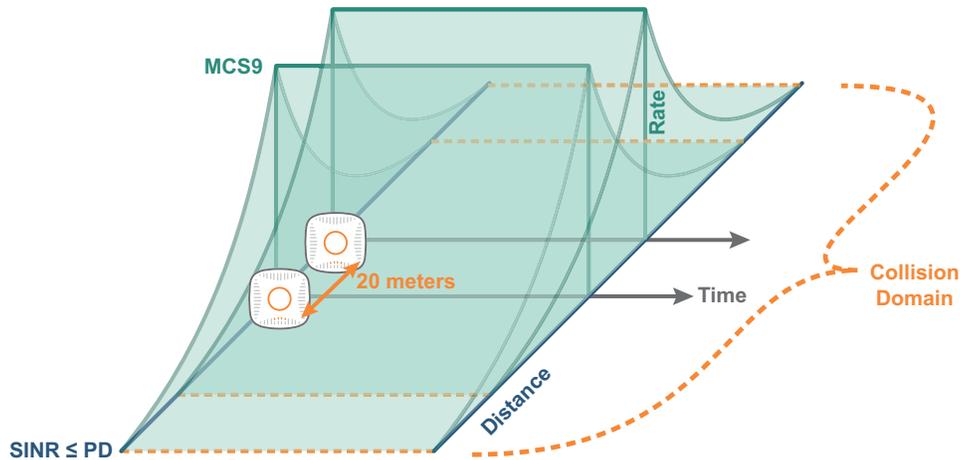


Figure T2-9 Overlapping Collision Domains Are One Channel

Here again the model shows its value as it clearly shows that these two cells are in effect a single collision domain. Even if the payload rates of data and control frames have been increased according to the best practices of this VRD, the two cells are one collision domain. Radio signal power decays exponentially, so it falls off quickly at first, but then proceeds for a very long distance.

Aruba knows that many customers believe that their network behaves more like [Figure T2-8](#), when in fact it is more like [Figure T2-9](#). Wireless architects must set proper expectations with customers when they design any WLAN, but especially VHD systems.

In [Chapter EC-2: Estimating System Throughput](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide you learned to perform capacity planning using the total system throughput (TST) methodology. The *Reuse Factor* term in the TST formula is a *measure of collision domain overlap*. A low value of 1 indicates that all same-channel APs exist within the same collision domain and are therefore a single channel from a capacity perspective. Higher values imply an expectation that there is some degree of independence. As has been stated many times in this guide, it is virtually impossible to obtain collision domain independence in VHD environments of 10,000 seats or less, even when specialized antennas and mounting strategies are used.

Collision Domains of Stations

One obvious simplification in this model is that it considers only the collision domain of the AP. We have not considered the problem of STAs in between the two APs. STAs must follow the same rules as APs, and therefore their collision domains are also relatively large. (The exception is that some STAs use reduced transmit power to increase battery life.)

The truth is that if we look at all of the APs and STAs on a given channel frequency in a modern dense WLAN, it is completely impossible to draw definite boundaries. Collision domains are *relative* to the

transmissions in progress on that specific center frequency at that specific instant in time. To visualize this point, Let us revise Figure T2-4 to show a single collision domain with the relative instantaneous shape of the space-time collision domain based on which device currently controls that channel.

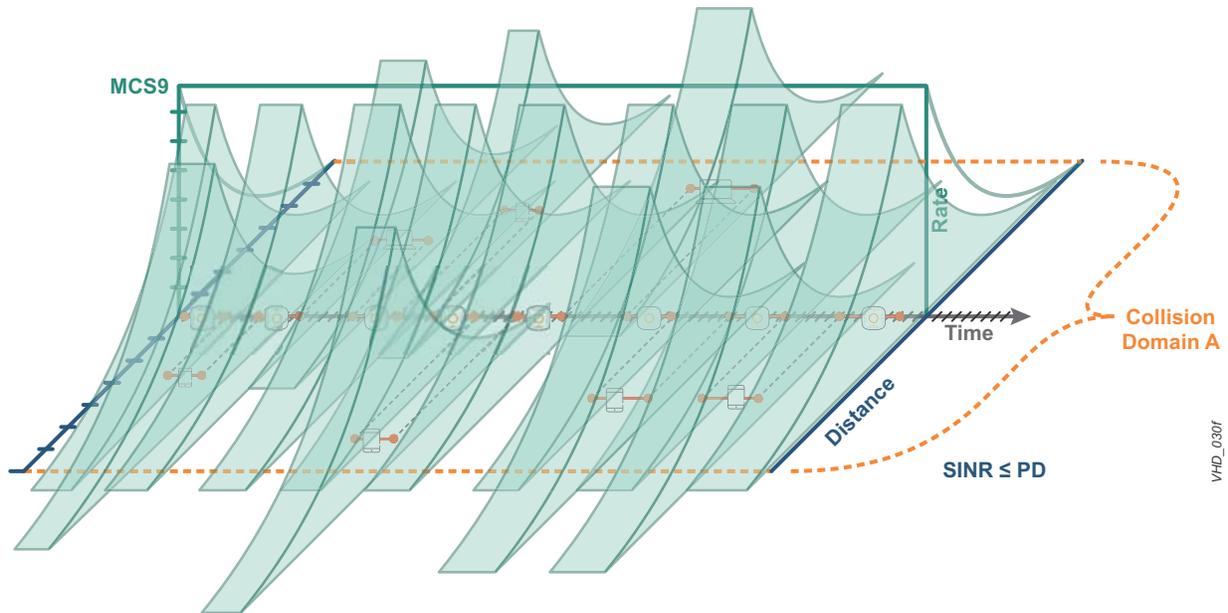


Figure T2-10 Dynamic Collision Domain Model with STAs and APs

Again, time plays an important role. From moment to moment, collision domains split apart or merge together depending on which AP or STA has won the channel during the arbitration process.

Figure T2-10 might seem confusing or overly complicated, but this is precisely what is happening to the collision domain on that channel from moment to moment. Collision domains are constantly changing in both space and time. Remember that the goal of the preceding exercise simply is to expand your vision and understanding of the mechanics of these environments. When you can see radio in both space and time, it is a simple matter to apply that awareness to specific physical facilities.

Take the Red Pill

The remainder of this guide describes each of the three axes of the collision domain model in great detail.

[Chapter T-3: Understanding Airtime](#) explores the concept and reality of airtime. [Chapter T-4: How Wi-Fi Channels Work Under High Load](#) looks at the efficiency of data rates, especially the impact of control and management traffic in VHD areas. [Chapter T-5: Understanding RF Collision Domains](#) considers the physical boundaries of collision domains beyond the simplified model that we just explored.

Even the most experienced WLAN architect will be surprised by some of the material presented in these chapters. After you read the whole guide, you will never again look at 802.11 in quite the same way.

"This is your last chance. After this, there is no turning back. You take the blue pill - the story ends, you wake up in your bed and believe whatever you want to believe. You take the red pill - you stay in Wonderland and I show you how deep the rabbit-hole goes."

Morpheus to Neo, The Matrix (1999)

Chapter T-3: Understanding Airtime

This chapter builds on the foundation laid in the last chapter by studying airtime and techniques that you can employ to control it to your advantage. In [Chapter P-3: RF Design](#) of the *Very High-Density 802.11ac Networks Planning Guide*, we stated that one of the four over-arching radio design responsibilities of a wireless architect in a very high-density (VHD) network is to protect every microsecond of airtime on every available channel from being used unnecessarily or inefficiently. By the end of this chapter you will understand why protecting airtime is so important.

The critical question that this chapter seeks to answer is this: *to send any arbitrary amount of data payload, what is the true price that must be paid in airtime?* Everyone knows that it takes time to send data; that is not in dispute. But you may be surprised by the magnitude of the price. When you know the true cost, you will want to learn airtime management techniques to reduce it.

If you succeed at RF design and fail at airtime management, your VHD network will likely fail to meet capacity expectations. Conversely, with good airtime management, even an suboptimal RF design can carry a significant amount of traffic.

What is Airtime?

Mastery of airtime begins with a clear idea of exactly what airtime is.

At the highest level, 802.11 airtime can be thought of a continuous series of alternating idle and busy periods on a given channel (or collision domain). The length of each period is measured in a unit of time, such as milliseconds (ms) or microseconds (μ s). When viewed at this high level, the time required for each period varies constantly.

[Figure T3-1](#) shows alternating idle and busy periods on three adjacent, nonoverlapping channels.

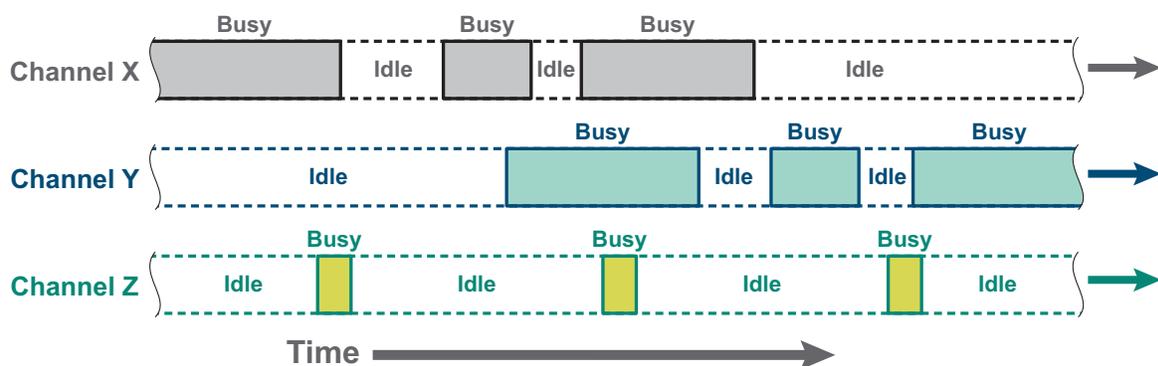


Figure T3-1 Alternating Idle and Busy Periods on Three Channels

In the time period covered by the diagram, channel X is in the middle of a transmission on the left, and then ends up in an idle state. Channel Y begins in idle and ends busy. Channel Z is idle except for periodically repeating transmissions, such as an access point (AP) beacon.

Airtime Structure

Now we zoom in and focus on how a single 802.11 radio channel is organized at a more granular level in the MAC layer. The PHY layer has an even finer structure that is not relevant to this discussion.

Wi-Fi employs a technology called carrier sense multiple access with collision avoidance (CSMA/CA or just CSMA) to order the channel. With CSMA, airtime is divided into busy units called transmit opportunities (TXOPs) and idle time. Idle time is further broken down into arbitration periods and truly idle time (when no device has anything to send). Stations contend with one another to gain control of the channel in a process called arbitration. There is no central scheduler. The station that wins the arbitration process becomes the TXOP holder and has exclusive use of the channel up to the TXOP limit.



Figure T3-2 Two Types of Idle Time In 802.11

This process is quite different from the many centrally-scheduled radio technologies that are based on some form of time division multiplexing (TDM). With TDM, the airtime is broken down rigidly into very short blocks and allocated between users by a master scheduler using precisely synchronized clocks. TDM is common in cellular networks, which have exclusive use of specific spectrum and control over both the base stations as well as the mobile terminals. TDM is also found in long-haul wireless bridges. By contrast, Wi-Fi uses unlicensed frequency bands that are open to anyone, which necessitates a standardized negotiation process for medium access.



NOTE

This chapter is intentionally focused on data transmissions, and therefore TXOPs. 802.11 has many other control and management transmissions that still require arbitration and acknowledgments, but do not use the TXOP format. When you are clear on the airtime consumption of TXOPs, the immense airtime impact of these other frame types will be self-evident.

Arbitration

Arbitration is necessary because on any given channel, only one station can transmit at the same time within the same 802.11 collision domain. In addition, not all stations are equal. Wi-Fi includes quality of service (QoS) capabilities that allow for up to four different transmission priority queues (Voice, Video, Background, and Best Effort). These queues are enforced via the arbitration mechanism.

For readers who are new to arbitration, it is a form of interframe space prior to a TXOP where stations compete to become the TXOP holder. This period has two parts:

- **Arbitration interframe space (AIFS):** variable, but fixed within a class of service (CoS) (from 34 to 79 μ s)
- **Contention window (CW):** variable based on CWmin/CWmax and CoS (from 0 to 9,207 μ s)

If no station has data to send, then no timers are decrementing and the channel truly is idle. Arbitration actually begins when a station with data to send begins counting down its AIFS timer. Each station that is preparing to transmit chooses the AIFS duration that is appropriate for the CoS of the data that it has to send. If it is voice data, it uses a [VO] CoS. The overwhelming majority of data sent in VHD areas uses an unprioritized [BE] CoS.

After the AIFS timer expires, the CW timer begins. Each station chooses a random number of “slot times” that must elapse before it can begin to send. If someone else’s timer expires first and the channel goes busy, then the station stops counting until the next arbitration period. If the station counts to zero and the 802.11 clear channel assessment (CCA) reports the medium is idle, then the station turns on its radio and begins to transmit. When the channel goes idle again, the remaining stations that have data to send can resume their CW timer countdowns from where they left off.

Figure T3-3 shows three different devices on the same channel, each with data to send.

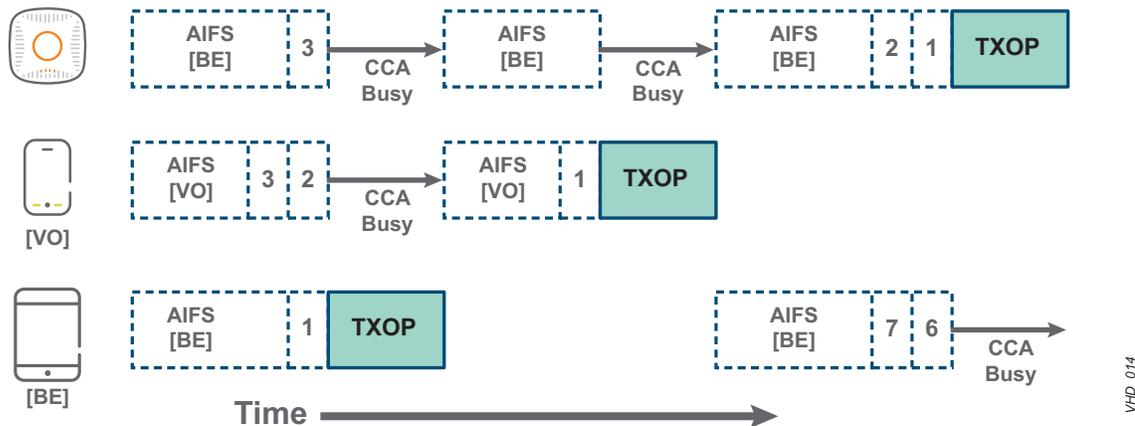


Figure T3-3 Arbitration Between Three QoS Stations in 802.11

The dashed boxes represent the CSMA arbitration period that precedes every transmission in Wi-Fi. In this example, the tablet wins arbitration and is able to send first (even though it is in the [BE] queue). A new arbitration period then begins, which is won by the smartphone. Finally the AP wins. This process continues indefinitely for all clients in a collision domain. The TXOP time durations are not drawn to scale.



For a much more complete discussion of arbitration, see one of the textbooks listed in the bibliography at the end of the chapter.

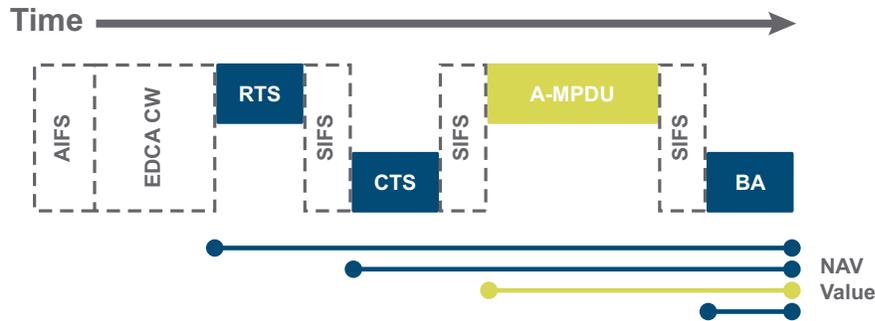
TXOP Structure

An 802.11 TXOP technically begins from the moment that any station on the channel wins arbitration. WLAN architects must understand the structure of a TXOP. In 802.11ac, all data payloads must be sent using this format. A basic 802.11ac TXOP is shown in Figure T3-4. It consists of these components:

- Ready to Send (RTS) frame preceded by an arbitration period
- Clear to Send (CTS) frame preceded by a SIFS
- Aggregated MAC protocol data unit (A-MPDU) data frame containing one or more MPDUs preceded by a SIFS
- Block acknowledgment frame preceded by SIFS

You can see that a TXOP is basically a time-limited conversation between two stations. With 802.11ac Wave 2 and multiuser multiple input, multiple output (MU-MIMO), a TXOP may expand to include up to four stations plus the AP for downstream data transmissions. In 802.11ac, virtually all TXOPs begin with an RTS/CTS exchange to allow the dynamic channel bandwidth function to sense how many subchannels are clear.

Figure T3-4 also shows how each of the four successive frame types sent during the TXOP updates the Network Allocation Vector (NAV). The NAV is the mechanism that the CCA process uses to set the virtual carrier sense to busy. The CCA process is described in depth in [Chapter T-4: How Wi-Fi Channels Work Under High Load](#).



VHD_033

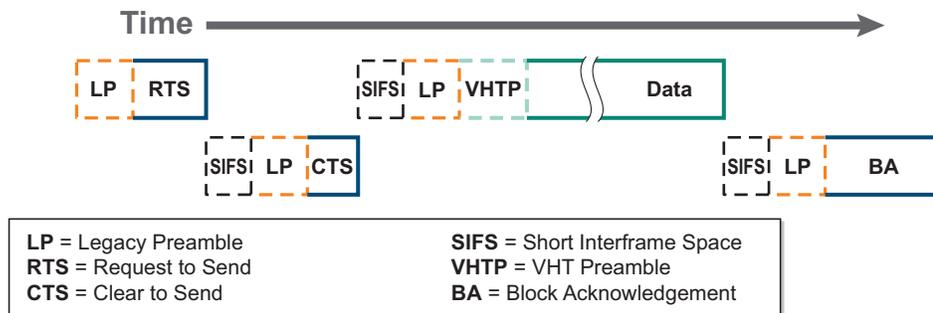
Figure T3-4 Structure of a TXOP In Time Domain



This guide does not explain 802.11 protocol operation in any greater depth. Interframe spaces, QoS access categories, frame aggregation, MU-MIMO, and other core aspects of MAC operation are outside the scope of the VRD. Numerous excellent textbooks cover these topics in great depth. A bibliography is provided at the end of this chapter. Aruba strongly recommends that WLAN architects familiarize themselves with the 802.11 protocol at this level.

Frame Preambles

Going into deeper detail, we must zoom into the airtime structure even further to discern that every 802.11 frame is actually composed of two or more parts. These parts are the preamble(s) and the payload. Figure T3-5 shows this breakdown and shows that there is more than one type of preamble.



VHD_035

Figure T3-5 Structure of a TXOP Including Preambles

The preamble is the tool that is used by the radio to bootstrap each and every frame. The preamble contains various elements used by the radio to lock onto the transmission, as well as a number of data fields that describe how the payload should be processed. There are several kinds of preambles, and the two that this guide describes are the legacy preamble (LP) and the Very High Throughput (VHT) preamble (VHTP). VHTPs are preceded by LPs to ensure compatibility with legacy stations. As you will learn shortly, preambles consume significant amounts of airtime.

From the perspective of the preamble, even control or management frames can be thought of as a form of data frame because ultimately an RTS or a CTS or a beacon is simply a payload type that consists of a fixed sequence of bytes sent at a certain rate.

Optimizing TXOPs

Without going any deeper into TXOPs, you should see that *the key to performance in VHD areas with many stations is to minimize busy time and maximize idle time*. Idle time is maximized when:

- Unnecessary TXOPs are avoided
- Necessary TXOPs are completed in as few microseconds as possible
- Retransmissions of failed TXOPs are minimized or avoided altogether

You can boil this entire VRD down into these three principles. Whether you are serving hundreds or thousands of clients, every busy period takes away capacity from someone else. As the wireless architect, you must take a ruthlessly critical view of all airtime consumption.

Your ability to deploy successful VHD networks depends on how well you understand and how firmly you enforce these principles.

Data Rates for 802.11 Data MPDUs

Understanding the structure of a TXOP does not tell us anything about how much airtime one consumes. For that, we must turn our attention to the PHY layer and data rates.

802.11ac Data Rate Table

802.11ac introduces significant additional complexity to the data rate table. Some of this complexity is obvious, and some is hidden. See [Appendix EC-B: 802.11ac Data Rate Table](#) in the *Very High-Density 802.11ac Networks Engineering and Configuration* guide for a complete list of rates up to four spatial streams.

The most obvious changes come from the addition of new 256-QAM modulations and wider channels. 802.11n had eight modulation and coding scheme (MCS) values for each spatial stream, but 802.11ac can have up to 10. However, in a few cases 802.11ac has only nine. Of particular relevance to VHD areas, MCS9 is not available for 1SS or 2SS devices in a VHT20 channel.

Each new channel width requires a full set of data rates for every MCS. Within each channel width are 400-ns and 800-ns guard intervals. The result is that, for three spatial streams, the 84 data rates that were defined in 802.11n have grown to 208 data rates with 802.11ac!

As explained in [Chapter EC-3: Airtime Management](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide, Aruba recommends using only the 20-MHz channel width to improve overall performance. One helpful by-product of this decision is that it reduces the rate table ([Table T3-1](#)) to something much easier to remember.

Table T3-1 802.11ac Data Rates for 20-MHz VHT Operation

MCS	Modulation	Bits per Symbol	Coding Ratio	1 Spatial Stream		2 Spatial Streams		3 Spatial Streams		4 Spatial Streams	
				SGI	No SGI	SGI	No SGI	SGI	No SGI	SGI	No SGI
MCS 0	BPSK	1	1/2	6.5	7.2	13.0	14.4	19.5	21.7	26.0	28.9
MCS 1	QPSK	2	1/2	13.0	14.4	26.0	28.9	39.0	43.3	52.0	57.8
MCS 2	QPSK	2	3/4	19.5	21.7	39.0	43.3	58.5	65.0	78.0	86.7
MCS 3	16-QAM	4	1/2	26.0	28.9	52.0	57.8	78.0	86.7	104.0	115.6
MCS 4	16-QAM	4	3/4	39.0	43.3	78.0	86.7	117.0	130.0	156.0	173.3
MCS 5	64-QAM	6	2/3	52.0	57.8	104.0	115.6	156.0	173.3	208.0	231.1
MCS 6	64-QAM	6	3/4	58.5	65.0	117.0	130.0	175.5	195.0	234.0	260.0
MCS 7	64-QAM	6	5/6	65.0	72.2	130.0	144.4	195.0	216.7	260.0	288.9
MCS 8	256-QAM	8	3/4	78.0	86.7	156.0	173.3	234.0	260.0	312.0	346.7
MCS 9	256-QAM	8	5/6	N/A	N/A	N/A	N/A	260.0	288.9	N/A	N/A

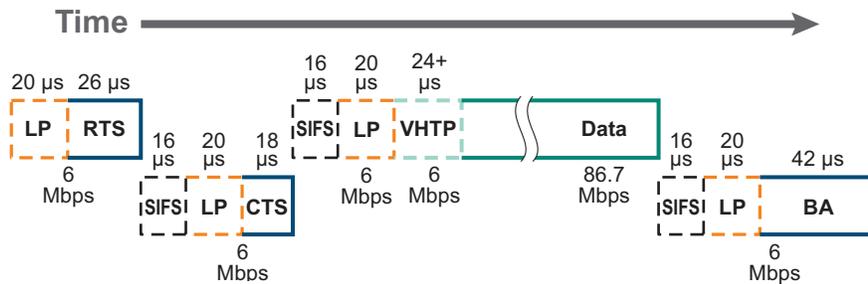
As always, the maximum data rate that can be used between an AP and a STA depends on the signal-to-interference-plus-noise ratio (SINR). The faster the data rate, the greater the SINR needed to successfully demodulate that rate. The new 256-QAM rates generally require a minimum of 30 dB to as much as 35 dB SINR. This ratio is possible only within a few meters of the radio in free space. In a VHD area packed with people, this distance can drop to just 1 – 2 m. So in practice we do not engineer for MCS8 or MCS9. We are very happy to obtain it when we can, but as you saw in [Chapter EC-2: Estimating System Throughput](#) in the *Very High-Density 802.11ac Networks Engineering and Configuration* guide, we engineer for a much lower impaired value for all users. This method leaves open the possibility of bursting much faster at some times if the channel is lightly loaded or the user is close to the AP.

Preamble Rate vs. Payload Rate

However, it is not true that an 802.11 frame is sent at one data rate. In fact, each frame that is transmitted by a Wi-Fi radio is sent at two different data rates, as shown in [Figure T3-6](#).

- Legacy and VHT preambles – Required to be sent at 6 Mbps BPSK rate
- PHY Service Data Unit (PSDU) payload – Sent at chosen data payload rate

Now we update Figure T3-5 with the TXOP structure to reflect all of the component parts and data rates. For now, we assume that the control data rate is the same as the preamble rate: 6 Mbps. Let us also ignore the AIFS and contention window for the time being. Figure T3-6 shows these changes. Note that we have also added the number of microseconds required for each individual transmission element and interframe space.



VHD_034

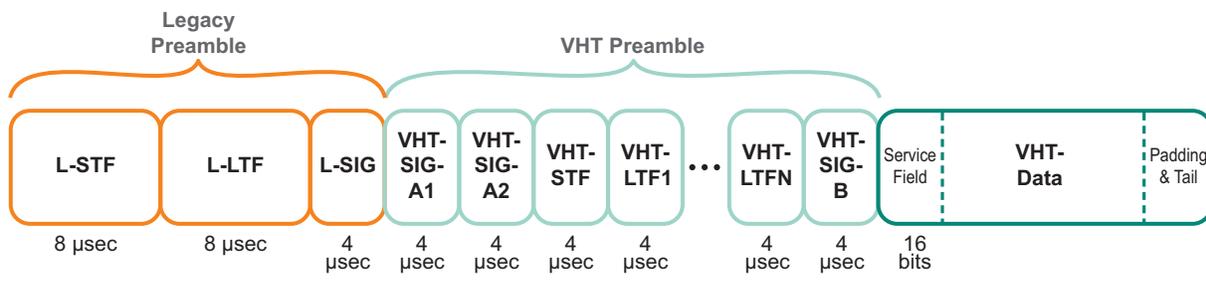
Figure T3-6 Detailed TXOP Structure with Preamble Data Rates

Study this figure carefully, especially the airtime required for each part. Even if we achieve MCS8 or MCS9 for our payload rate because our SINR is high enough, the first 20 μs or more of every single frame is consumed by the legacy preamble, which is sent at the slowest 6 Mbps rate. This data rate is hardwired into 802.11 and cannot be changed.

Preamble Airtime vs. Payload Airtime

20 μs may not sound like a lot of time, but with faster and faster data rates, the legacy preamble can actually consume more time than the payload takes to send. Particularly because the average frame size on most WLANs is no more than 500-600 bytes. The math may surprise you.

A legacy preamble requires 20 μs. A VHT preamble requires a minimum of 24 μs and could be even longer if additional long training fields (LTFs) are required. In 802.11ac, there is generally one LTF required per spatial stream. VHT frames require an LP and a VHTP for a total minimum preamble airtime of at least 44 μs.



VHD_036

Figure T3-7 Preamble Format with Symbol Durations

How does this time compare to the time required to send data payloads? It is very simple to calculate.

$$\text{Payload Airtime } (\mu\text{s}) = \frac{\text{Payload Size (bytes)} * 8 \text{ bits/byte}}{\text{Data Rate (Mbps)}}$$

Consider one of the most common MPDU frames sent on any wired or wireless network – a 90-byte TCP acknowledgment. Assume it is sent by a 1SS VHT20 device at the maximum rate of MCS8. The airtime that is required is $90 * 8 / 86.7 = 8.3 \mu\text{s}$. **Compared to the 44 μs VHTP, the preamble requires 5.3 times more airtime than the TCP ack payload!** And this time does not include arbitration or the rest of the TXOP structure.

This contrast is even more stark when viewed graphically. To highlight the magnitude of the potential spread between preamble and payload airtime, we have computed the airtime required for a range of five common MPDU payload sizes and charted it in Figure T3-8. The MPDU sizes pictured from left to right are 64, 512, 1,024, 1,514, and 3,028 bytes.

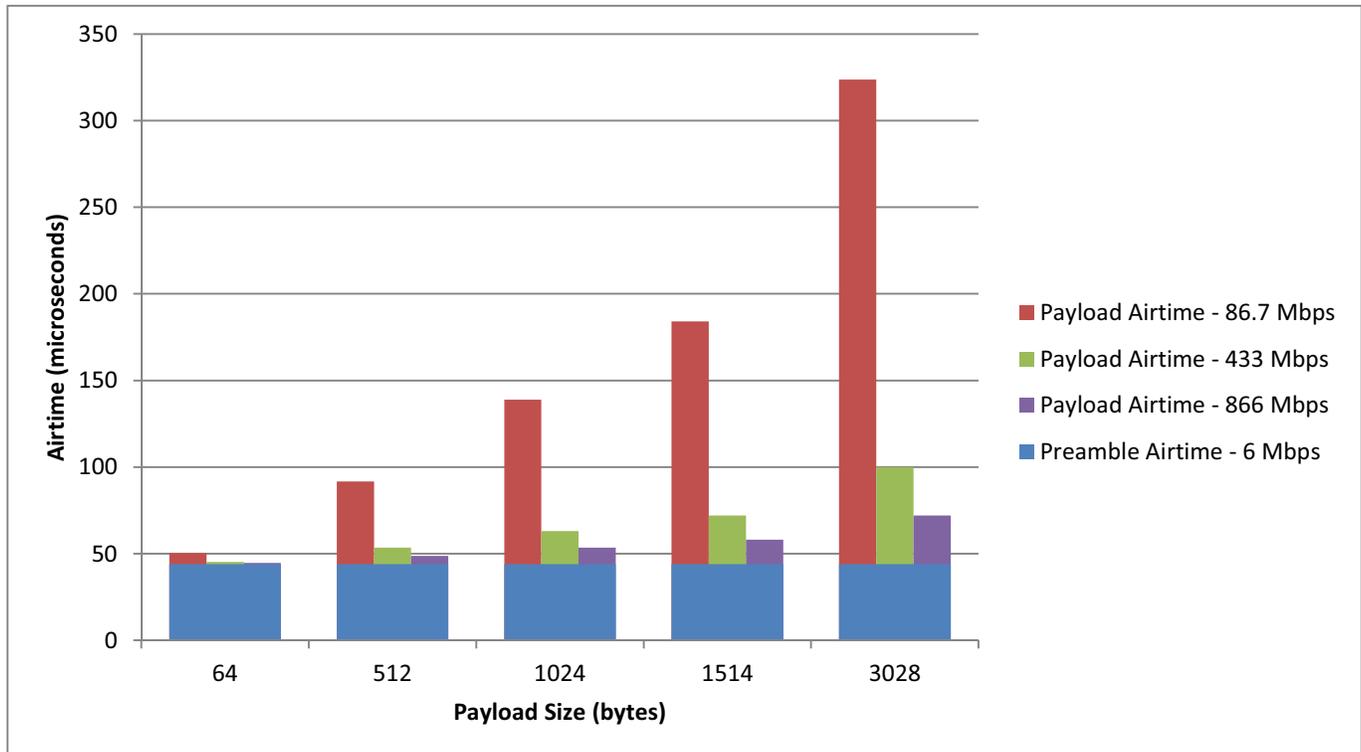


Figure T3-8 Preamble vs. Payload Airtime for Various Payloads and Rates

For each payload size, the chart shows how the payload airtime changes with three different and common data rates:

- Blue: Legacy + VHT preambles (6 Mbps)
- Red: Payload rate of 1SS VHT20 MCS8 (86.7 Mbps)
- Green: Payload rate of 1SS VHT80 MCS9 (433.3 Mbps)
- Purple: Payload rate of 2SS VHT80 MCS9 (866.6 Mbps)

The preamble time is constant at 44 μs . The payload airtime varies depending on the frame size and the data rate selected. Astonishingly, the preambles require almost 57% more airtime to send than a 3,028-byte frame at the 866 Mbps rate! For very small 64-byte frames, which are extremely common on WLANs, **the preamble towers over payload by a factor of 7X at 86.7 Mbps and by more than 70X at 866 Mbps!**

Data Rates for 802.11 Control Frames

In addition to the data MPDU, a TXOP is composed of three 802.11 control frames including RTS, CTS, and Block Ack (BA). We will calculate the airtime required for these frames in this section.

The payload portion of these control frames are sent at a default rate of 6 Mbps. In [Chapter EC-3: Airtime Management](#) in the *Very High-Density 802.11ac Networks Engineering and Configuration* guide we strongly advocated increasing this rate to 24 or even 36 Mbps in some cases. Control frames are all preceded by an LP at 6 Mbps, which requires 20 μ s. We can create a similar type of airtime chart just for these frames.

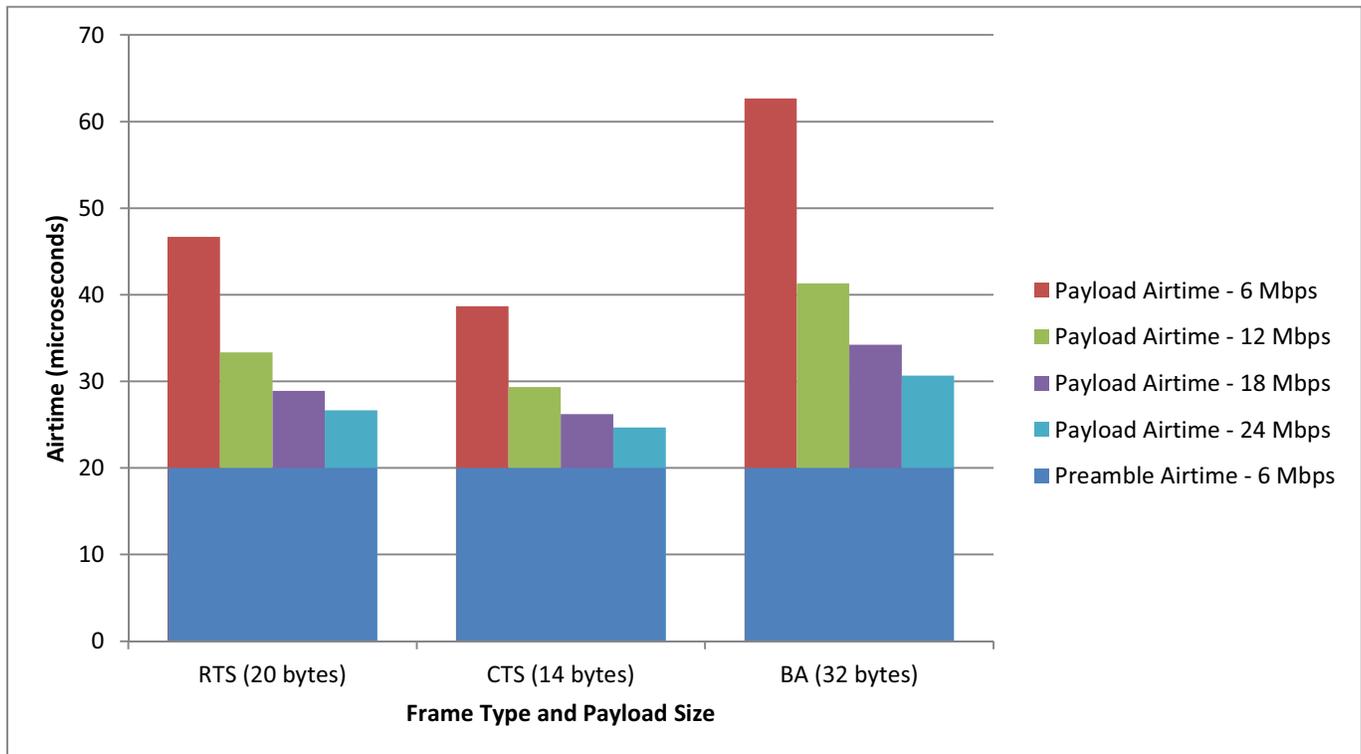


Figure T3-9 Preamble vs. Payload Airtime for Various Control Frames

An 802.11 RTS is always 20 bytes, a CTS is 14 bytes and a BA is 32 bytes. These frames must use legacy 802.11 OFDM rates for backward compatibility. [Figure T3-9](#) follows the same format as the chart for the data frames, but in this case we plot four different legacy rates. The default 6 Mbps rate is in red on the left, and the 24 Mbps rate is in light blue on the right.

The absolute magnitude of the preamble vs. payload delta is not as dire as with the data frames. This difference is solely because of the small byte size of the control frame payloads. However, when one realizes that at least one of each of these frames is required to send every MPDU, the total overhead percentage for the entire TXOP is clearly staggering for small payload sizes.

This same effect applies to beacon rates. In [Chapter EC2](#) we advocated raising beacon rates to 24 Mbps or higher. Recall the colored output from the [Table EC3-11 on page 48](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide showing the 75% reduction in airtime consumption by making this change.

The most critical takeaway you should have from [Figure T3-9](#) is *how much airtime you can recover by raising the control frame rate!* The same RTS+CTS+BA that requires 88 μ s at the default rate drops to just 22 μ s with a 24-Mbps control rate. This change yields 66 μ s savings for every single TXOP! This example is exactly

what is meant when we refer to “creating capacity” by recovering airtime. Multiplied over millions of TXOPs, this savings can become an enormous amount of extra airtime available for serving more users.

Effective TXOP Data Rate

Given that a TXOP is composed of multiple different types of frames and preambles, all going at different rates, then what is the actual “effective” data rate (EDR) of a TXOP? It must be significantly lower than the data payload MCS rate that most engineers use to talk about network performance.

Building a Frame Time Calculator

From [Figure T3-6](#) and the preceding discussion, we have enough information to construct a frame time calculator. Start by laying down each component of the TXOP in the left column in the order it occurs. Then populate the fixed duration elements including SIFS (16 μ s), LP (20 μ s), VHTP (24 μ s).

Then add rows to calculate airtime for each of the four data payloads. Add the fixed-byte totals for RTS (20), CTS (14) and BA (32). Add a dynamic field for the MPDU; here is where we will vary the payload size to do “what if” analysis. For now, assume that the data payload is 512 bytes. In 802.11ac, every MPDU is preceded by a 4-byte MPDU delimiter. Multiply the number of bytes times 8 to get bits, and then divide by the data rate.

The final step is to add a column for data rate for the five frame types with payloads. We set the control frames to the default of 6 Mbps and the data MPDU to 86.7 Mbps. If you build all this in a spreadsheet, it should look just like [Table T3-2](#).

Table T3-2 TXOP Airtime Calculator (512-B payload, 6-Mbps Control Rate)

MAC Unit	Payload Bytes	Payload Bits	Data Rate	μ sec	% Airtime
Legacy Preamble			6 Mbps	20.00	7.0%
RTS	20	160	6 Mbps	26.67	9.3%
SIFS				16.00	5.6%
Legacy Preamble			6 Mbps	20.00	7.0%
CTS	14	112	6 Mbps	18.67	6.5%
SIFS				16.00	5.6%
Legacy + VHT Preambles			6 Mbps	44.00	15.3%
A-MPDU Delimiter	4	32	86.7 Mbps	0.37	0.1%
Data Frame Payload	512	4096	86.7 Mbps	47.24	16.4%
SIFS				16.00	5.6%
Legacy Preamble			6 Mbps	20.00	7.0%
BA	32	256	6 Mbps	42.67	14.8%
Airtime for TXOP only (excluding arbitration)				287.61	100.0%
Effective TXOP rate for TXOP only (excluding arbitration)				16.2	

In this example, the entire TXOP requires 287.61 μs to send, exclusive of arbitration. Of this period, the 512 bytes of data payload take just over 16% of the airtime to send. The remaining 84% of airtime is consumed by the MAC protocol overhead and framing.

We can then calculate the EDR across the entire TXOP by dividing payload bits by total airtime like this:

$$\text{Effective TXOP Data Rate (Mbps)} = \frac{\text{MPDU Payload Size (bits)}}{\text{TXOP Airtime (}\mu\text{s)}}$$

As a result, the EDR for this TXOP is just 16.2 Mbps! Most wireless engineers think that the “MCS data rate” is the speed they should get on the medium when transmitting any kind of payload. Looking at packet captures tends to reinforce this idea because data MPDUs are always displayed with their actual TX data rate. But the truth is that the MAC overhead dramatically diminishes the EDR for most common on-air traffic except for file transfers and streaming applications.

This calculation also assumes there are no retries. In the case of a retry, some or all of the entire previous TXOP becomes additional overhead. The EDR for retried frames is further reduced as a result.

Performing What-If Analysis

Now that we have constructed the calculator, we can perform various kinds of “what-if” analysis on the scenario. Let us change the control rate from 6 Mbps to 24 Mbps. These changes are highlighted with red boxes in Table T3-3. We see that while the effective TXOP data rate has jumped by only 5 Mbps, we have reduced the airtime by 23% by recovering 66 μs. As already explained, that reclaimed airtime will add up to huge gains because it is saved for every TXOP.

Table T3-3 TXOP Airtime Calculator (512-B payload, 24-Mbps Control Rate)

MAC Unit	Payload Bytes	Payload Bits	Data Rate	μsec	% Airtime
Legacy Preamble			6 Mbps	20.00	9.0%
RTS	20	160	24 Mbps	6.67	3.0%
SIFS				16.00	7.2%
Legacy Preamble			6 Mbps	20.00	9.0%
CTS	14	112	24 Mbps	4.67	2.1%
SIFS				16.00	7.2%
Legacy + VHT Preambles			6 Mbps	44.00	19.9%
A-MPDU Delimiter	4	32	86.7 Mbps	0.37	0.2%
Data Frame Payload	512	4096	86.7 Mbps	47.24	21.3%
SIFS				16.00	7.2%
Legacy Preamble			6 Mbps	20.00	9.0%
BA	32	256	24 Mbps	10.67	4.8%
Airtime for TXOP only (excluding arbitration)				221.61	100.0%
Effective TXOP rate for TXOP only (excluding arbitration)				21.0	

You can also vary the amount of data to send to study the relative efficiency of small and large MPDUs. If you want to see a TCP ack, just plug in 90 bytes instead of 512. If you want to see an A-MSDU of 2 for full-buffer traffic, plug in 3,000 bytes instead.

For the more visually-inclined reader, it can be even more compelling to turn the calculator into a stacked bar chart. This gives us a way to literally see the time taken by the entire TXOP. In [Figure T3-10](#), we have used the calculator to compare the airtime required for a 90 byte vs. 3,000 byte MPDU size. This method makes it quite easy to visualize the relative efficiency of the two transmissions, as well as to perceive just how much of the time the channel is quiet instead of in a transmitting state.

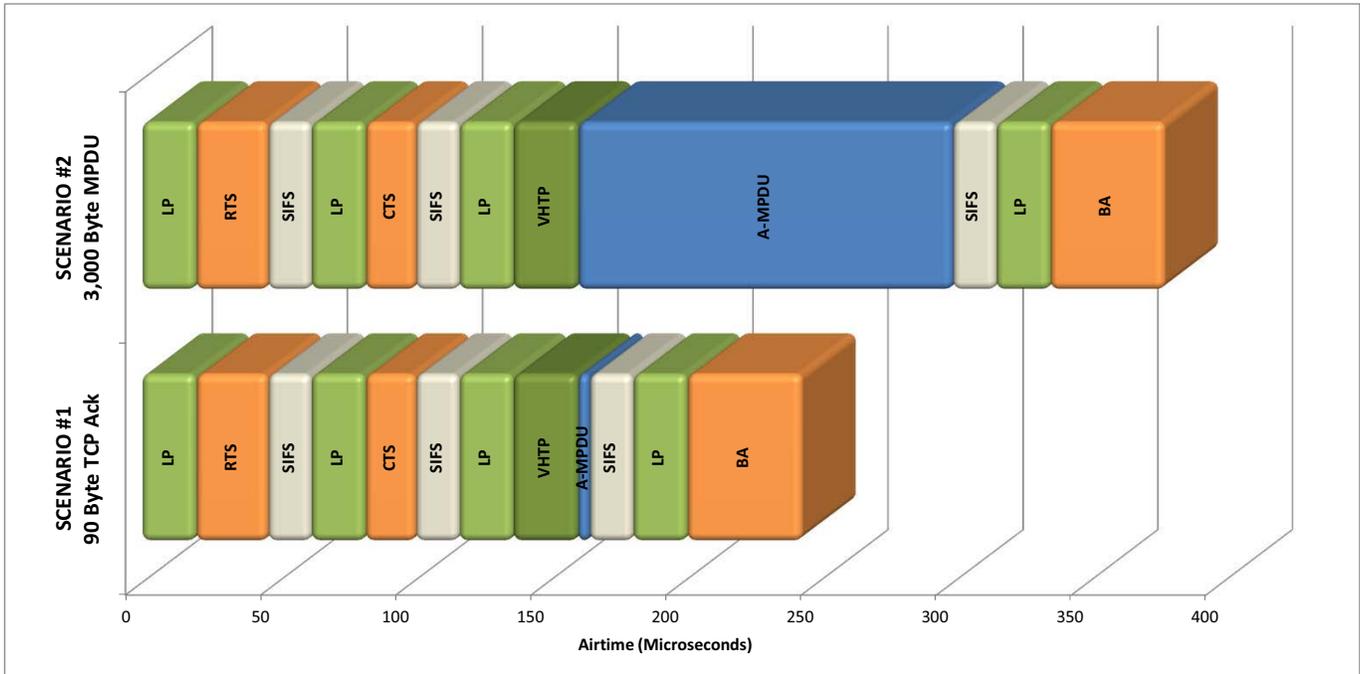


Figure T3-10 Visual Comparison of Airtime for 90B and 3,000B MPDUs (Excluding Arbitration)

You can do other things from a what-if perspective. You can add additional rows for more MPDUs in an A-MPDU burst. (Remember that every MPDU has a 4-byte delimiter). Finally, you can change the data rate of the MPDU itself. If you want to see how long a 2SS VHT80 station requires, plug in an 866.7 Mbps rate. The possibilities are endless!



Aruba is providing an airtime calculator as part of this VRD. It is available for download from the VRD page of the Aruba Networks web site.

Effects of Arbitration

Until now, our analysis has excluded the airtime required for the arbitration process. This process adds additional time to each TXOP and changes the result significantly enough that it is worth studying on its own. Let us update [Figure T3-6 on page 22](#) to show:

- The fixed AIFS period at the beginning of the TXOP for the [BE] queue
- The variable length contention window
- Use of enhanced 24-Mbps control frame payloads

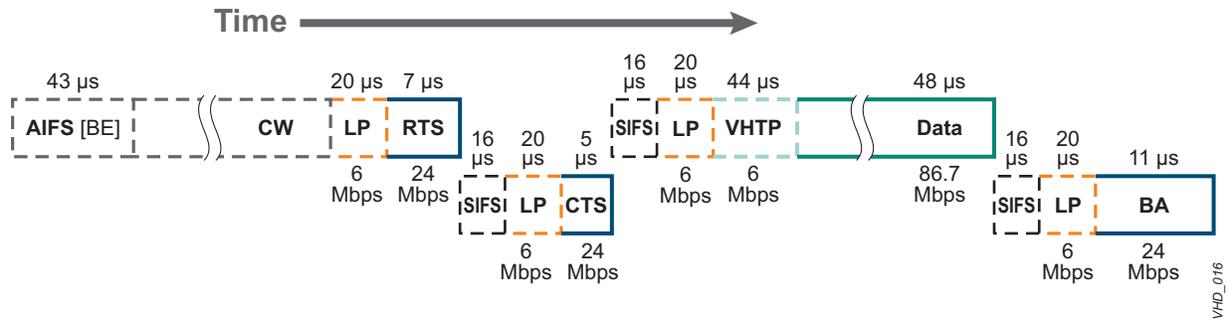


Figure T3-11 Full TXOP Structure Including Arbitration Period

We also must add two new rows to the top of our frame time calculator: one for the AIFS value and one for the contention window. Both should be adjusted for the CoS that is used. Table T3-4 adds these rows to the calculator. The calculator has been further adjusted with a data payload of 90 bytes to simulate a TCP acknowledgement.

Table T3-4 TXOP Airtime Calculator with Arbitration

MAC Unit	Payload Bytes	Payload Bits	Data Rate	μsec	% Airtime with CSMA	% Airtime TXOP Only
AIFS[BE]				43.0	14.5%	
Contention Window [BE]				72.0	23.9%	
Legacy Preamble			6 Mbps	20.0	6.7%	10.9%
RTS	20	160	24 Mbps	6.7	2.2%	3.6%
SIFS				16.0	5.4%	8.8%
Legacy Preamble			6 Mbps	20.0	6.7%	10.9%
CTS	14	112	24 Mbps	4.7	1.6%	2.6%
SIFS				16.0	5.4%	8.8%
Legacy Preamble			6 Mbps	20.0	6.7%	10.9%
VHT Preamble			6 Mbps	24.0	8.1%	13.1%
A-MPDU	94	752	86.7 Mbps	8.7	2.9%	4.7%
SIFS				16.0	5.4%	8.8%
Legacy Preamble			6 Mbps	20.0	6.7%	10.9%
Block Ack	32	256	24 Mbps	10.7	3.6%	5.8%
Total Airtime including CSMA		1,280		297.7	100.0%	100.0%
Effective TXOP rate including CSMA				4.3		
Total Airtime for TXOP only				182.7		
Effective TXOP data rate for TXOP only				7.0		

In this example, the AIFS for [BE] is fixed at 43 μs . This quiet period is nearly as long as a CTS (40.7 μs) or a BA (46.7 μs) *including the SIFS and legacy preamble*. For the CW value, the random timer can begin anywhere between 0 μs and 279 μs for a first transmission attempt. We choose 72 μs (or eight slot times) as an arbitrary fixed value for the calculator. Together, the AIFS plus the CW total 115 μs , **which on a percentage basis is about 40% of the total TXOP!** Collectively, the arbitration period plus the three SIFS equal 55% of the TXOP during which nothing is actually happening on the channel.

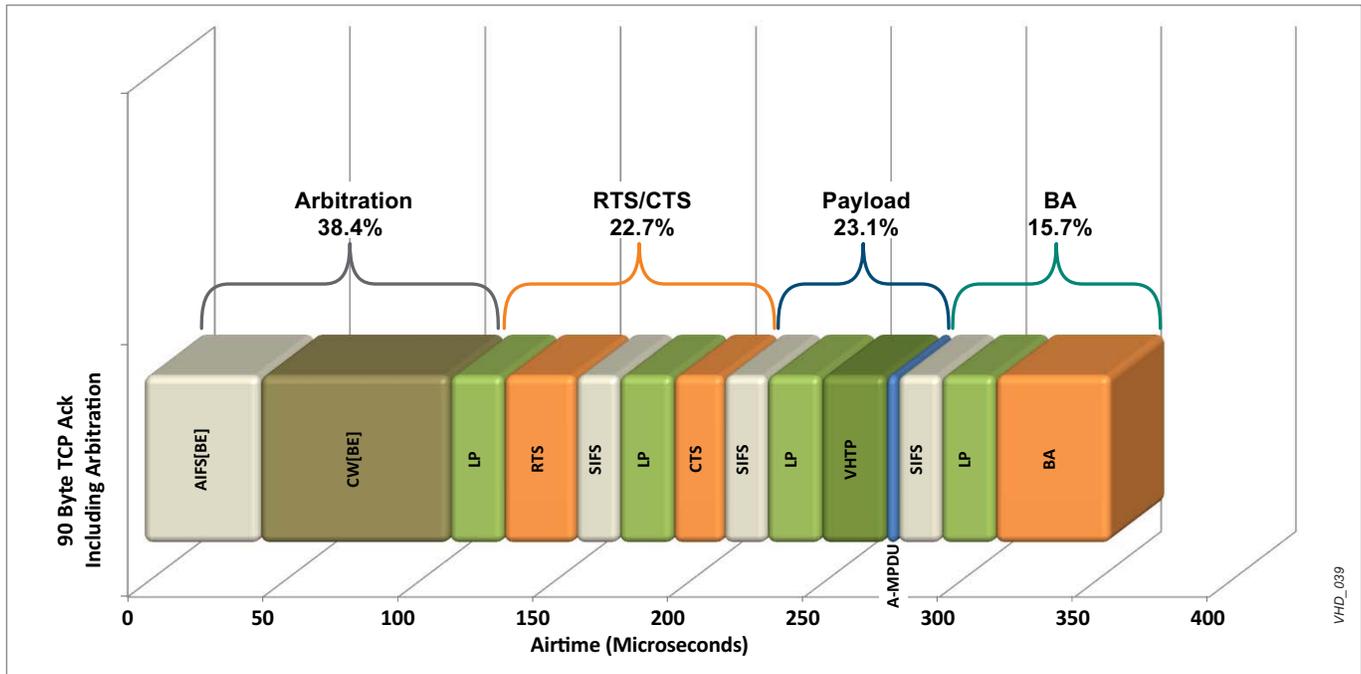


Figure T3-12 Visual Analysis of TXOP Airtime Including Arbitration

The metaphor that comes to mind when looking at this chart is sending a rocket into space. One rule of thumb in rocket design is that a maximum of between 1% and 4% of the total launch mass can be payload, depending on the ultimate destination of the vehicle.^{1 2} The remaining launch mass is made up of fuel and the vehicle itself, without which the payload cannot be delivered. Just like the capsule at the top of a rocket, the 90 byte MPDU shown in Figure T3-12 is essentially cargo. From Table T3-4, we see that the A-MPDU requires about 3% of the airtime to send at MCS8. The other 97% of the total airtime required by the TXOP is analogous to the rocket vehicle and its fuel. And like the rocket, the TXOP airtime is thrown away during the transmission and cannot be reused.

Admittedly, this is scenario a conservative example. The CW is actually a random value that could be less than eight slot times (and most likely will be with multiple STAs contending). But the CW time could also be significantly more. We are using this value to make the point about the airtime impact of arbitration, as well as the very poor airtime efficiency of common frame sizes.

When the entire TXOP duration is factored in for this example, the EDR drops from 7 Mbps to just 4.3 Mbps, even though the data payload for the TCP ack is sent at the full MCS8 rate of 86.7 Mbps.

1. http://en.wikipedia.org/wiki/Payload_fraction
 2. http://en.wikipedia.org/wiki/Saturn_V

As you can see, every transmission has an airtime cost. The most “expensive” parts are the repeating control frames. Conventional, non-VHD areas do not need to worry about this overhead because their overall duty cycles are generally low. However, the opposite is true in VHD areas. So you should scrutinize each and every transmission to determine whether it is necessary, whether it could be sent at a faster rate, and how to minimize the amount of retries that may occur.

Average Frame Size Measurements In Live Environments

This guide has asserted several times that average frame sizes are quite small in VHD areas, on the order of 500 bytes or less. In fact, average frame sizes are small in the vast majority of WLANs most of the time. Only those networks whose traffic is comprised primarily of file transfers, video streaming, or speed tests will show larger values. Normally, these traffic types are not a significant amount of the load in most WLANs.

Aruba Administration Building

To quantify the real range of frame sizes on live networks, we took multichannel packet captures over 30 minutes in a busy part of the Aruba main administration building. Traffic was captured on channels 36+, 44+, 132+, and 157+ on the 5-GHz band, and on channels 1 and 11 in the 2.4-GHz band. These channels were chosen to get some coverage on each major band.

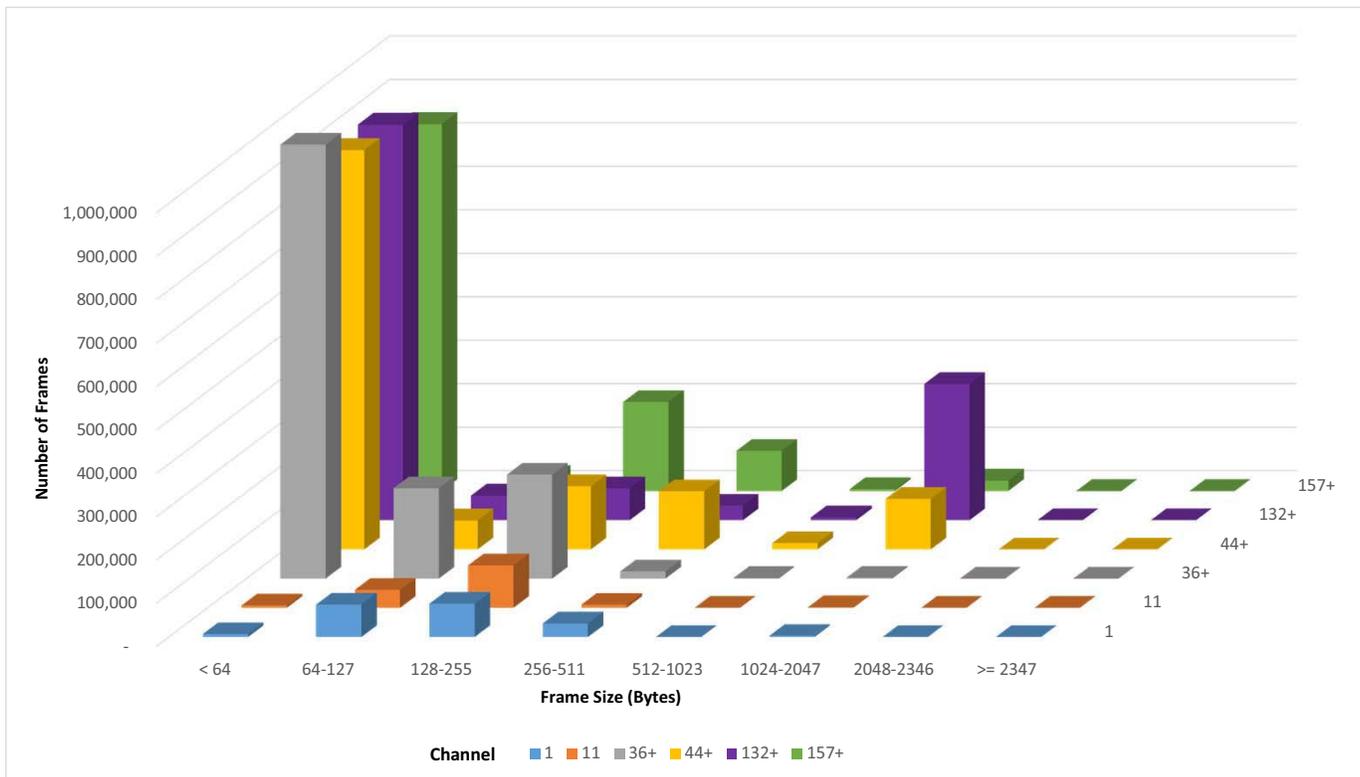


Figure T3-13 Frame Size Distributions in Office Environment (6 channels, 30 minutes)

Figure T3-13 shows histograms for each of the six channels. The overall average frame size was 201 bytes in 5-GHz and 191 bytes in 2.4-GHz. Over 80% of frames in both bands were under 256 bytes. Figure T3-14 shows the aggregate results by band in pie chart format.

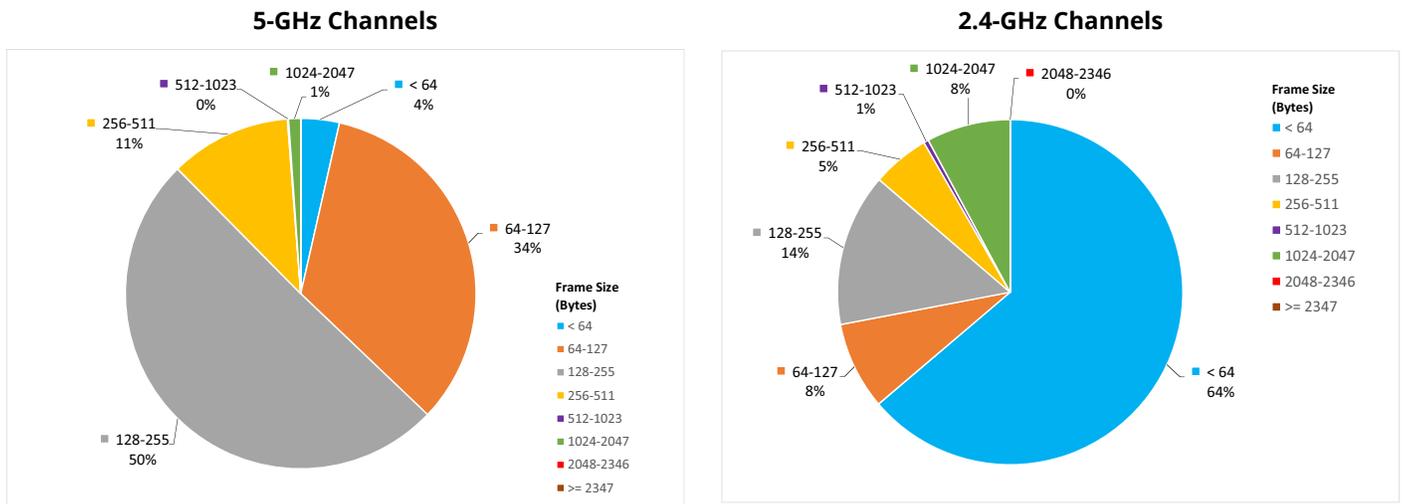


Figure T3-14 Frame Sizes by Band in Office

Football Stadium

Aruba has also taken measurements in a range of VHD environments. The data in Figure T3-15 was taken over a 10-minute period in the third quarter of a football game in a 70,000 seat stadium. Traffic was captured on eight channels.

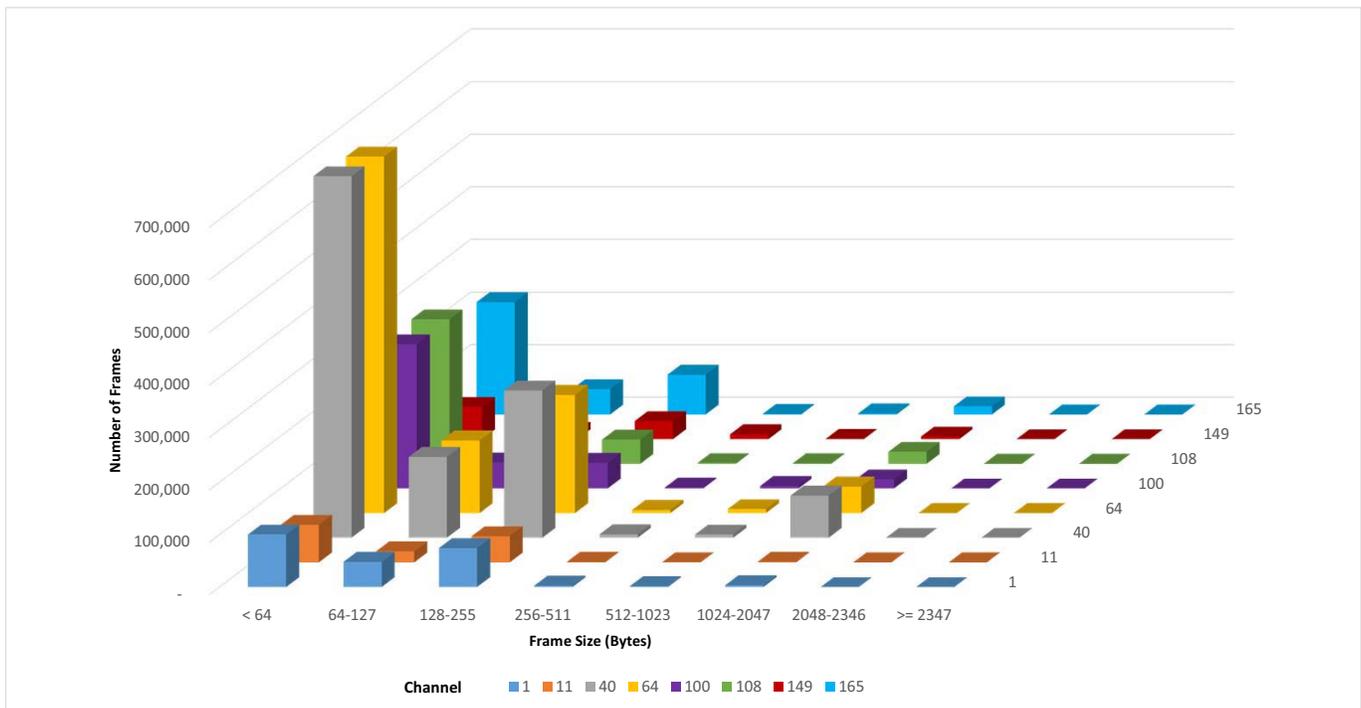


Figure T3-15 Frame Size Distributions in During Football Game (8 channels, 10 minutes)

As with the office example, it's clear that the vast majority of traffic is under 256 bytes. The overall average frame size was a mere 160 bytes in 5-GHz and just 125 bytes in 2.4-GHz. Over 80% of frames in both bands were under 256 bytes. [Figure T3-16](#) shows the aggregate results by band in pie chart format.

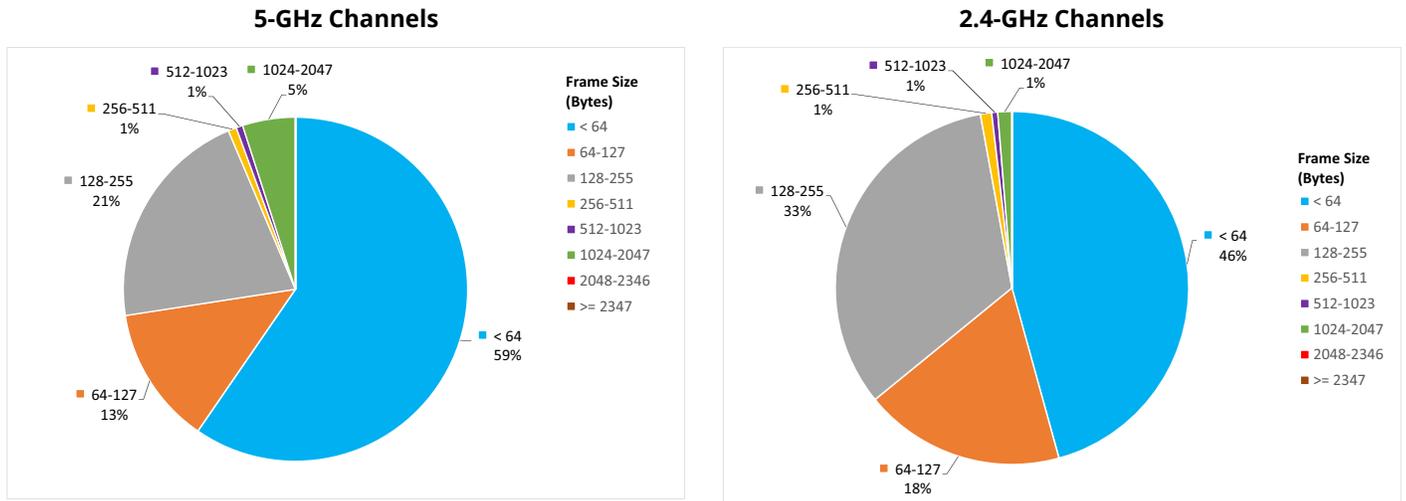


Figure T3-16 Frame Sizes by Band During Football Game

The huge amount of sub-64 byte frames strongly suggests that these are 802.11 control frames. We can verify this with the frame type breakdown in the packet capture tool see whether this is true. In fact, we find that 58% of total frames during the measurement were control frames. You may be surprised to learn that data frames were less than 25% of total traffic during the period (see [Figure T3-17](#)).

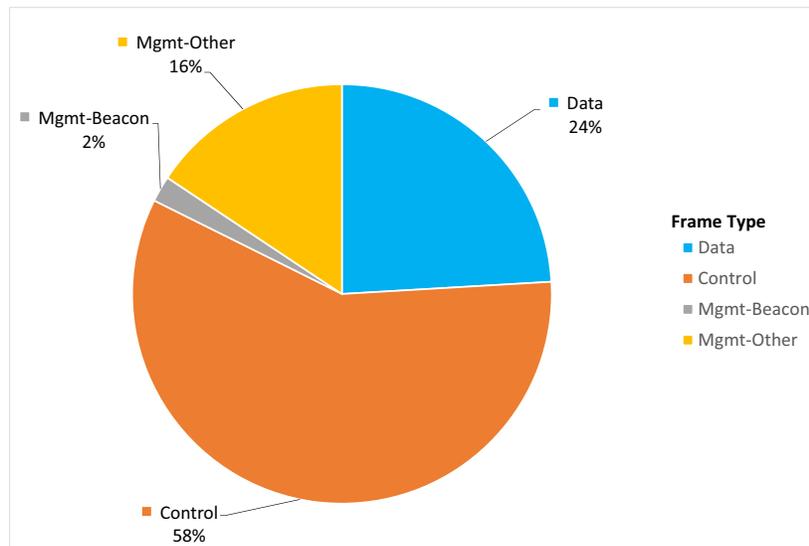


Figure T3-17 Frame Type Distribution in Football Game

These results, combined with the earlier analysis of TXOP structure, paint a sobering picture of the enormous potential user capacity that is lost to management and control traffic overhead in an 802.11 system. The data is saying that the air is very busy, but it does not necessarily carry useful data payload. This is probably just fine when there are not many users present, or the duty cycles are low. However, when a large spike in traffic happens due to some event, the system needs all the latent capacity it can get to absorb the spike. This example shows why you must learn to become almost fanatical about airtime recovery and enforcing good airtime management practices.

What is the Relationship Between Airtime and Bandwidth?

For this chapter, we use the term bandwidth to mean the amount of data transferred in a given amount of time. Bandwidth usually is expressed in bits per second. For example, a Wi-Fi speed test might generate 60 Mbps upstream and 80 Mbps downstream.

One corollary of the TXOP EDR analysis is that **data bandwidth can never exceed the EDR of the TXOPs used to send the data**. In fact, upper-layer protocol overhead as well as Layer 2 retransmissions further reduce the usable bandwidth below the EDR.

Furthermore, speed tests tend to overestimate the usable bandwidth of the channel. Such tests involve sending continuous full-buffer traffic, which allows the network driver to employ frame aggregation to drive down overhead a certain percentage of the TXOP and raise the EDR for that specific traffic.

But speed tests are just a specialized and infrequent type of load in a VHD network. Most normal traffic in VHD wireless networks consists of small, transactional, upper-layer packets. Their data payloads are small and often cannot be aggregated. So for that type of traffic, the actual usable bandwidth of the channel is more like the examples just presented.

Why is Wired Bandwidth Fixed but Wireless Bandwidth Varies?

Wired networks have a fixed relationship between bandwidth and time. Wired interfaces send at well-known, fixed PHY rates: 10 Gbps, 1 Gbps, DS-3, T-1, and so on. Furthermore, most wired network topologies are:

- Effectively point-to-point (for example, switched Ethernet and fiber links)
- Full duplex
- Collision-free due to lack of contention and direct medium sensing
- Free from external interference
- Served by aggregating equipment at all layers (access, distribution, or core) with considerably higher backplane bandwidth than any individual interface.

As a result, the data bandwidth of any given speed test is basically equal to the link speed. An iPerf test between two laptops with Gigabit Ethernet interfaces should produce just under 1 Gbps of bandwidth, regardless of whether the test is run for 1 second or 60 seconds. The limiting factor of course is the CPU utilization of each machine.

By contrast, Wi-Fi differs from wired networks in these important ways:

- A radio channel is hub, not a switch. It is shared between all users who can hear (decode) one another's transmissions.
- Only one user can send at one time in the same RF collision domain.
- Collisions cannot be directly sensed, so a listen-before-talk method must be used, which consumes time (reduces capacity).

- Protocol overhead to take control of the channel reduces the usable capacity. This overhead can vary with load and external interference.
- The data rate for any single data frame payload can vary by over 2 orders of magnitude based on a dizzying array of criteria (for example, from 6 Mbps to 1.3 Gbps).
- The maximum data rate of a given client varies widely based on the capabilities of its hardware (principally its Wi-Fi generation and number of spatial streams).
- All transmissions must be acknowledged or they are assumed to have failed. Acks are sent at a very low data rate, which reduces overall channel efficiency.

The result is that it is utterly impossible to know from a simple speed test result what the actual conditions of the test might have been. For example, here are just four of many possible scenarios that could produce a speed test of 100 Mbps over Wi-Fi:

- **Good:** 1 spatial stream 802.11n smartphones in an HT40 channel (max PHY rate of 150Mbps)
- **Average:** 2 spatial stream 802.11ac tablets in a VHT20 channel (max PHY rate of 173.3Mbps)
- **Poor:** 2 spatial stream 802.11n laptops in an HT40 channel (max PHY rate of 300Mbps) with some co-channel interference from nearby APs
- **Awful:** 3 spatial stream 802.11ac laptops in a VHT80 channel (max PHY rate of 1.3Gbps) with significant interference

Use these examples when you work with others to explain some of the unique dynamics of Wi-Fi performance.

Summary

This chapter has two main goals:

- To open your eyes to the enormous amount of unproductive overhead that goes into radio communication
- To make you worried enough about it to become completely paranoid and relentless about how airtime is used in the VHD environments for which you are responsible

Your single most important strategy to increase VHD capacity is to maximize the efficiency of the airtime you have on each and every channel.

In addition, the preceding discussion is intended to create or enhance your awareness of these aspects of Wi-Fi operation in VHD areas:

- The basic structure of an on-air transaction in 802.11ac
- How time is consumed during an on-air transaction
- The massive amount of overhead needed to send a basic data packet
- The need to avoid **all** unnecessary TXOPs and associated overhead
- The need to maximize data rates of data frame payloads
- The impact of raising 802.11 control rates to reduce busy time
- The relationship between airtime and actual throughput

These examples intentionally leave out aggregation. MPDU aggregation can help swing the efficiency back the other way. Unfortunately, the vast majority of frames sent in VHD environments are small single-packet TCP exchanges that can only be aggregated some of the time. In practice, aggregation is only impactful for video sessions and speed tests.

Bibliography

- Westcott, David A., David D. Coleman, Ben Miller, and Peter Mackenzie. *CWAP Official Study Guide*. Sybex, 2011. Available online at <http://www.wiley.com/WileyCDA/WileyTitle/productCd-0470769033,miniSiteCd-SYBEX.html>.
- Perahia, Eldad and Robert Stacey. *Next Generation Wireless LANs: 802.11n and 802.11ac*. Cambridge University Press, 2013.
- Gast, Matthew. *802.11ac: A Survival Guide*. O'Reilly Media, 2013. Available online at <http://shop.oreilly.com/product/0636920027768.do>.

Chapter T-4: How Wi-Fi Channels Work Under High Load

Now that you have a solid understanding of airtime for individual transmit opportunities (TXOPs), consider how an entire channel performs when hundreds of devices attempt to use it at the same time.

The key question this chapter addresses is this: *Why does the total capacity of an 802.11 channel decrease as the number of stations trying to use it increases?*

In the 2010 edition of this VRD, Aruba published research showing that the total aggregate throughput of an 802.11 channel declines as up to 50 devices are added to a test. Since then, other major WLAN vendors have reported similar findings.

For this edition, Aruba set out to increase the testbed size to 100 simultaneous devices. We decided to explore the effect of different numbers of spatial streams. We also set a goal to explain the underlying mechanism of this effect.

To achieve these goals, we built a VHD test lab with 300 802.11ac devices. The test lab has three pools of 100 devices each. One pool is 1SS smartphones, another is 2SS laptops, and the third is 3SS laptops. With this equipment, we can study a wide range of device combinations to answer important questions about VHD performance. For detailed information on the testbed, see [Appendix T-A: Aruba Very High-Density Testbed](#).

One of the three important variables in the TST capacity planning formula presented in [Chapter EC-2: Estimating System Throughput](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide is *average channel bandwidth*. As a WLAN architect, one of your responsibilities in VHD design is to choose an appropriate value for this term. But if the bandwidth value you choose itself depends on load, how do you decide? Our goal in doing the research and writing this chapter is to give you a clear understanding of what is happening in the channel so that you can successfully apply the methodology to your own deployments.

Channel Capacity Is Inversely Proportional to Client Count

One of the most important phenomena governing the performance of a Wi-Fi® channel in a high-density environment is that capacity decreases with load.

For example, [Figure T4-1](#) shows results for 100 station tests of a 1SS phone, a 2SS laptop, and a 3SS laptop in a VHT20 channel. On this chart, the horizontal axis is the number of STAs in the test. Notice that total throughput decreases from left to right, as we increase the number of STAs from 1 to 100. You may recall seeing something similar on the test results presented in [Chapter P-3: RF Design](#) of the *Very High-Density 802.11ac Networks Planning Guide*, and [Chapter EC-2: Estimating System Throughput](#) and [Chapter EC-3: Airtime Management](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide.



As a reminder, the peak PHY rate in a VHT20 channel is 86.7 Mbps for a 1SS device, 173.3 Mbps for 2SS, and 288.9 Mbps for 3SS.

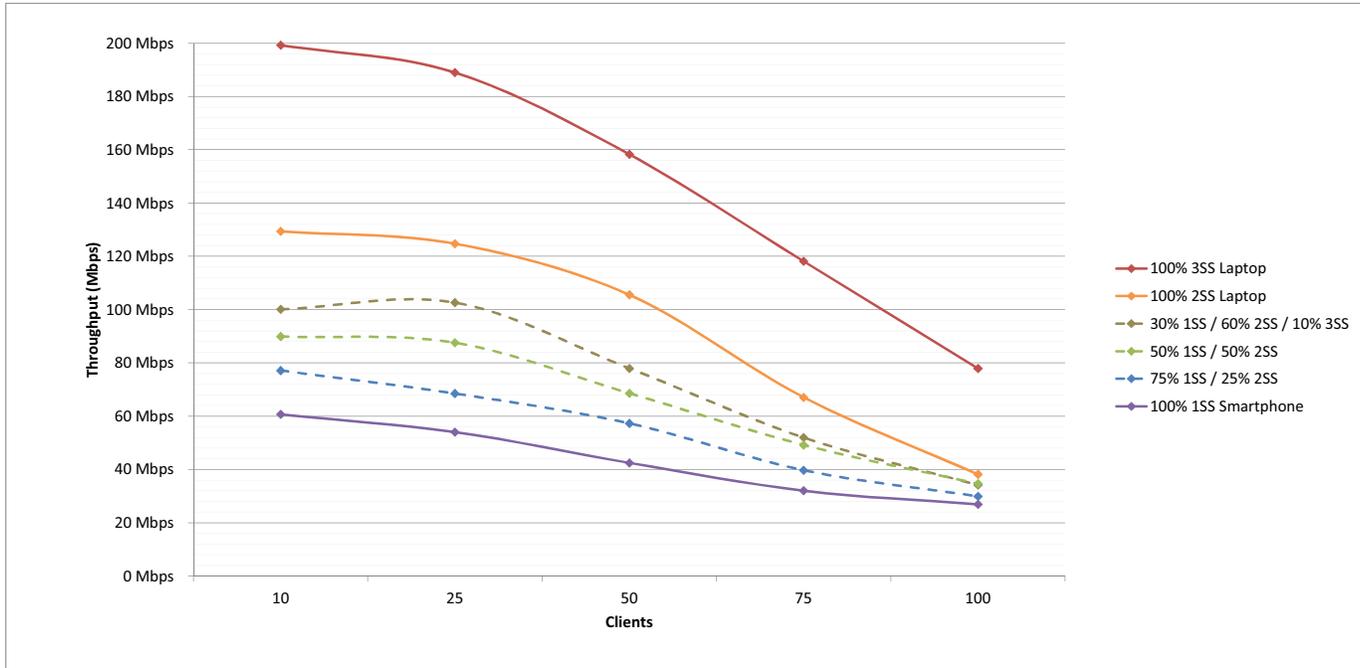


Figure T4-1 Capacity Decreases as Client Count Increases (AP-225, VHT20, TCP Bidirectional)

This behavior is a fundamental property of 802.11. You will see this type of curve from every WLAN vendor, with every client vendor, and in every channel width. This behavior is true of 802.11a, 802.11n, and 802.11ac. As you can see from [Figure T4-1](#), it is true with multiple input, multiple output (MIMO) regardless of the number of spatial streams. You can reproduce this result in your own lab if you gather 50 or more clients to test.

It's probably not even that surprising to you (although the magnitude of the roll-off beyond 50 STAs may raise eyebrows). Intuitively, we know that Wi-Fi is a shared medium that is prone to collisions. More users means that each station gets less airtime. So it is reasonable to expect that we would get less total goodput with 100 stations than with 5 or 10 stations.

Or is it? Does the collision hypothesis stand up under closer scrutiny?

- Why should simply "cutting the pie" into more slices shrink the entire size of the pie by nearly 60%?
- Why would there be significantly higher collisions in a clean test environment with a single BSS and a well-ordered channel?
- Why is the drop so similar for a 3SS laptop that can move over 3X the data in the same airtime as a 1SS smartphone?

If collisions are not the primary cause of this effect, then what is? Can anything be done to control or limit this effect to recover some of the lost capacity? These questions will all be answered in this chapter.

Defining the Contention Premium

To begin our analysis, we need a way to normalize results from very different kinds of tests. Different channel bandwidths and different spatial stream counts produce very different absolute throughput numbers. So do different generations of equipment.

Aruba has found that a good comparison method is to replot client scaling test results on a percentage basis, using the single-station throughput value for that test as 100%. Each of the other data points in the test run are then expressed as a percentage of the single-station throughput. If we do this with the results for the three different client types shown in [Figure T4-1](#), we can get an idea of the scale of the drop. [Figure T4-2](#) shows that data in the percentage format.

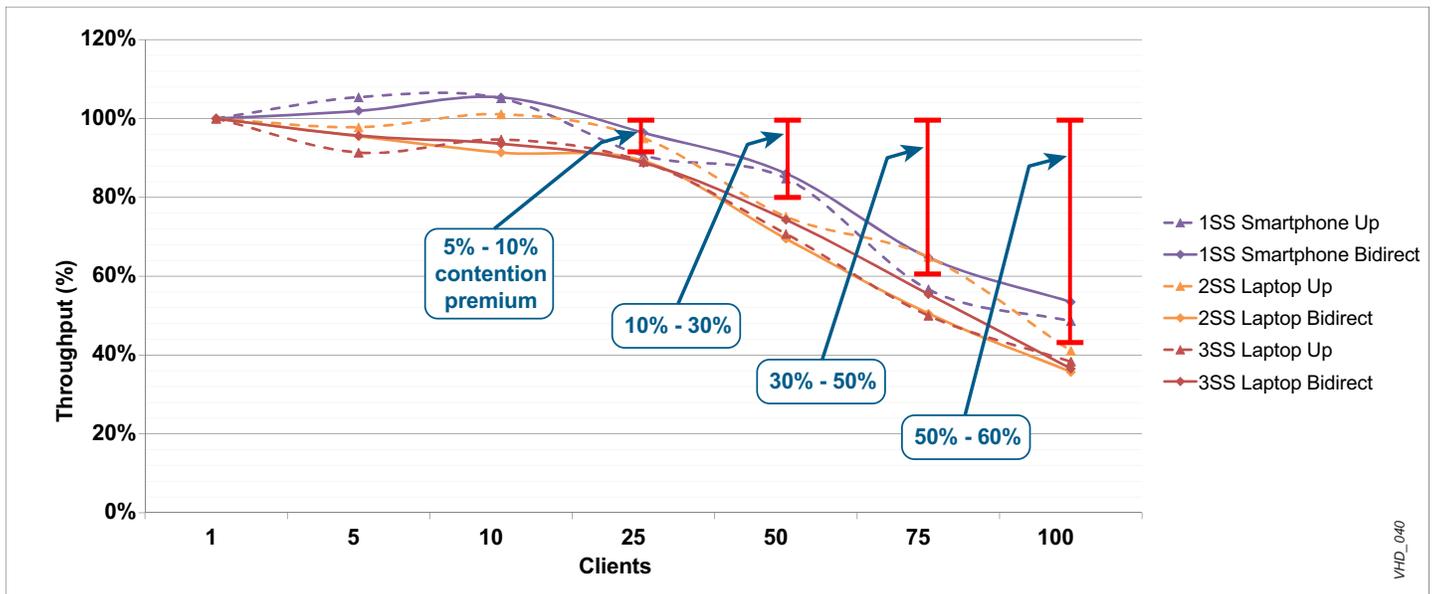


Figure T4-2 Contention Premium as a Percentage of One Station Throughput

Aruba defines the term “contention premium” to mean the difference between total aggregate throughput for one station as compared with a larger number of stations. For the test shown in [Figure T4-2](#), the contention premium increases from an average of about 10% at 25 stations to about 60% at 100 stations. Though we see some variation from client to client and run to run, the consistency of the overall trend is quite clear.

Explaining the Contention Premium

There are at least four possible explanations for the contention premium phenomenon. They are:

- Collisions and retries
- Increase in downward rate adaptation
- TCP windowing
- MAC layer framing and airtime consumption

All four candidates are at work all the time in any WLAN. We want to determine if any one of these is the primary cause.

To get to the bottom of what is happening, we must study packet captures of these tests. The Aruba VHD testbed was designed for both wireless and wired packet captures; see [Appendix T-A: Aruba Very High-Density Testbed](#) for more information on how this was done. After studying the captures, we found that what is going on in the channel is rather different than one might expect.

Collisions and Retries Are Not the Cause

First we dispose of the collision hypothesis. Packet captures prove that corruption and collisions are not a significant factor in the contention premium. They certainly occur, but not in sufficient volumes to produce the effect in the lab environment. (In real world networks, collisions and retries are a major factor and will aggravate the results documented in the lab.)

We processed captures from tests with similar numbers of clients, and analyzed the volume of retries. A retry is indicated by a bit in the 802.11 MAC header. [Figure T4-3](#) shows a breakdown of the retry status of packets from an upstream 2SS laptop test similar to [Figure T4-1](#).

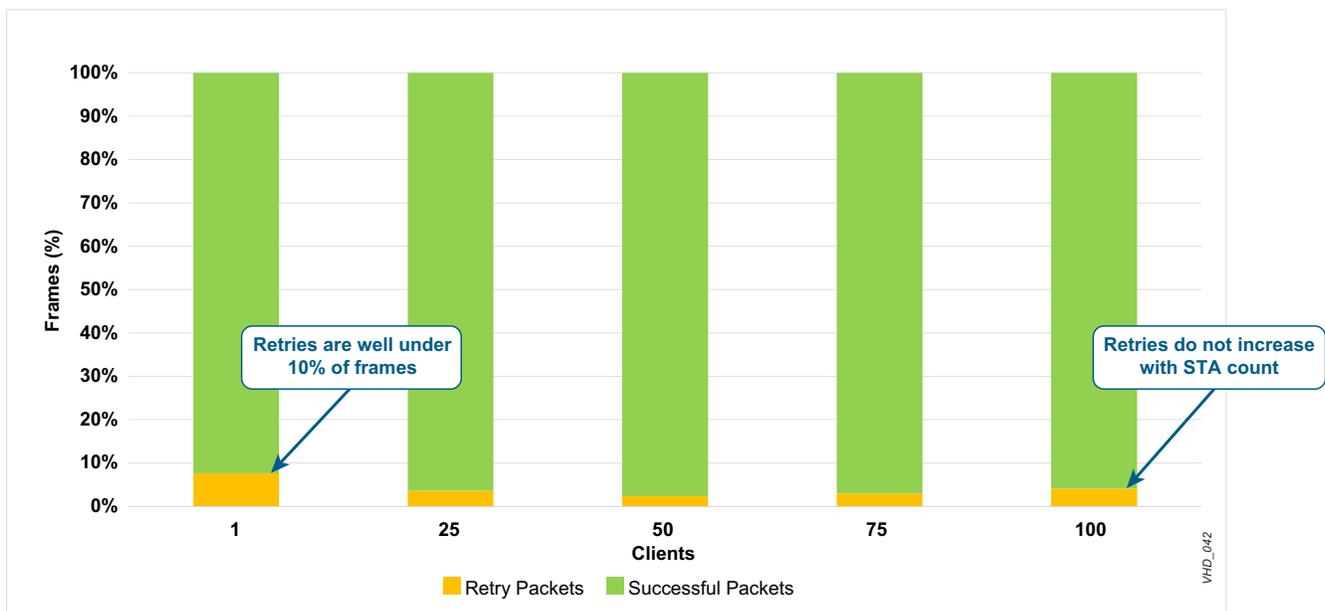


Figure T4-3 *Retry Comparison with 100 Station Scaling Test (AP-225, 20-MHz Channel, TCP Up)*

This test is done with 100 2SS MacBook Airs (MBAs) sending full-buffer TCP traffic upstream to an Aruba AP-225 3SS 802.11 access point in a 20-MHz channel. We see retries under 10% at all station counts. Of particular significance is that the retry rate does not increase as STAs are added to the test. [Figure T4-3](#) is broadly typical of our findings in all traffic directions and channel widths. We also see the same result for 1SS and 3SS stations.

Downward Rate Adaptation Is Not the Cause

When an acknowledgment is not received after a frame is sent, most Wi-Fi rate adaptation algorithms reduce the modulation and coding scheme (MCS) used for subsequent retries. If the packet fails at MCS7, the radio tries again at MCS6, then MCS5 and so on in the belief that perhaps the client needs a lower SNR or more robust modulation to recover the packet. Most radio drivers rate-adapt after just one or two retries.

After the rate drops, typically it stays down for communication with that client for some period of time. Periodically, a radio attempts a higher rate to probe whether conditions have improved. Meanwhile, all communication with that STA slows down as if a slow driver pulled in front of you on the highway.

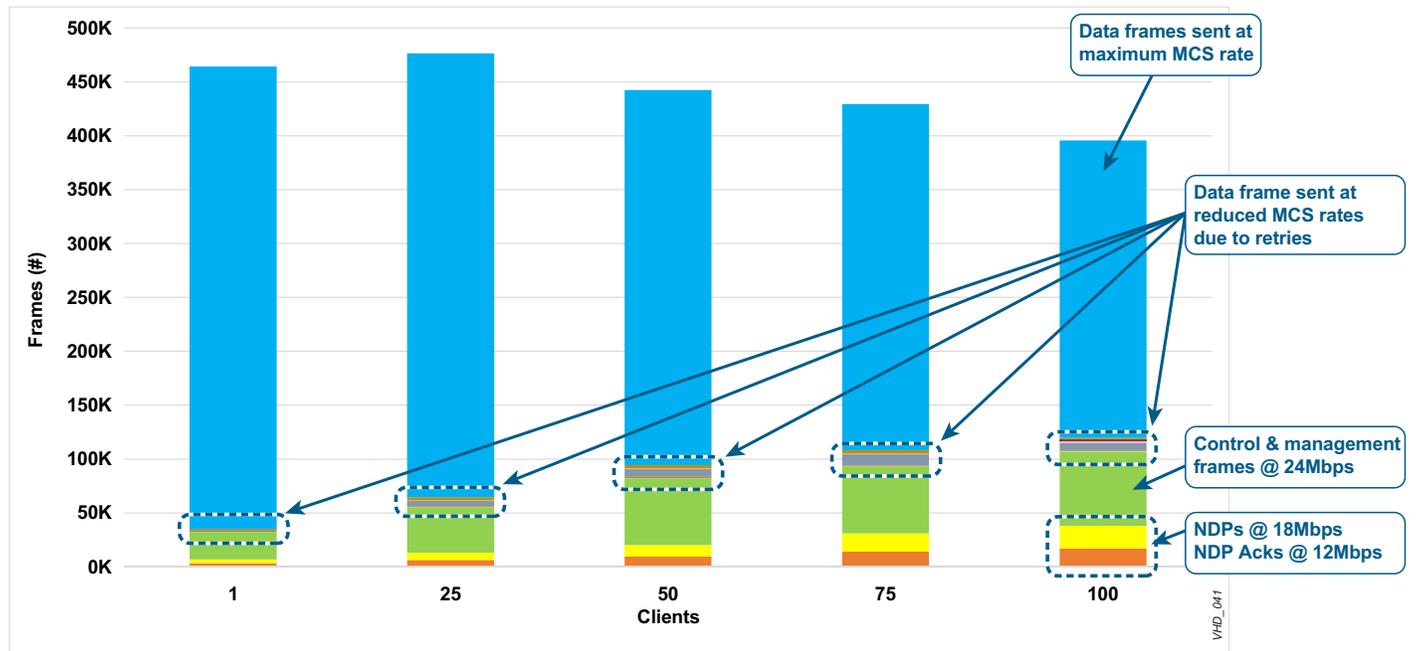


Figure T4-4 Data Rate Distribution in Scaling Test (AP-225, 20-MHz Channel, TCP Up)

We have already established that the overall level of retries is very low and does not grow as stations are added. Therefore, you would not expect to see unusual levels of retries. Figure T4-4 basically confirms this expectation.

The figure was produced by using the packet capture software to measure the distribution of data rates for all of the frames sent during the test. A limited amount of downward rate adaptation occurs in each test, which is considered normal as a percentage of the frames sent. However, the vast majority of data frames go at the maximum MCS rate that is usable by the 2SS clients in the test.

Recall that our Service Set Identifier (SSID) configuration uses a control rate of 24 Mbps. Figure T4-4 shows a significant amount of traffic at 18 Mbps and 12 Mbps rates. This traffic is not derating of the control frames. Rather, it is power-save state signaling carried in NDP frames, which are sent at 18 Mbps and the resulting Acks, which are sent at the next lowest rate.

Ruling Out TCP Windowing

Astute engineers with deep TCP protocol experience might immediately suspect some type of windowing limitation. Especially because maximum TCP window size is governed by the operating system and they remain unconscionably low even in 2015! As of the date of writing, the Windows default TCP window remains 65 KB where it has been for many years, while MacOS is just 128 KB.

However, Layer 4 protocols are not related to the contention premium effect. You easily can check this conclusion by running UDP tests and seeing whether the throughput degradation has approximately the same magnitude.

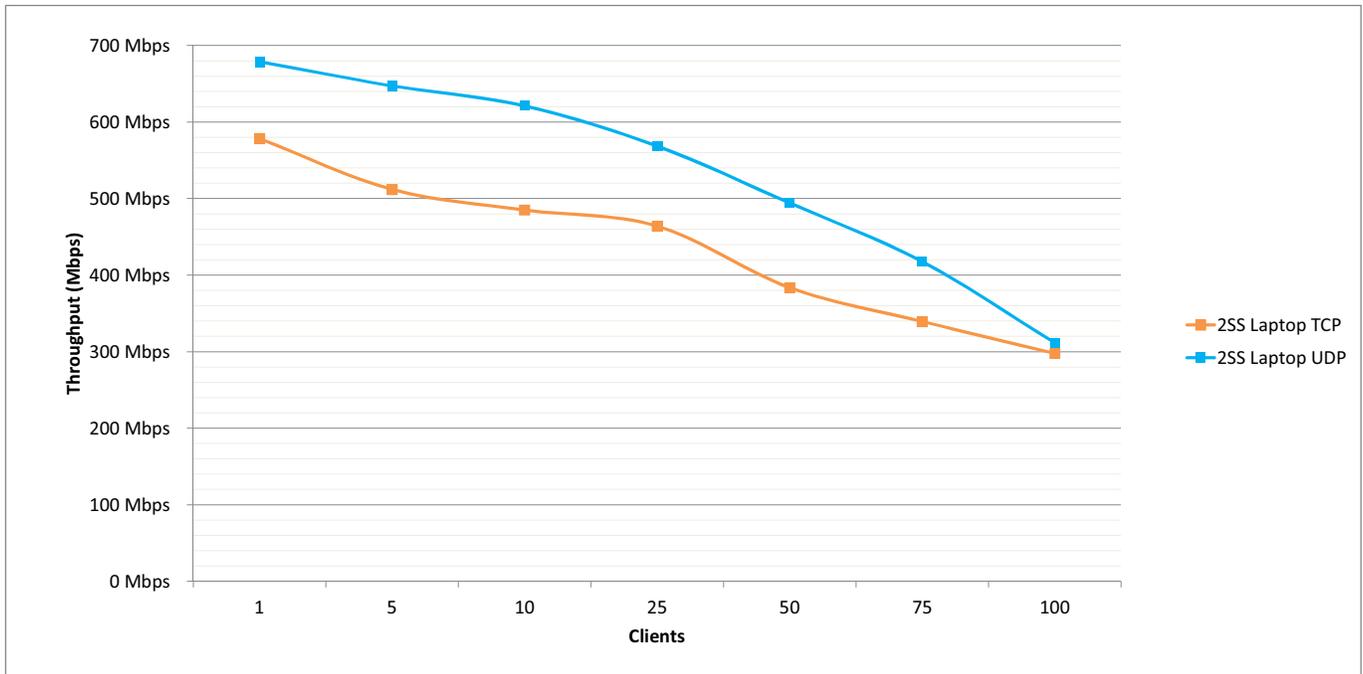


Figure T4-5 UDP vs. TCP Throughput Test in VHT80 Channel

Figure T4-5 shows an 80-MHz test with the same 100 MBAs. The TCP run is shown in orange, and the UDP run is shown in blue. As you can see, UDP is 10-20% faster than TCP across most of the range, but it still suffers from the contention premium. At 100 STAs, the TCP and UDP lines converge. This implies that the contention premium effect may accelerate with UDP at high STA counts.

Control Frame Growth Is the Critical Factor

Having ruled out three of the four potential explanations listed earlier, let us turn to an analysis of how the 802.11 MAC layer performs during these tests. Our tests indicate that the proximate cause of the contention premium is a significant increase in 802.11 control frames. Figure T4-6 shows the frame type distribution for the same series of AP-225 tests shown in Figure T4-1.

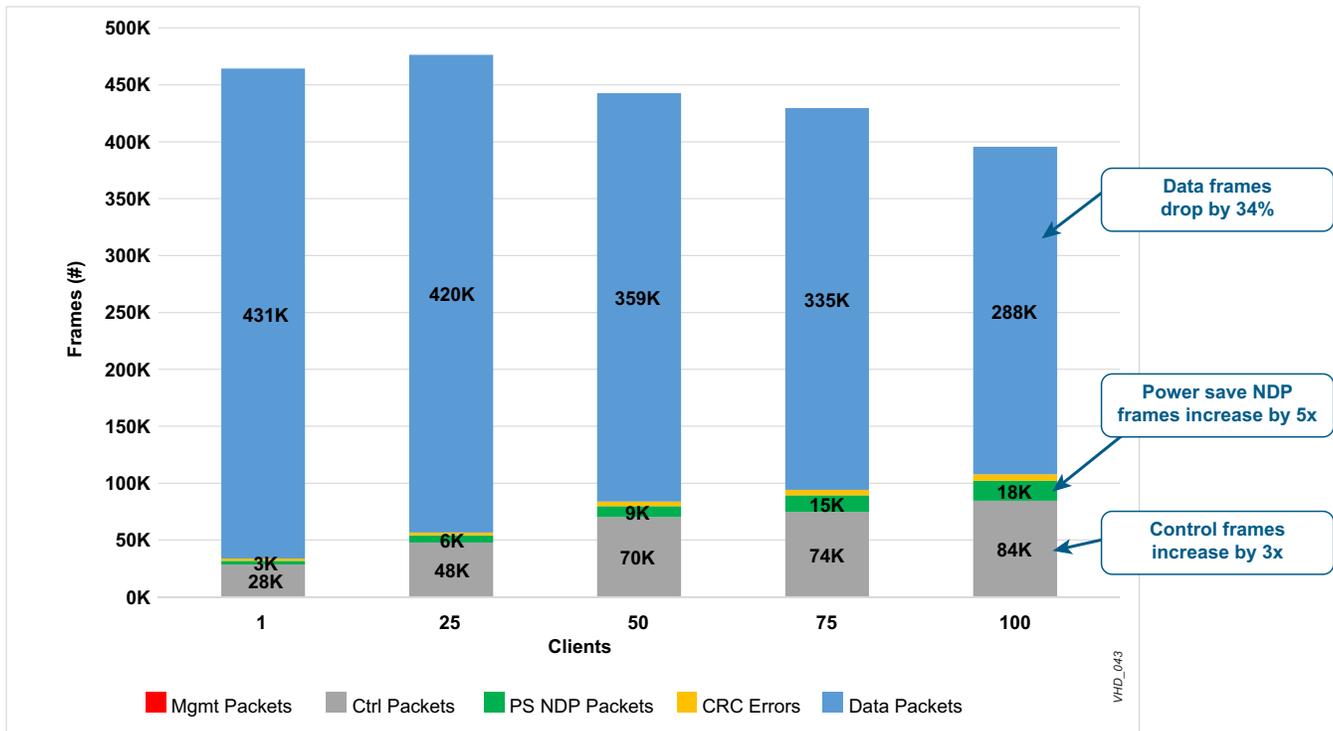


Figure T4-6 Increase in 802.11 Control Frames (AP-225, 20-MHz Channel, TCP Up)

With 100 MBAs and the 802.11ac AP-225, we observe that total volume of data frames decreases by 33% from 431,000 for one STA to 288,000 for 100 STAs. At the same time, 802.11 control packets increase by a factor of 3X from 28,000 to over 84,000.

We also observe that null data packet (NDP) power save (PS) signaling traffic increased by a factor of 5X, from 3,000 to over 18,000. NDPs are technically 802.11 data frames, however functionally they serve as 802.11 control frames to inform the AP that a STA is moving in and out of power save state. This signaling in turn regulates the flow of traffic to a PS STA.

Analysis

These frame type distributions directly explain the throughput loss with increasing STA counts in the client scaling tests. Reduced amounts of high-rate data frames, combined with increased amounts of low-rate control frames, can have only one result: a significant drop in relative throughput.

Aruba has measured the same phenomenon with both 802.11n and 802.11ac, and in all three channel widths. We have measured it with both AP and client radios from completely different manufacturers. This fact is strong evidence that the contention premium effect is a fundamental property of 802.11.

However, though the frame type distributions explain throughput loss, they do not tell us what the underlying mechanism is. Nor are they sufficient to prove causation. In other words, are the data frames decreasing because of the control frame increase, or the power save activity, or something else? For that we must continue to climb further down the rabbit hole.

Average Frame Size Decreases with Load

We can study packet size distributions and further break down the data frames by type, as a cross check on the conclusion that decreasing data frame volumes are the proximate cause of the throughput drop. The breakdown is expressed in two ways: by absolute numbers of frames on the left, and as a percentage of the total on the right. Both views contain important information.

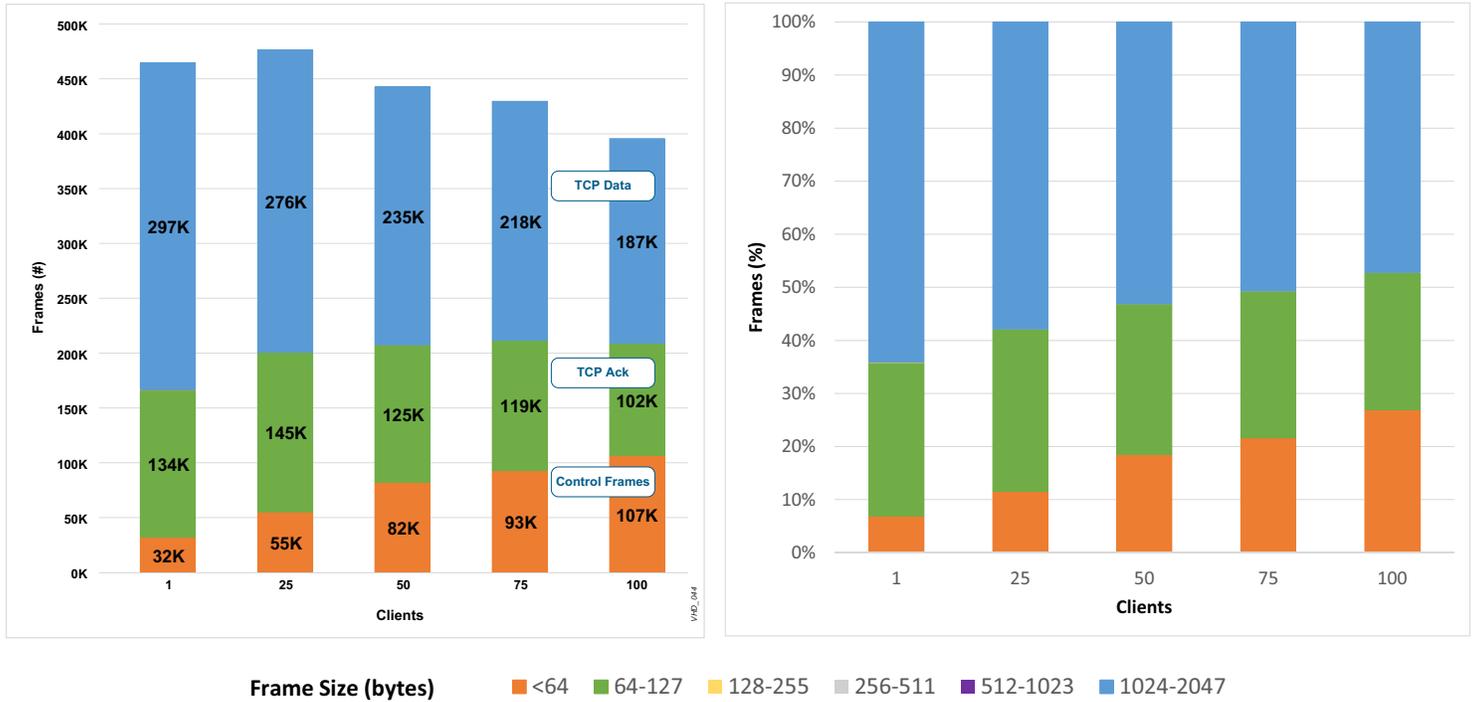


Figure T4-7 Decrease in Average Frame Size (AP-225, 20-MHz Channel, TCP Up)

In this presentation, we can now see that the data packets shown by the blue bar in Figure T4-6 can be further broken down into very large and very small frames. These frames are the 1,500-byte full-buffer TCP MPDUs sent by IxChariot followed by the 90-byte TCP acknowledgment for those payloads. This conclusion was confirmed by inspecting the actual packets in the trace.

On the left, we see that on an absolute basis, the data frame drop is even more substantial than seen in Figure T4-6. It is clear from the new chart that though both the payloads and acknowledgments are declining, payload frames are decreasing faster. Payloads drop by over 40% from 1 STA to 100 STAs, as compared with an overall 33% drop in all data frames. By contrast, acknowledgments drop by just 23%.

On the right, it is interesting to see the drop in payload frames from about 62% of the total with 1 STA to about 49% at 100 STAs. But an even more important conclusion can be drawn from the small frames. *Together, small frames account for over 51% of all frames sent during the 100 STA test.*

In Chapter T-3: Understanding Airtime, you learned the enormous airtime cost to send a small frame. Now that you have learned to “see” time, take another look at Figure T4-7 with this concept in mind. Think about the preambles and interframe spaces that are implied in the figure. The relative airtime efficiency of the small data frames and the control frames is extremely poor. Using the airtime calculator tool, you can compute that payload frames are less than 4% of the total airtime. While this chart looks bad from an absolute frame count perspective, if you were to replot it according to airtime, it would look much worse.

Analysis

In summary, the frame size breakdown reinforces the conclusion that the throughput drop is directly attributable to a total reduction of data payload frames sent on the network. The frame size breakdown does not prove causation. However, [Figure T4-7](#) strongly implies that this reduction is a by-product of the control frame growth.

Causes of Control Frame Growth

There are nine different 802.11 control frame types, so we must look deeper into the control frame component of [Figure T4-6](#) just as we did with the data frames. A breakdown of the control frame distribution from that test is shown in [Figure T4-8](#). As stated earlier, we are treating PS NDP frames as 802.11 control frames for purposes of this analysis.

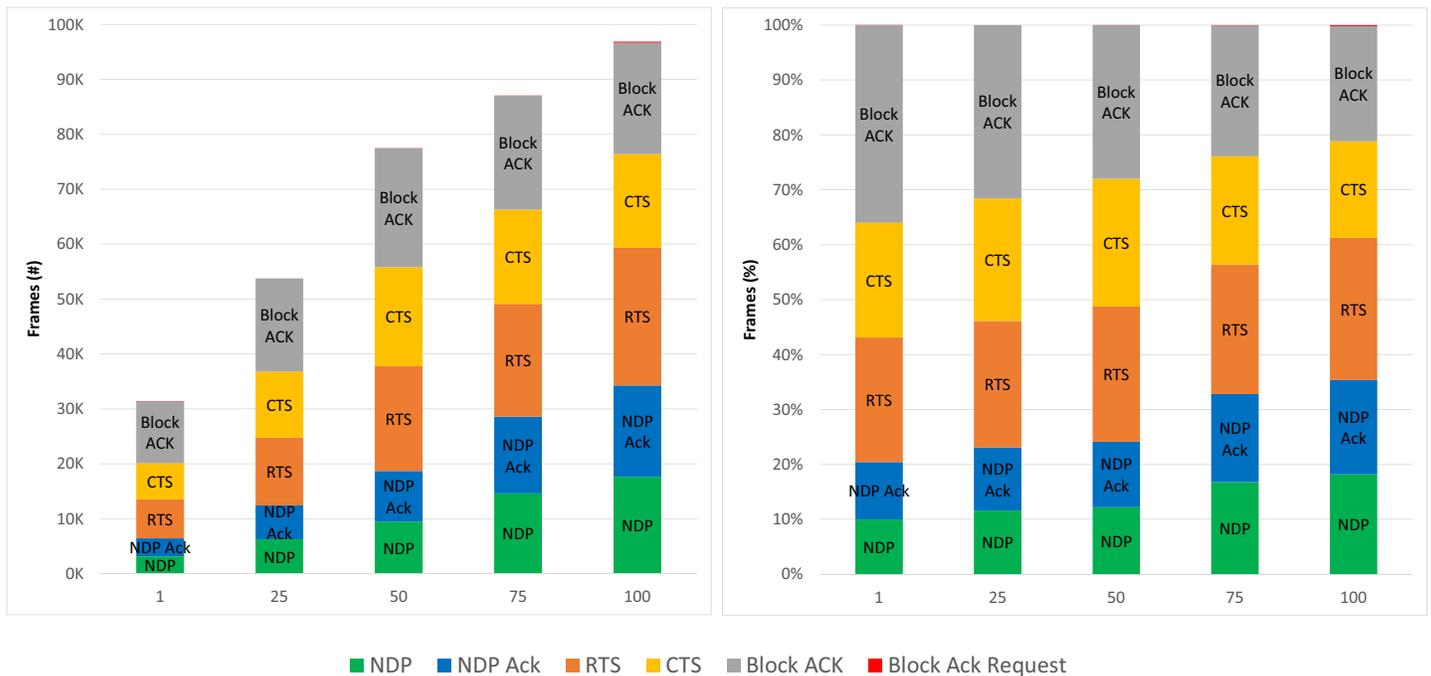


Figure T4-8 Breakdown of 802.11 Control Frame Types (AP-225, 20-MHz Channel, TCP Up)

This chart is extremely interesting. Let’s look at the details:

- 320% Increase in Arbitration:** Of the five control frame types shown, only RTS and NDP require a full arbitration because they initiate a transmission. Clear-to-send, block ack, and Ack are preceded by a SIFS time. RTS+NDP increase by 320% from about 10,000 combined frames for 1 STA, to over 42,000 frames at 100 STAs.
- TXOP Growth:** The total volume of RTS+CTS+BA frames increases by 200% from about 24,000 with 1 STA to over 63,000 at 100 STAs. You learned in [Chapter T-2: What Is “The Channel?”](#) that these frames are the components of a TXOP. They are growing at about the same rate, as would be expected if they were being sent together. This growth implies a 200% increase in the total number of TXOPs from left to right. This result makes sense considering 100 different STAs are all attempting to send full-buffer upstream traffic in this test.

- **80% Decrease in Ratio of Data Frames per TXOP:** If the RTS+CTS+BA are all part of a TXOP, and the number of total data frames is declining, then this must mean that the average number of data frames (for example, MPDUs) per TXOP is decreasing. In fact, we can compute it if we divide the RTS total from this figure into the data frame totals from [Figure T4-6](#). The MPDU-per-TXOP ratio is approximately 60 at 1 STA, and drops by over 80% to about 11.5 at 100 STAs.
- **NDP+Ack Growth:** The PS-NDP frames and Ack frames are related. NDPs are a form of data frame that must be acknowledged. If an Ack is not received the NDP will be resent. The combined volume of these frames grows by over 430% from left to right, from 6,400 to 34,200.
- **Increased Ratio of Power Save Transitions to TXOPs:** The right side of [Figure T4-8](#) shows that the rate of NDP growth is increasing relative to the rate of RTS growth. NDPs increase from 30% of the NDP+RTS total at 1 STA to over 42% at 100 STAs. Stated differently, PS activity is growing faster than TXOPs.

Analysis

Now we have enough information to understand cause and effect. Simply stated, here is the basic chain of causation that we think is happening:

- With increasing numbers of stations in the test, the packing efficiency of the A-MPDUs drops precipitously (for example, each STA is able to send fewer and fewer MPDUs per TXOP).
- This drop drives up the number of TXOPs required to send any given amount of payload data.
- Each TXOP requires a full arbitration, so the total amount of airtime that is consumed by channel acquisition increases linearly as a multiple of the TXOP total.
- The payload airtime fraction of each TXOP is also dropping, which means that each TXOP is less productive.
- Meanwhile, the TXOP increase drives a parallel increase in PS activity, because each STA has to wake up more often to send smaller and smaller amounts of data.
- Additional airtime is lost to power save transitions, each of which requires a full arbitration. This further reduces airtime for TXOPs.

In effect, the channel has to work harder and harder to send less and less data. These effects accumulate at a non-linear rate, and they explain the contention premium.

MIMO Works!

We have descended into the minutia of packet captures in pursuit of our target, and now we climb back up and survey a remarkable aspect of the whole scene.

Look again at [Figure T4-1](#), and notice that it shows something really remarkable. Namely, MIMO works and works very well. Not only that, but MIMO works independent of the number of stations contending for the medium. Different signals can indeed travel different paths at the same time and be successfully recovered at a receiver.

To help this fact stand out, we have used the percentage technique to replot the data. In this case, we will normalize using the 3SS MacBook Pro (MBP). For downstream, upstream, and bidirectional, we used the MBP result as the reference value of 100%. We then divide the 2SS and 1SS clients to obtain relative percentages of the 3SS throughput.

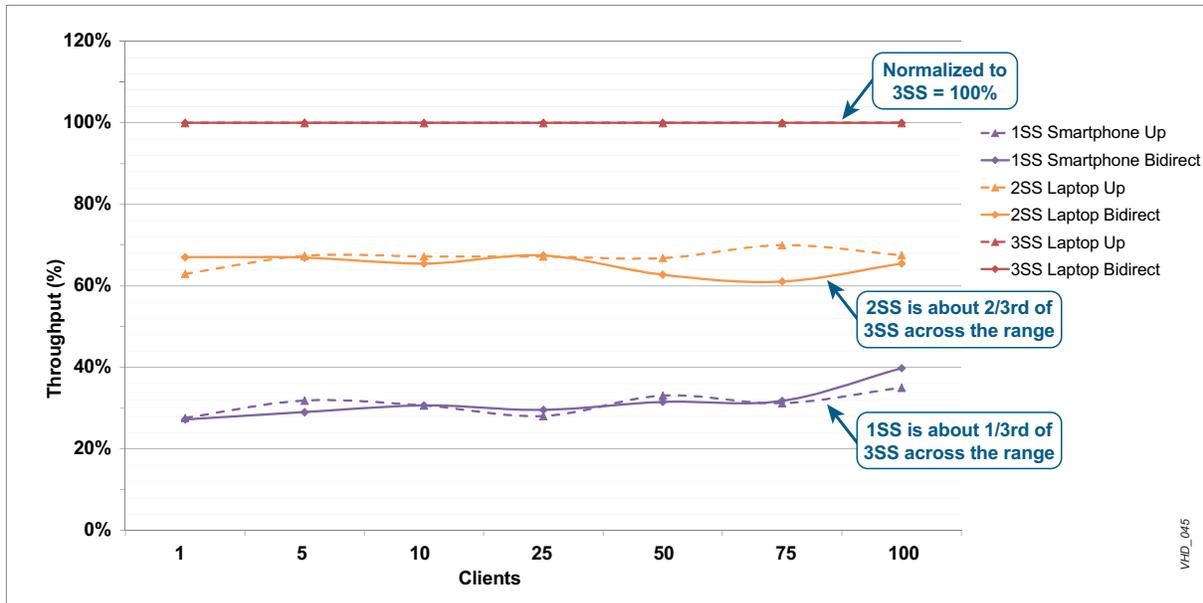


Figure T4-9 Relative Throughput of 1SS and 2SS Clients Compared to 3SS

If each client is fully achieving the maximum data rate for its antenna chain count, then the 2SS clients should be getting about two-thirds of the 3SS result. The 1SS clients should be getting about one third of the 3SS result. These values are exactly what we find. This means that MIMO works across all STA counts in a heavily loaded channel.

It is deeply reassuring to see that MIMO technology is this consistent across a wide range of loads. This consistency is particularly important because many of the most important planned benefits of 802.11ac Wave 2 and future generations depend on MIMO doing what it claims to do.

Conversely, there is a cautionary tale here as well. MIMO is a two-edged sword for VHD environments. MIMO radios are engineered specifically to recover bounced signals, and they are very good at it as [Figure T4-9](#) shows. But in VHD areas, we often do not want any kind of bounce whatsoever, especially in large arenas and outdoor stadia, with carefully chosen external antennas. [Figure T4-9](#) helps explain the point that has been made repeatedly in this VRD, that RF spatial reuse is difficult, if not impossible, to achieve in most VHD environments.

Per-Client Throughput

We have examined aggregate channel throughput for large numbers of devices as a group. How does this throughput translate at the individual device level?

As already stated, most customers who purchase a VHD system define their requirements in terms of minimum per-seat or per-device throughput. When video or other high-bitrate services are required, some customers even attempt to guarantee such minimums contractually. The TST methodology attempts to provide a good-faith estimate that is ultimately derived from the Aruba VHD testbed data published in [Chapter EC-2: Estimating System Throughput](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide. So this section briefly reviews that data.

Figure T4-10 is an alternate view of the aggregate results in Figure T4-1, beginning with 10 STAs and showing the average per-STA throughput for each type of device.

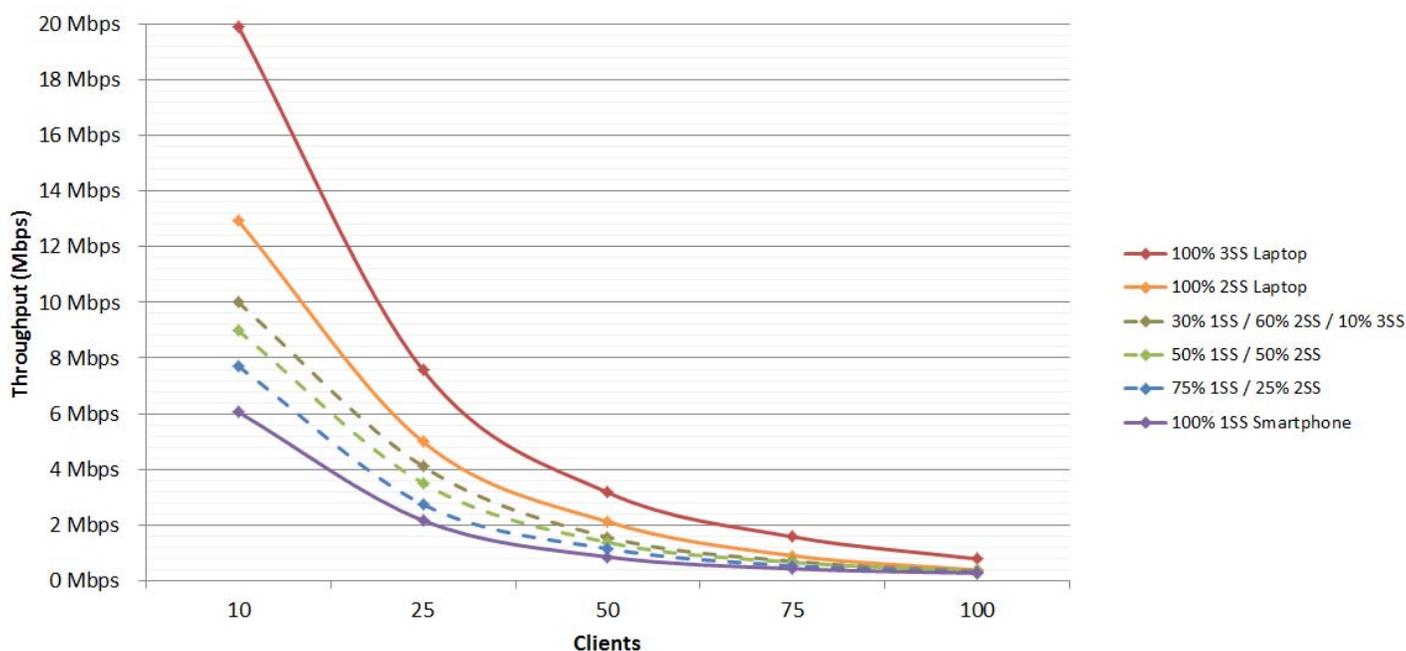


Figure T4-10 Average Per-Device Throughput (AP-225, VHT20, TCP Bidirectional)

This chart was obtained by taking the total throughput for each data point and dividing by the number of devices in the test. Therefore the curves are *averages*. Some devices actually did better and some did worse.

Here are some of the principal insights that may be drawn from the figure:

- It is important to understand the expected user duty cycle in the environment. The per-device throughput varies tremendously with load, so the user experience will be quite different if 25 devices are attempting to use the channel than if 100 devices are.
- The spatial stream capabilities of the devices matter a lot. Multistream-capable clients can experience significantly higher throughput in a favorable channel model.
- At 50 simultaneous devices in a VHT20 channel, most devices can achieve an average of 1 Mbps times their spatial stream count (for example, 1 Mbps for 1SS, 2 Mbps for 2SS, 3 Mbps for 3SS, and so on).
- SLAs that require more than 1 Mbps per device cannot be achieved when more than 50 devices contend at the same time.
- At 100 simultaneous devices in a VHT20 channel, average throughput is measured in *kilobits per second*. A 1SS smartphone will not see any more than 250 Kbps.
- Aruba APs are stable at high concurrent device loads. You should not have any concern about reliability with 100, 200, or even 255 users contending for the medium. As you have learned in this chapter, the channel will run out of capacity before the AP does.

Again – to be crystal clear – these results are unimpaired values obtained in lab conditions with no external interference and a well-ordered channel. The TST methodology in [Chapter EC-2: Estimating System Throughput](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide requires that you apply an impairment factor to these values based on the specific type of environment you are forecasting.

Chapter T-5: Understanding RF Collision Domains

This Theory guide began with a new conceptual approach to thinking about collision domains in an 802.11 system. We defined the physical edge of a collision domain as the point where the signal-to-interference-plus-noise ratio (SINR) of a received signal falls below the 4 dB threshold necessary to decode the Binary Phase Shift Keying (BPSK) preamble of an 802.11 frame. [Chapter T-3: Understanding Airtime](#) and [Chapter T-4: How Wi-Fi Channels Work Under High Load](#) then described in great detail the structure and operation inside a collision domain from both an airtime and a Layer-7 throughput perspective. We complete this guide by returning to collision domains, specifically how they are defined in a physical sense at the radio level.

After you maximize the efficiency of airtime use, the next best strategy to increase total airtime is to achieve spatial reuse. If two or three devices can use the same RF spectrum at the same time, you can increase your available airtime by two or three times. This increase, in turn, produces equivalent increases in overall system capacity.

The basic question this chapter answers is this: *what are the isolation requirements to create truly independent collision domains?* To achieve RF spatial reuse, a collision domain must be independent in time, independent in preamble detection, and independent in energy detection.

We have stated repeatedly that RF spatial reuse is extraordinarily difficult to achieve in practice in very high-density (VHD) areas due to co-channel interference (CCI). To understand how to mitigate CCI, and adjacent-channel interference (ACI) you must first understand the mechanisms by which they degrade performance.

How the 802.11 Clear Channel Assessment Works

When an 802.11 station has data to send and begins the arbitration process, it first uses the clear channel assessment (CCA) mechanism to determine whether the channel is presently idle.

Unlike Ethernet, where collisions can be physically detected, when two or more frames collide on the air, they leave no evidence. 802.11 employs a two-part solution to this problem. A virtual carrier sense and a physical carrier sense must report an idle channel before an 802.11 station initiates the Enhanced Distributed Channel Access (EDCA) contention window process.

- **Physical carrier sense:** For the channel to be idle, the radio must report that no energy is detected above a defined threshold. No kind of radio transmission, Wi-Fi or non-Wi-Fi, can be detected. Per the 802.11ac standard, the energy detection (ED) threshold is -62 dBm for a 20-MHz channel width. This threshold is increased by 3 dB for each doubling of channel width up to 80-MHz.
- **Virtual carrier sense:** For the channel to be idle, the Network Allocation Vector (NAV) must be zero. The NAV essentially is a timer that is always counting down. As long as the NAV is greater than zero, the virtual carrier sense knows that the medium is busy. When any Wi-Fi station decodes a frame with a valid Layer 1 or Layer 2 duration field, it sets the NAV to that value.
 - **Layer 1 Duration:** The L-SIG field in every 802.11 legacy preamble includes a length field that tells other stations how much time the current frame will take on the air. The preamble detection (PD) threshold is 20 dB below the ED threshold.
 - **Layer 2 Duration:** The Ready to Send/Clear to Send (RTS/CTS) frames that begin each data transmit opportunity (TXOP) include a duration field that indicates the total expected length of the entire TXOP including all subframes and the acknowledgement.

Unfortunately, and by design, the virtual carrier sense applies to every frame that any station can decode. Per the 802.11 standard, the lower limit of detection for 802.11ac transmissions is listed in [Table T5-1](#).

Table T5-1 Detection Minimums for 802.11 Clear Channel Assessment

Channel Width	Preamble Detect Threshold (Primary Channel)	Preamble Detect Threshold (Secondary Channel)	Energy Detect Threshold
20 MHz	-82 dBm	-72 dBm	-62 dBm
40 MHz	-79 dBm	-72 dBm	-59 dBm
80 MHz	-76 dBm	-69 dBm	-56 dBm
160 MHz	-73 dBm	n/a	n/a

However, while the standard requires only -82 dBm in the primary channel for PD, in practice, modern radios are vastly improved from this. For example, the Aruba economical 802.11ac AP-205 features a receive sensitivity for legacy OFDM BPSK of -93 dBm. Considering that each 6 dB of power corresponds to a doubling of distance in free space, this 11-dB improvement equates to nearly a 4X increase in PD interference radius as compared with the -82 dBm required by the standard.

Aruba is recommending the use of 20-MHz channels only for VHD areas, so you need not consider secondary channel detection levels for PD or ED.

When Wi-Fi stations go through the EDCA process to count down the random backoff value that they chose for arbitration, they continuously poll the CCA to check that the channel is still idle. If CCA reports that the channel has gone busy, the station is forced to suspend its arbitration until CCA reports that the channel is idle again.

How Co-Channel Interference Reduces WLAN Performance

CCI is simply the assertion of NAV due to detection of an Layer 1 or Layer 2 duration field by the radio.

CCI has an enormous negative impact on overall performance in VHD areas. The impact is large even when channels are not reused inside the VHD area itself, because those same channels typically are reused by nearby APs outside. Walls and floors may provide some isolation, but even highly attenuated 802.11 legacy preambles often can be decoded by the increasingly sensitive radios in modern NICs.

The important concept behind CCI is that any Wi-Fi device that detects an 802.11 preamble on the air is inhibited from transmitting or receiving any other transmission until that frame has ended. It does not matter if the transmitting and receiving stations are part of the same Basic Service Set (BSS). As long as they are on the same channel and can decode the legacy preambles that precede one another's frames, this limitation exists. It also does not matter if the frame payload is corrupted or not, so long as the L-SIG field in the preamble can be successfully recovered.

If two devices that want to transmit at the same time are sufficiently isolated from one another so that they cannot decode one another's legacy preambles, then they may transmit. Figure T5-1 shows both situations and the resulting effect on overall capacity.

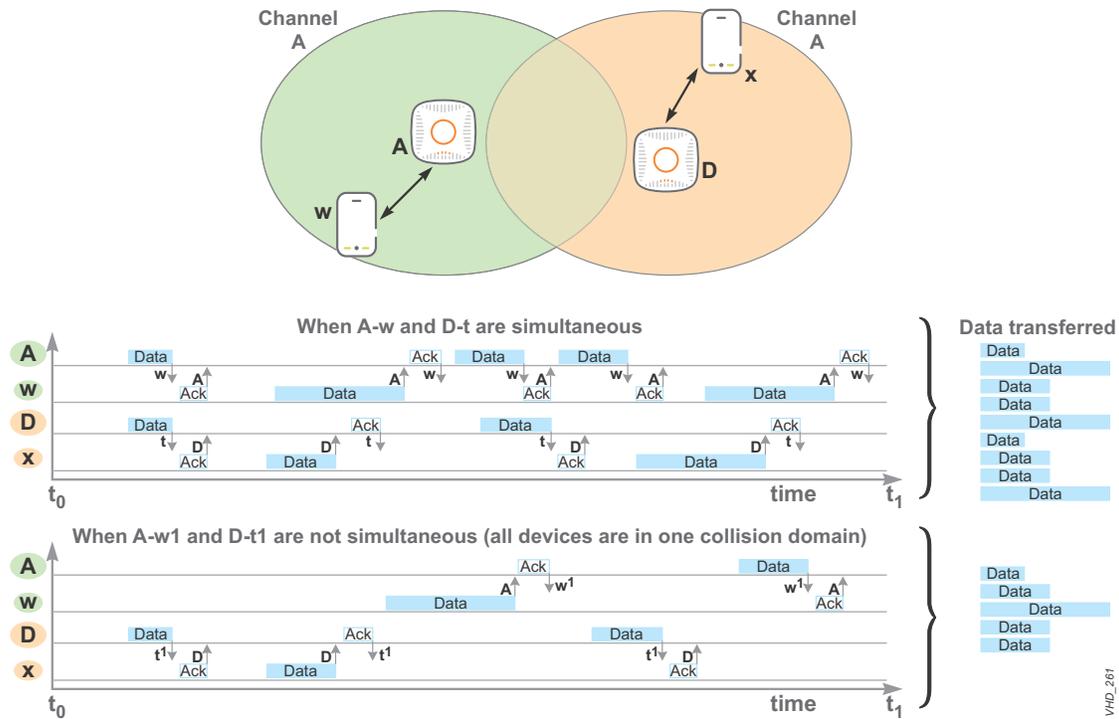


Figure T5-1 Behavior of Two Radio Cells With and Without CCI

This effect is very easy to measure, and is a great home lab project for any WLAN architect. Set up two APs on the same channel, each with one client. Run a speed test on one AP at a time. Then run both APs together. You will find that the total bandwidth of the combined test is about the same as the solo tests, but has been split between the two APs.

How Adjacent-Channel Interference Reduces WLAN Performance

The spectral mask of an 802.11 transmission in the frequency domain allows for significant energy outside the main channel bandwidth. The mask is shown in Figure T5-2. Though it possible to design radios with more precise filters, the resulting increase in cost and physical size of the radio are prohibitive for typical Wi-Fi products.

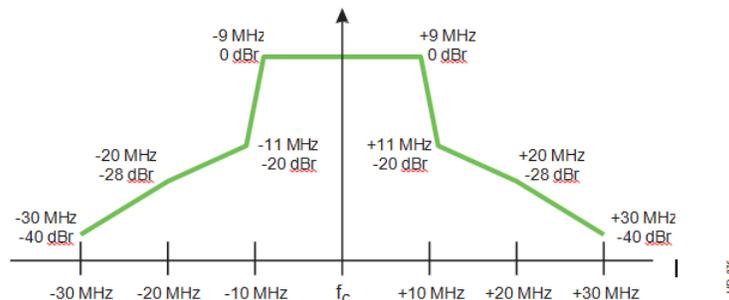


Figure T5-2 802.11 Spectral Mask for a 20-MHz Channel Width

Energy outside the nominal envelope can directly block the channels on either side if it is strong enough, or merely induce noise and increase errors. In most enterprise deployments, ACI is not a factor because APs on adjacent channels are separated by at least 20 m (65 ft). The expected free-space propagation loss at that distance is at least 80 dB in the 5-GHz band, which provides adequate isolation to minimize or avoid ACI performance impacts.

However, in a VHD WLAN with multiple adjacent channel APs and user devices spaced close together, Wi-Fi signals may be received at sufficiently high power levels to cause the ED mechanism to assert CCA busy. In this situation, adjacent channels have effectively become part of the same collision domain. This problem is even more significant for adjacent clients that are even more numerous and more tightly packed than the APs. Therefore, at the densities required for HD WLANs, so-called “non-overlapping” 5-GHz channels actually may overlap.

ACI Interference Example

Consider the VHD WLAN in [Figure T5-3](#), which has three pairs of APs and clients, each one on an adjacent 20-MHz channel. Pairs 1 and 3 transmit heavy-duty cycle traffic such as a video stream. All six stations are configured to use maximum equivalent isotropic radiated power (EIRP).

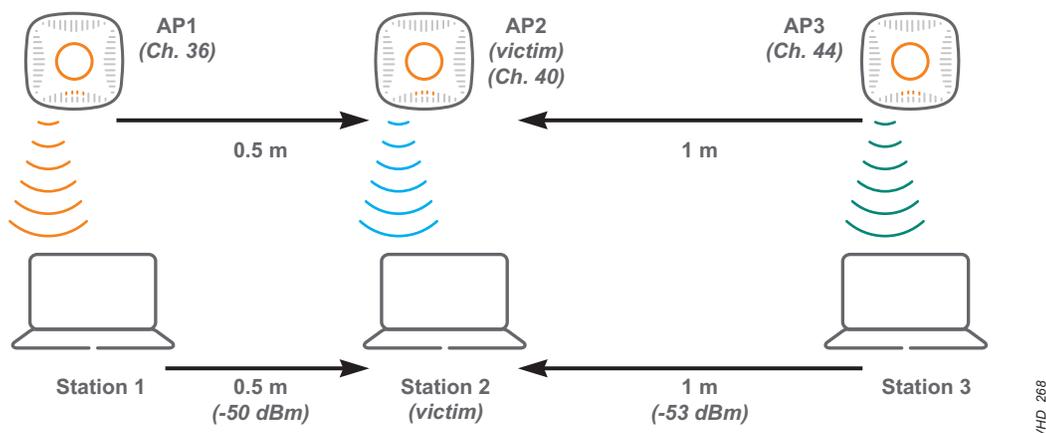


Figure T5-3 ACI Example with APs and Clients at Short Range

VHD_268

AP2 and station 2 on channel 40 now want to transmit and perform a CCA. Pair 1 is only 0.5m (1.5 ft) away, so their transmissions are received at -50 dBm, but signals from pair 2 travel 1m (3 ft) and are received at -53 dBm. Neither AP2 nor station 2 are allowed to transmit because the detected energy exceeds the CCA threshold, even though no one else is using channel 40. [Figure T5-4](#) shows the overlap of the transmit skirts.

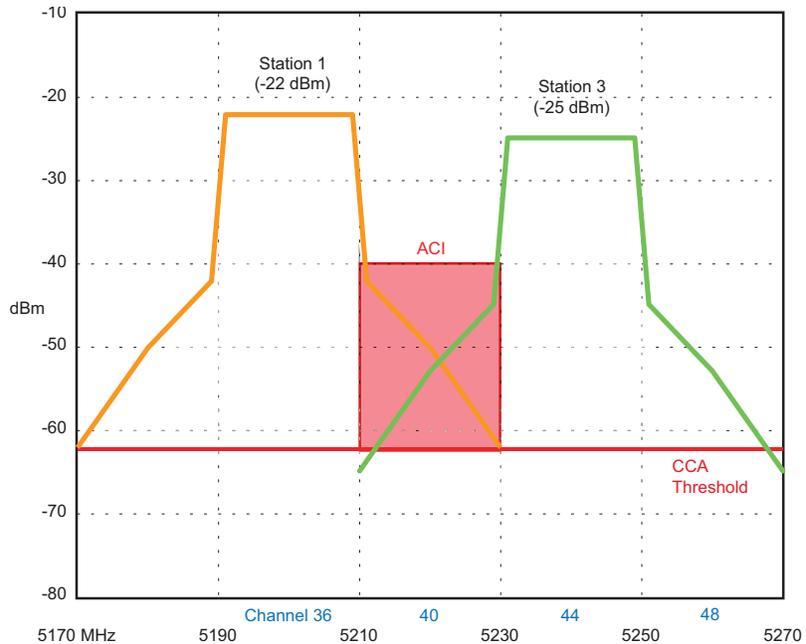


Figure T5-4 Frequency Domain Illustration of ACI at Short Range

Especially inside indoor VHD areas with high multipath conditions, with minimal free space propagation loss between stations, the edge of the skirt can easily be -70 dBm or higher.

Measuring the ACI Impairment

To quantify this effect, Aruba tested ACI in our VHD lab. We subdivided each group of 100 stations into quadrants of 25 devices each, as shown in [Figure T5-5](#). Four AP-225s were configured on adjacent channels (100, 104, 108, and 112). The APs were located on the ceiling approximately 3 meters apart. Maximum EIRP of +23 dBm per chain was used.

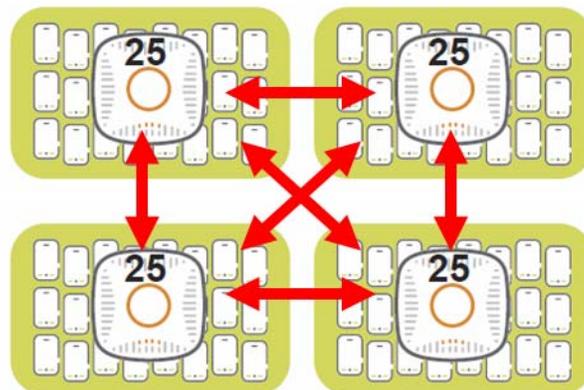


Figure T5-5 Testbed Layout for ACI Test

For each type of client device (1SS smartphone, 2SS laptop, and 3SS laptop), we tested each quadrant individually and added up the results. Then all four quadrants were run at the same time. Figure T5-6 shows the results for the smartphones.

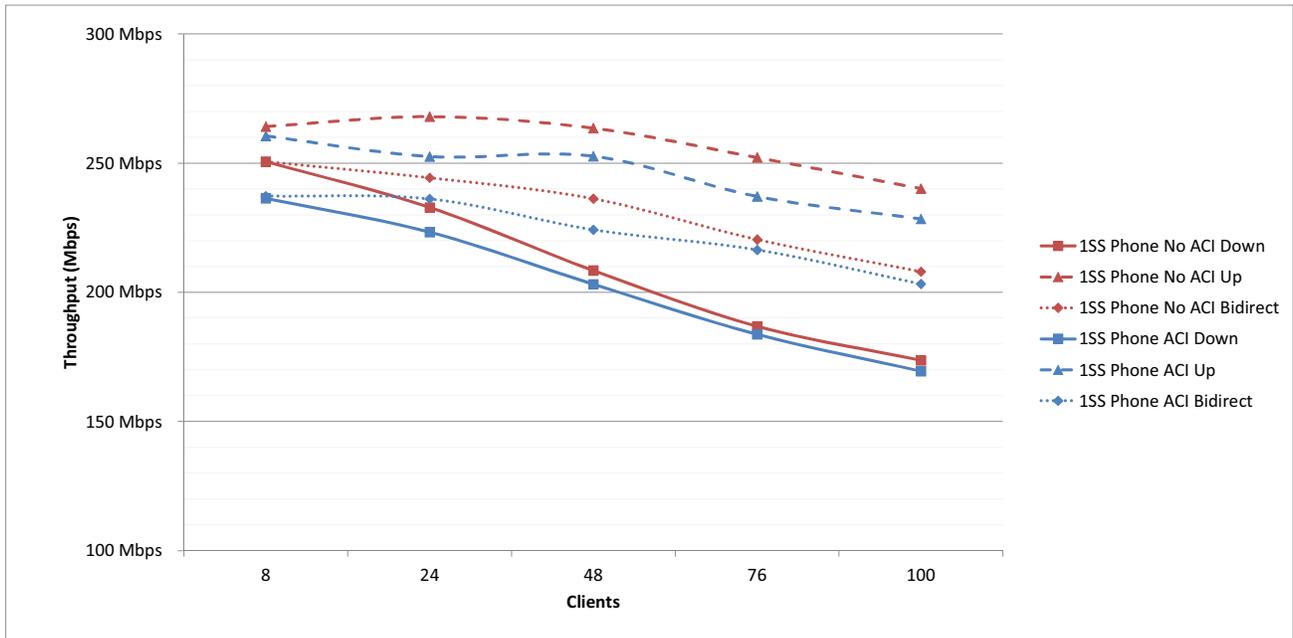


Figure T5-6 ACI Test Results for 1SS 802.11ac Smartphones (AP-225, 80-MHz Channel)

The ACI impairment for the 1SS phones is quite real. Degradation was observed in all tests, with a range of 2% to 10%. Upstream degradation was the most pronounced and show that the individual quadrants were blocking one another at the device level. This makes sense because the client devices are much more tightly packed than the APs.

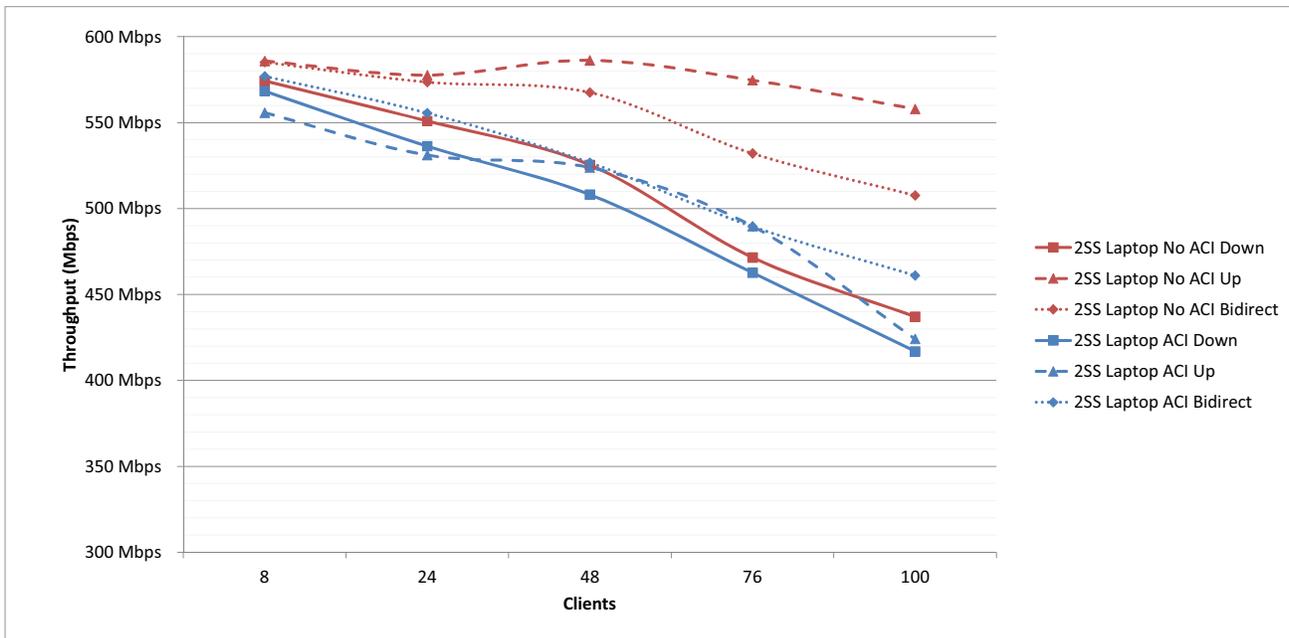


Figure T5-7 ACI Test Results for 2SS 802.11ac Laptops (AP-225, 80-MHz Channel)

Turning to the MacBook Airls (MBAs), we see the effect of the enhanced sensitivity due to multiple receiver chains. Again, upstream traffic is the most heavily affected, and the impairment increases with STA count. Downstream was the least affected, suggesting that the APs were sufficiently isolated at 3 m (10 ft) separation to avoid blocking one another.

The photo of the testbed in Appendix T-A gives you an idea of the physical dimensions of the space and the probable channel model. Our results apply to any intermixed group of clients in an indoor VHD environment with low-to-medium ceilings. Virtually all lecture halls and theaters will be subject to the same type of ACI degradation. ACI impairments are already considered in the suggested impairments provided in step 4 of the total system throughput (TST) process in [Chapter EC-2: Estimating System Throughput](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide.

Interference Radius of Energy Detect and Preamble Detect

The maximum interference distance of the energy detect threshold is significantly different than that for the NAV. You must understand this distinction when you plan VHD areas with more than one AP on the same channel. While ED essentially is a short-range phenomenon, PD is a very long-range phenomenon.

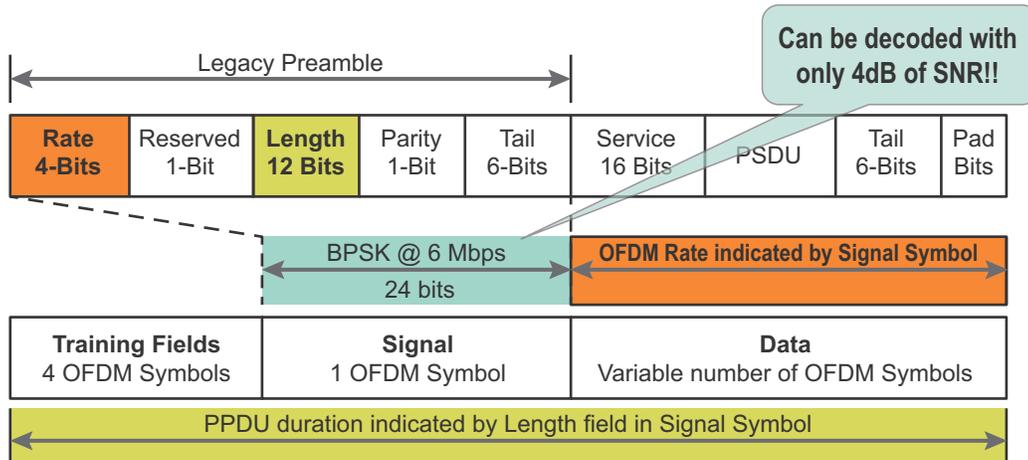


Figure T5-8 Different Data Rates Are Used in Preamble and Payload

We have seen that the legacy preamble uses the 6 Mbps BPSK modulation, which requires just 4 dB of SINR to decode. The Layer 1 duration field is contained in the L-SIG field in this preamble. As a result, it can be decoded to an extraordinarily far distance, and even the rest of the frame payload uses a much higher data rate. For example, a preamble that arrives at -86 dBm results in CCA asserting busy with a noise floor of -90 dBm.

A far more concrete and compelling way to think about PD-based CCI interference is to define it in terms of the cell edge RSSI. It is a widespread best practice to design 802.11 cell edges to be -65 dBm, which yields an SINR of 25 dB with a -90 dBm noise floor.

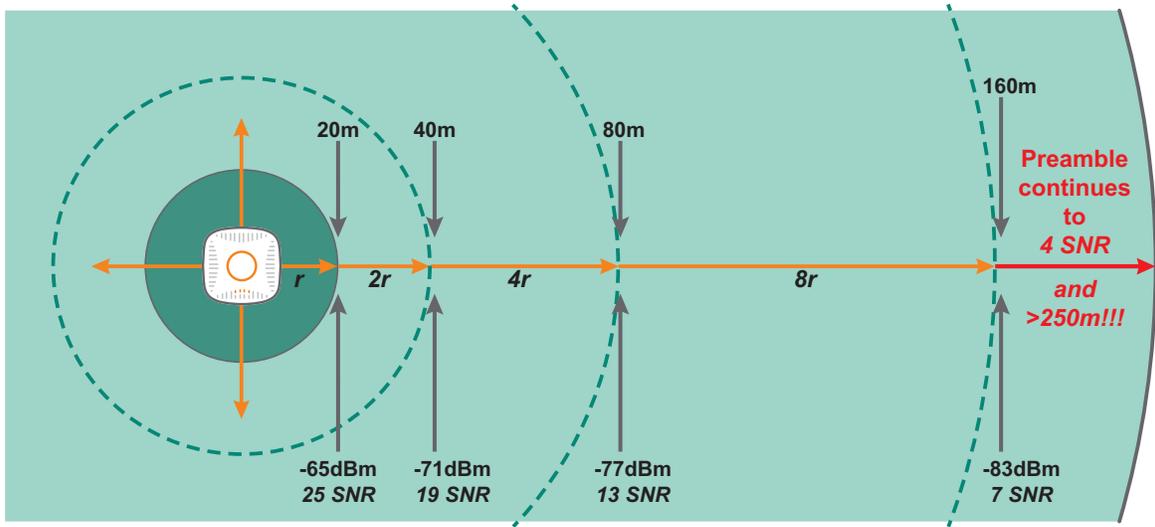


Figure T5-9 Interference Radius of PD Relative to AP Cell Edge (assuming -90 dB NF)

We know from the 6 dB rule that distance doubles for every 6 dB of power increase. Applying this rule, you can see in [Figure T5-9](#) that any cell that is designed to the -65 dBm cell edge criteria will have a NAV interference radius of more than 250 m (820 ft) in free space!

Contrast the PD interference distance with interference radius of the ED threshold. From the free space path loss formula, we can calculate that the maximum ED interference range is about 4 m (13 ft) in the 5-GHz band and 8 m (25 ft) in the 2.4-GHz band.

A Real World Example of 802.11 Radio Power

As an architect you have performed and reviewed the results of many RF site surveys. However you don't often get to survey a large open air facility like a football stadium. Radio signals travel at the speed of light (about 3 nanoseconds per meter) so it takes just about 300 ns to cross a football field. [Figure T5-10](#) shows an actual 802.11 heatmap from a football stadium. The AP was placed underneath a seat at the mid-field line. At full power of +23 dBm EIRP per chain, the AP beacons can be detected as strongly as -75 dBm at the very top of the highest row in the upper sections (and a +15 dB SINR).

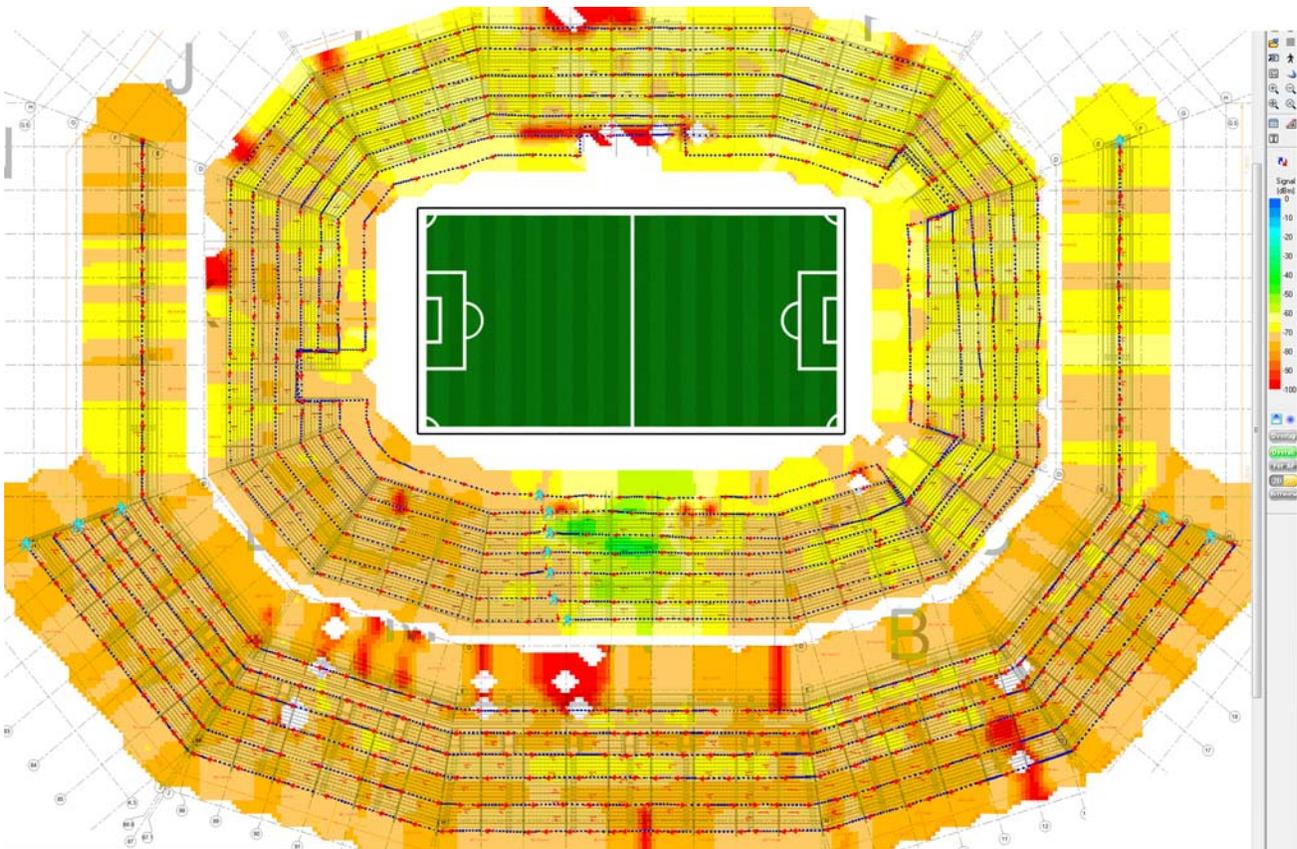


Figure T5-10 Received Signal Power of Midfield AP in a Football Stadium

This survey software is measuring beacons that are being sent out at the 6 Mbps BPSK rate. Therefore it is also effectively measuring the signal strength of legacy preambles! You can see that PD interference is quite real and extraordinarily powerful. Imagine your results in a large indoor environment with walls and a roof to enhance the multipath conditions.

Containing CCI By Trimming Low Data Rates Is a Myth

One of the greatest myths repeated by engineers across the Wi-Fi industry is that the size of a cell can be “shrunk” by eliminating low data rates from the BSS transmit rate set. This is not true, at least from a PD interference perspective. [Chapter T-2: What Is “The Channel?”](#) explained this myth in detail.

Removing the 6, 12, and even 18 Mbps data rates from the BSS has no effect on the legacy preamble rate, which must always use BPSK. So removing low data rates has no effect on the PD/NAV interference radius.

The true purpose of removing the low rates is to:

- Force mobile clients to roam sooner than they otherwise might by removing options from their rate adaptation algorithm.
- Reduce airtime consumption by 802.11 control frames by forcing stations in a BSS to use a higher minimum rate.

Minimum Requirements to Achieve Spatial Reuse

One of the main goals of this chapter is to convey the technical reasons why spatial reuse is so difficult to achieve in practice with Wi-Fi. The design of the 802.11 CCA mechanism intentionally uses a robust modulation to ensure that the widest number of stations decode each transmission. However, this fact is buried deep in the standard, and most Wi-Fi engineers have never been exposed to it.

That said, spatial reuse is real and can be achieved in specific types of facilities and crowd conditions. Having a good understanding of the actual technical challenges is the first step.

The second step is a good RF design that follows these principles:

- Use as many channels as possible, including DFS channels, to reduce the overall amount of channel reuse needed and increase the distance between same-channel APs in the channel plan.
- Choose a coverage strategy that minimizes CCI from other APs near the VHD area, and consider the construction of the building.
- Ensure that the facility meets the minimum requirements for spatial reuse:
 - Large physical volume (at least 10,000 seats)
 - Suitable mounting locations for APs and external antennas
 - Site survey to validate feasibility of spatial reuse with a crowd present
- Engage your Aruba systems engineer or an experienced wireless integrator who has the training and tools to properly design it.

The third step is to use VHD configuration best practices, combined with the Cell Size Reduction feature in ArubaOS to limit exposure to CCI.

Controlling ACI

The primary method to control ACI is to ensure maximum possible physical separation of adjacent channel APs. This requirement is the reason we recommended to evenly distribute APs throughout the coverage area in [Chapter P-3: RF Design](#) of the *Very High-Density 802.11ac Networks Planning Guide*, and to ensure a well-distributed channel plan in [Chapter EC-4: Channel and Power Plans](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide.

A secondary solution may be to use the minimum amount of transmit power necessary for the size of the VHD area. However, it is far more likely that reducing power results in lower SINRs, which in turn drop the data rates for many clients. This result is worse than taking the ACI penalty. Also, it is of critical importance that you run the 5-GHz radios by +6 dB or +9 dB higher than the 2.4-GHz band to improve self-steering. This requirement further constrains your flexibility to play with power at the AP.

On the client side, the sad truth is that most major operating systems on the market today do not respect the 802.11h TPC power constraint message. This reason is why Aruba does not recommend that you enable it. Virtually nothing can be done about client-side transmit EIRP beyond leveraging crowd loss and structural loss in your RF design.

Appendix T-A: Aruba Very High-Density Testbed

Results from the Aruba very high-density (VHD) testbed have been presented in each of the three main guides of this VRD. The VHD testbed was built specifically for the authoring of this VRD to validate our design recommendations. This appendix explains the testbed design and test plans for those who want to replicate our results.

Testbed Justification

The need for real-world, open-air performance data when planning a VHD wireless network cannot be overstated. Such data takes out much of the guesswork, but it can be expensive and time-consuming to obtain because it requires hundreds of devices, dedicated network hardware, skilled engineers, shielded test facilities, and specialized measurement tools.

802.11ac greatly magnifies the need for this data because of the many new options for things like data rates, spatial streams, channel widths, aggregation, and beamforming.

Recognizing this challenge and the broad-based marketplace need, Aruba undertook a research program into client performance in VHD environments as part of its industry leadership efforts. Our goal is to assist our customers, our partners, and our own engineers to better understand and succeed at very high-density deployments.

Testbed Design

The testbed is shown in [Figure T-A1](#), and is made up of 300 brand new, native 802.11ac devices:

Make	Model	Radio	Spatial Streams	Quantity
Samsung	Galaxy S4	BRCM 4335	1SS	100
Apple	MacBook Air	BRCM 4360	2SS	100
Apple	MacBook Pro	BRCM 43460	3SS	100

The devices are placed in eight rows with spacing between units of 15-30cm (6-12 in). All three device types are intermingled with consistent spacing between them.



Figure T-A1 Aruba VHD Testbed with 300 Stations

Topology

The topology of the network is shown in [Figure T-A2](#). We employ two parallel sets of APs, each with its own Aruba 7220 controller. One set of APs was used as the data plane to carry test traffic. The other set of APs was used for packet capture to perform analytics. An Aruba S2500-48P switch was used to power the APs. The controllers were connected to the switch via 10G Ethernet links.

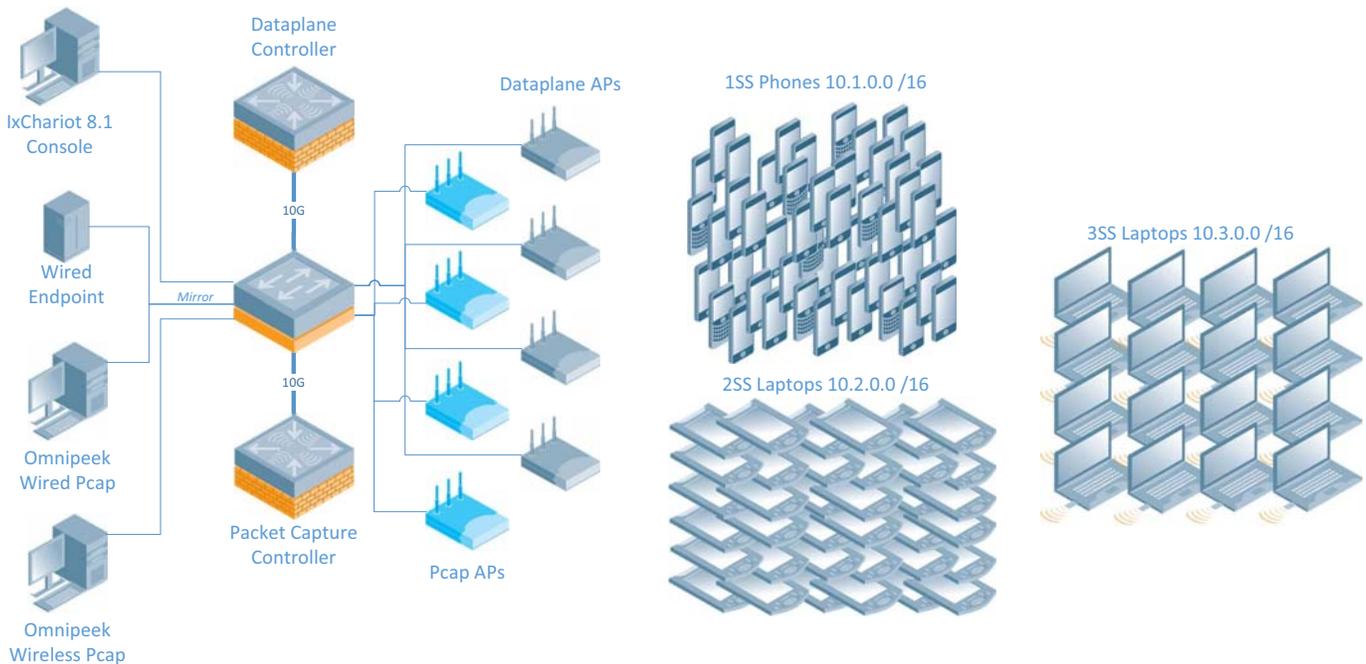


Figure T-A2 Aruba VHD Testbed Topology

All results that are presented in this VRD were run with jumbo frames enabled and frame aggregation enabled.

Channels

Our testbed has exclusive use of 80 MHz of spectrum from 5490 – 5570 MHz. This corresponds to these channelizations:

- VHT20 / HT20 - 100, 104, 108, 112
- VHT40 / HT40 - 100+, 108+
- VHT80 - 100E

The building with the testbed is in a remote part of the Aruba campus with little nearby interference. Channels are swept daily to ensure cleanliness.

SSID Configuration

The charts presented in the VRD were taken with the recommended SSID configuration from [Chapter EC-2: Estimating System Throughput](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide. That configuration is:

```
wlan ssid-profile "hdtest100-ssid"
  essid "HDTest-5"
  a-basic-rates 24 36
  a-tx-rates 18 24 36 48 54
  max-clients 255
  wmm
  wmm-vo-dscp "56"
  wmm-vi-dscp "40"
  a-beacon-rate 24
!
wlan ht-ssid-profile "HDtest-htssid-profile"
  max-tx-a-msdu-count-be 3
!
rf dot11a-radio-profile "hdtest100-11a-pf"
  channel 100
!
```

Automation

Ixia IxChariot 8.1 was used as the automation platform for the tests. IxChariot generates repeatable IP traffic loads and provides a control plane for the tests. With the 8.X release, Ixia has moved to an OVA-based deployment model with the control software running on a dedicated virtual machine. A dedicated laptop was used as wired endpoint that was connected via Gigabit Ethernet to the switch. The IxChariot endpoint was installed on all of the stations in the testbed.

For TCP tests, the *throughput.scr* script was used. For UDP, the *UDP_throughput.scr* was used. 30 second durations were used.

The number of flows or streams used on each client varied according to test objective and the number of stations in the test. Aruba conducted upstream, downstream, and bidirectional test cases for most tests.

What is a Client Scaling Test?

Aruba calls this type of test a “client scaling test.” Client scaling tests measure performance with increasing numbers of real clients in open air to characterize behaviors of interest to a wireless engineer.

For this VRD, test runs generally included scaling with 1, 5, 10, 25, 50, 75, and 100 clients.

Multiple test runs make up a test case. Each test case changed one aspect of the testbed at a time to study how that particular variable affects performance.

Some of the major variables we studied include:

- AP-205, AP-215, AP-225, and AP-275 access point models
- 20-MHz, 40-MHz, and 80-MHz channel widths
- One, two, and three spatial stream clients
- Open authentication vs. WPA2 encryption
- TCP vs. UDP
- 64-byte, 512-byte, 1024-byte, 1514-byte, 3,028-byte, and 4,500-byte frames
- Airtime fairness enabled and disabled

Scaling clients for each variable provides an intrinsic consistency check on the data because occasional bad runs are quite obvious.

Why No High Throughput or Legacy Clients This Time?

The 2010 edition of this VRD studied not only the new 802.11n High Throughput (HT) technology, which was brand new, but also older 802.11a/g clients. Combinations of HT and non-HT clients were also studied because most environments had a mix of both.

These combinations were necessary because there were fundamental differences at both the PHY and the MAC layer between 802.11a/g and 802.11n. HT modulations, multiple input, multiple output (MIMO), frame aggregation, and wider channels are just a few examples.

By contrast, 802.11ac is more of an incremental extension of 802.11n:

- The PHY data rates of 802.11n are identical to 802.11ac. Only 256-QAM is new (MCS8 and 9).
- 802.11ac clients fall back to standard 802.11n modulation and coding scheme (MCS) rates when 256-QAM is not available
- MIMO has been enhanced from 4 streams in 802.11n to a limit of 8 streams.
- Aggregated MAC protocol data unit (A-MPDU) aggregation is extended and made mandatory.

There are certainly important differences between 802.11ac and 802.11n. However, for the purposes of this VRD, they do not alter our overall conclusions or our design recommendations.

Comparing with Other Published Results

The performance charts that are published with this VRD cannot be directly compared to marketing white papers from Aruba or other vendors.

Our goal was to replicate a real-world high-density channel to better characterize how it performs and how best to optimize it.

As a result, we made certain specific configuration decisions that hurt our results. We also made changes that helped our results. Few, if any, of these changes are part of typical marketing performance reports.

Critical differences that negatively impacted our results include:

- **250 STAs associated at all times:** Generally, the test radios were fully loaded with associated clients, even if the test itself used a much smaller number of stations. These extra associated stations produced additional low rate 802.11 power-save, management and control traffic. This traffic would be typical of a VHD environment with multiple same-channel APs. This traffic is more realistic, but it reduced our results.
- **Narrow channels were used:** Most of the charts reprinted in the VRD were taken in a 20-MHz bandwidth in keeping with the recommendations of [Chapter EC-2: Estimating System Throughput](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide. The APs are capable of dramatically higher performance.

Critical configurations that positively impacted our results include:

- **Enhanced minimum data rates:** All tests published use our recommended VHD SSID configuration of 24 Mbps minimum data rate. This offset the loss from having extra stations associated.
- **Exclusive DFS channels:** All tests were conducted on channels 100 – 112, which are subject to DFS rules. Our primary reason was to obtain clean air and improve the repeatability of our results. However, clients behave differently on DFS channels, in particular they probe less. This behavior offset the loss from having extra stations associated.

In summary, the performance charts and tables provided in this guide are purely for the purposes of optimizing performance in VHD environments. They should not be compared to test reports using different conditions or having a different objective.