Spectral Coexistence of Wi-Fi Networks with Radar Systems

Benjamin Cizdziel

A thesis
submitted in partial fulfillment of the
requirements for the degree of

Master of Science

University of Washington

2015

Committee:

Sumit Roy

Thomas Henderson

Program Authorized to Offer Degree:
Electrical Engineering

University of Washington

Abstract

Spectral Coexistence of Wi-Fi Networks
with Radar Systems

Benjamin Cizdziel

Chair of the Supervisory Committee:
Professor Sumit Roy
Electrical Engineering

With crowded frequency bands and high mobile data demand, Wi-Fi networks would benefit from more available spectrum. Underutilized radar spectrum due to the directionality and rest periods of rotating radar makes it a viable candidate for spectrum sharing with Wi-Fi. But if operating on the same bands, radar signals would interfere with Wi-Fi receivers, reducing WLAN performance.

This thesis investigates Wi-Fi network performance when sharing spectrum with radar. Both link and system level software simulators were developed to study the effects of coexistence, providing highly configurable platforms with repeatable results. The link simulator comprises a signal level model of baseband Wi-Fi transmitter and receiver chains with injected radar pulse interference, and is used to determine the effects of radar interference on PHY communication performance. The ns-3 network simulator is used for full Wi-Fi system performance analysis, with enhancements added to model radar interference using mapped link simulation results. The source code for both simulators will be hosted on the UW Fundamentals of Networking Laboratory website at: **http://depts.washington.edu/funlab/**.

From analyzing the simulation results, insight was gained about the impact of radar interference on Wi-Fi and how different system parameters affect the performance. Techniques are proposed to mitigate Wi-Fi performance degradation amid radar interference.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

GLOSSARY

AGC: Automatic Gain Control

AP: Access Point

AWGN: Additive White Gaussian Noise

BPSK: Binary Phase Shift Keying

CCA: Clear Channel Assessment

CSMA/CA: Carrier Sense Multiple Access with Collision Avoidance

CTS: Clear to Send

CW: Contention Window

DCF: Distributed Coordination Function

DIFS: DCF Interframe Space

EM: Electromagnetic

FDM: Frequency Division Multiplexing

FEC: Forward Error Correction

FER: Frame Error Rate

FFT: Fast Fourier Transform

HT: High Throughput

ICI: Inter-Carrier Interference

IDFT: Inverse Discrete Fourier Transform

IFFT: Inverse Fast Fourier Transform

IEEE: Institute of Electrical and Electronics Engineers

INR: Interference-to-Noise Ratio

IOT: Internet of Things

ISI: Inter-Symbol Interference

MAC: Media Access Control

MCS: Modulation and Coding Scheme

MIMO: Multiple Input Multiple Output

MMSE: Minimum Mean Square Error

OFDM: Orthogonal Frequency Division Multiplexing

PHY: Physical Layer

PRF: Pulse Repetition Frequency

PRI: Pulse Repetition Interval

QAM: Quadrature Amplitude Modulation

QPSK: Quadrature Phase Shift Keying

RSS: Received Signal Strength

RTS: Request to Send

SCM: Single-Carrier Modulation

SIFS: Short Interframe Space

SISO: Single Input Single Output

SNIR: Signal-to-Noise-Plus-Interference Ratio

SNR: Signal-to-Noise Ratio

STA: Station

UDP: User Datagram Protocol

WLAN: Wireless Local Area Network

ACKNOWLEDGMENTS

The author wishes to express sincere appreciation and gratitude to Hossein Safavi, Professor Sumit Roy, and Professor Thomas Henderson, who have guided him throughout his past year of research.

Chapter 1

INTRODUCTION

Within the last decade, increased use of Internet-connected mobile devices and applications has led to an influx of data, which has been dubbed "big data", to be communicated across the global network. Additionally, many tech companies are providing services, applications, and devices which generate large amounts of user data that these companies analyze to gain insight on their users and improve their products. The current Internet of Things (IoT) boom is leading to many new types of connected devices for various novel purposes, which yields even more data and network load.

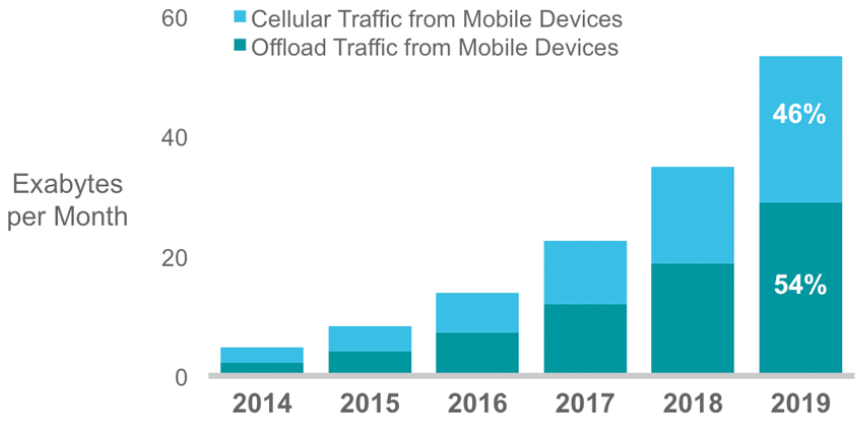

Figure 1.1: Mobile traffic and offload forecast. Retrieved from [1]

The majority of today's mobile devices, namely smartphones and tablets, rely on wireless technology such as Wi-Fi and LTE for their connections and do not even provide the capability of using a wired connection. This is causing the wireless infrastructure to become heavily loaded, as vast amounts of devices are downloading and uploading data. This increased load

causes heavy wireless interference and loaded channels, reducing the data rates that each user experiences. Cisco's future data traffic prediction for mobile devices is shown in Figure 1.1. It depicts that the amount of mobile traffic will increase exponentially over the next few years. The figure also illustrates the percentage of offload traffic steadily increasing and eventually overtaking cellular traffic, where offload refers to devices using Wi-Fi or small-cell networks for their data instead of cellular networks. Wireless local area networks (WLANs) are even taking on some of the load from traditional cellular voice calls, as voice over WLAN (VoWLAN) adoption is accelerating. It follows that WLANs, most being Wi-Fi networks, will be even more heavily utilized in the coming years than they are today.

$$C = B \cdot log_2 (1 + SNR) \tag{1.1}$$

One simple method for reducing wireless congestion and increasing performance is to allow more of the frequency spectrum to be used for wireless communications technology. Eq. (1.1) is the Shannon-Hartley Theorem, which specifies the upper bound on the channel capacity C (in bits per second) for a communications channel with additive white Gaussian noise (AWGN). It states that the channel capacity is directly proportional to the bandwidth B of the channel. Thus, by increasing the bandwidth used for a communications technology, in turn using more of the frequency spectrum, higher data rates can be achieved. However, the electromagnetic frequency spectrum is a crowded place. Figure 1.2 displays the current frequency allocations for the United States. It is clearly seen that the spectrum is packed tightly with wireless technologies, with many frequency bands being allocated for multiple uses. The majority of the spectrum is licensed for specific uses, with no other wireless technologies legally allowed to utilize those frequencies. The primary reason for these allocations are because technologies using the same frequencies of the spectrum will interfere with each other, so using many in the same frequency band can cause highly degraded performance for each one of them. Yet, some wireless technologies underutilize their allocated spectrum, leaving "wasted" spectrum that could be used for other needs, such as communications. One of the primary technologies of interest in this thesis, radar, is a prime example of this.

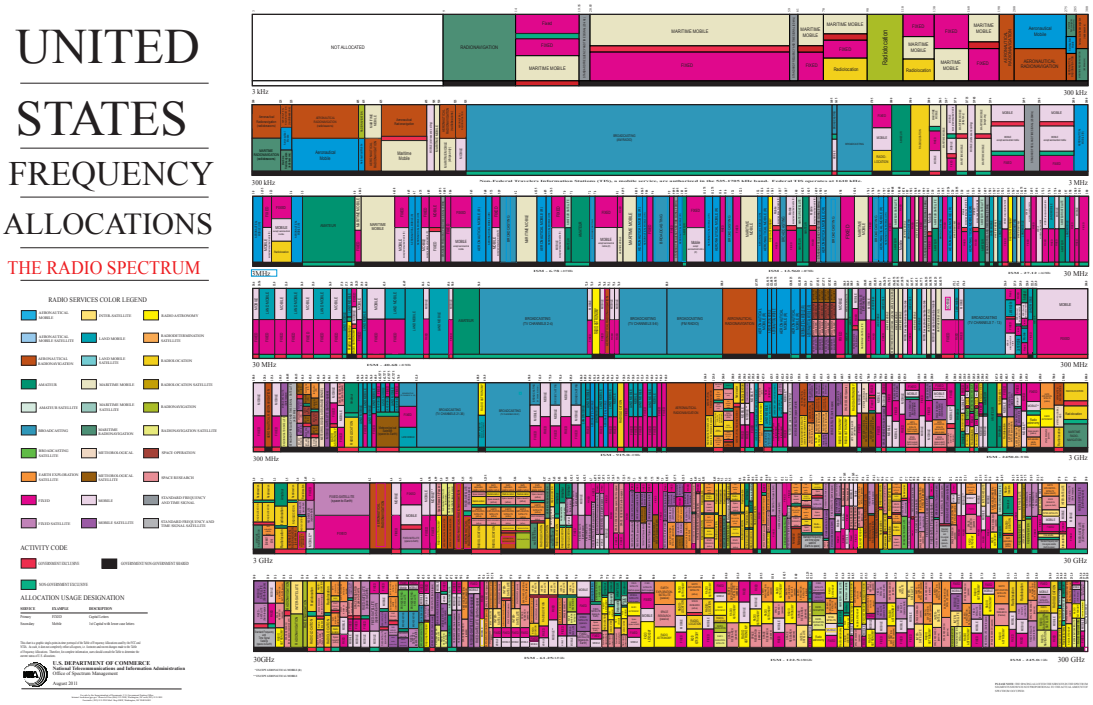

Figure 1.2: United States frequency allocations as of October 2011. Retrieved from [2]

Radar has two characteristics of interest to the author that make its licensed frequency bands potentially feasible for coexistence with other technologies: its directionality and its long listening period. Since radar antennas are typically highly directional and constantly rotating, the majority of interference coming from radar will be at times when its antenna is pointed near the direction of the wireless receiver of the other technology of interest. Much less interference will occur when the radar is directed far away from the receiver, which leaves the opportunity open for high performance spectrum utilization by the other technology. Radar systems detect objects by transmitting a short energy pulse in a direction and waiting, or listening, for its return after reflecting off of potential objects. Systems are typically configured with a listening period much longer than the pulse transmission period. Because of this, no matter what direction the radar is pointed in, it will not be occupying spectrum for the majority of the time. This leaves temporal openings in spectrum that

could be efficiently utilized for other wireless technologies. The United States Defense Advanced Research Projects Agency (DARPA) recently initiated the Shared Spectrum Access for Radar and Communications (SSPARC) program, whose goal is to make spectrum sharing between military radars and communications systems achievable [20]. This initiative further substantiates the prospect of radar band spectrum sharing.

As presented above, Wi-Fi networks are becoming increasingly utilized as more devices become Internet-connected, and would benefit from less loaded channels and more spectrum access. Current Wi-Fi technology primarily uses the unlicensed ISM (industrial, scientific, and medical) frequency bands at 2.4 GHz and 5 GHz. As they are unlicensed, many wireless technologies utilize these bands, especially 2.4 GHz which is home to Bluetooth, cordless phones, microwave ovens, medical devices, and many more. Multiple input multiple output (MIMO) antenna techniques such as beamforming and receiver diversity have been included in the latest Wi-Fi standards to combat performance degradation in crowded spectrum, but do not eliminate the underlying problem. With more spectrum availability, Wi-Fi networks would be less likely to interfere with each other and with other technologies. Additionally, the possibility would arise for Wi-Fi users to be allocated more spectrum bandwidth. Both of these phenomena would contribute to higher experienced data rate.

One can argue that millimeter-wave frequencies and above can be used for wireless communications technologies such as Wi-Fi to solve the problem of spectrum crowding and increased capacity demands, since these bands tend to be less crowded or even un-allocated. Unfortunately, this is difficult to accomplish with today's technology, as the majority of current WLAN hardware is designed for lower frequencies and extremely high frequency transmissions are more heavily impacted by channel effects such as fading and attenuation through obstructions. Thus, sharing bandwidth with licensed spectrum may prove necessary to keep up with data demands.

This thesis investigates spectrum sharing between radar systems and Wi-Fi networks. Specifically, the effect of radar interference on Wi-Fi networks using the radar S-band (2 - 4 GHz) is studied. Schemes are explored to mitigate the performance degradation experienced

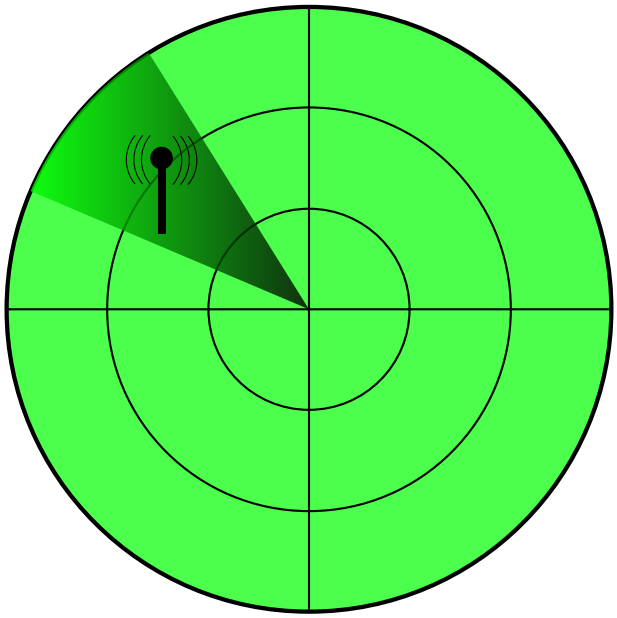

Figure 1.3: Rotating radar beam interfering with a Wi-Fi network.

by Wi-Fi networks (hereinafter *mitigation schemes*) in these scenarios. The aim is to provide a detailed analysis that helps pave the way for successful radar-Wi-Fi coexistence.

1.1 Contribution

Fully-realized hardware deployments to study wireless communications and networking performance are often expensive to put together and rigid in their configuration. Radar systems in particular tend to be extremely costly, have confidential inner workings (especially military radar), and require government licensing. In addition, it is typically difficult to obtain repeatable results from these experimental systems. Software simulators are popular in communications and networking research because with sufficiently accurate models, researchers are able to obtain useful, repeatable data by running easily modifiable experiments. Accordingly, software simulators are used in this research to study radar-Wi-Fi coexistence in the S-band.

The software simulations performed in this research are broken up into two categories:

signal level and system level. The signal level, or link simulation (hereinafter *link sim*), has been built from the ground up by the University of Washington (UW) Fundamentals of Networking Laboratory. It models both Wi-Fi transmitter and receiver chains, with blocks that implement specific functions for each. The transmitter generates a Wi-Fi packet waveform, sends it through the link sim's wireless channel model, and the receiver attempts to successfully process and decode the signal. Radar waveforms are also modeled by the link sim and interfere with the Wi-Fi signals at the receiver input. This enables evaluation of the performance impact of radar interference on Wi-Fi communications at the physical layer (PHY).

The open source network simulator ns-3 is used as the system level simulation platform. ns-3 is a discrete-event network simulator that models the full protocol stack for various Internet systems, including Wi-Fi [21]. It is used in this research to run simulations of Wi-Fi networks in different scenarios and measure useful network performance metrics. A Wi-Fi PHY model is included in ns-3, but it does not provide the same level of customization of the transmitter, receiver, waveforms, and interference as the link sim described above. Hence, error rate results from the link sim are mapped to the ns-3 model with additional mechanisms added to allow for detailed system simulations of Wi-Fi networks sharing spectrum with interfering radar systems.

The source code for both simulators will be hosted on the UW Fundamentals of Networking Laboratory website at: **http://depts.washington.edu/funlab/**.

In addition to the simulation and results analysis, mitigation schemes are proposed as suggestions for improving Wi-Fi network performance amidst radar interference. Two of the proposed mitigation schemes involve media access control (MAC) layer adjustments, while two pertain to the PHY layer.

1.2 Related Works

Related literature regarding coexistence between radar and wireless communication systems takes various approaches and focuses. Some works study the subject from the perspec-

tive of radar system performance. In [22], an estimation method is proposed for cancellation of interfering WiMAX signals from returning radar pulses at the radar receiver to improve detection performance. A technique based on RTS/CTS is explored in [23] to prevent Wi-Fi nodes from transmitting frames when a weather radar is directed toward it in order to reduce the chance of Wi-Fi signals perturbing radar measurements.

Other literature analyzes the results of hardware test bed deployments. A test bed is used in [26] to study the performance of LTE base station receivers that experience radar interference. Appendix A of [25] evaluates the performance of both downlink LTE and Wi-Fi with radar interference using hardware test beds for each.

An opportunistic spectrum sharing model for radar-cellular coexistence is investigated in [27], providing performance analysis of cellular devices avoiding transmissions that would interfere with radar above a tolerable level. In [16], downlink LTE communication performance amid radar interference is analyzed using link simulations of the same nature as the ones performed in this thesis. That work also proposes a mitigation scheme involving soft decoding at the LTE receiver. [24] explores the prospect of bit error pattern identification to help Wi-Fi nodes detect interfering radar pulses.

To the author's knowledge, at this time no published literature provides link and system level software simulations of radar-Wi-Fi coexistence. The simulators in this work are novel because of their high configurability, scalability, and repeatability that allow users to easily study performance of Wi-Fi networks sharing spectrum with radar systems in varying scenarios and configurations. This helps foster insights on the performance impact of different system parameters.

1.3 Thesis Overview

The remainder of this thesis is organized as follows. Chapter 2 is a background summary of key elements of the Wi-Fi MAC and PHY layers as well as general radar system characteristics and operation. Comprehension of this material is essential to understanding subsequent chapters. Chapter 3 covers the link sim by describing its major features.

An analysis of the link sim results is then presented. The system simulation is detailed in Chapter 4, with an overview of the ns-3 Wi-Fi module, an explanation of the additions and modifications to it, and an analysis of the system simulation results. Mitigation schemes are proposed and discussed in Chapter 5. Lastly, Chapter 6 consists of research conclusions and potential future work on this topic.

Chapter 2
BACKGROUND

2.1 Wi-Fi

Wi-Fi is the WLAN technology adhering to the Institute of Electrical and Electronics Engineers' (IEEE) 802.11 standards. A WLAN corresponding to any of the 802.11 standards and amendments is considered Wi-Fi. In short, Wi-Fi technology comprises MAC and PHY layer specifications to be used in WLAN operation. Many amendments to the base 802.11 standard exist, but most share the same fundamental principles. Some of the key Wi-Fi concepts and provisions are discussed in this section with an emphasis on IEEE 802.11n, which is currently very popular. Everything discussed in this section applies to 802.11n, the standard being simulated in this research, but most of the concepts apply to other versions as well.

2.1.1 OFDM

Orthogonal frequency division multiplexing (OFDM) is the most common PHY digital modulation technique used in modern Wi-Fi technology. Instead of the traditional method of modulating a single carrier wave to send information, OFDM uses multiple carriers with smaller bandwidths, or subcarriers, for parallel transfer. It is a subset of frequency division multiplexing (FDM), which is the technique of dividing the bandwidth of a communication channel into narrower sub-bands, each having its own subcarrier to be modulated with data.

FDM provides improved performance over single-carrier modulation (SCM) under frequency-selective fading conditions because each subcarrier occupies a smaller part of the channel bandwidth, making the FDM system less impacted by differences in signal attenuation for different frequencies. If each sub-band is smaller than the coherence bandwidth, or the

approximate bandwidth over which fading is nearly constant in the channel, then each subcarrier essentially experiences flat fading, which is much easier to account for in the receiver chain. FDM is also more robust against narrowband interference, because a narrowband interferer would only impede the subcarriers that it overlaps with in frequency, instead of the whole signal as it would in SCM [28]. Since FDM provides multiple parallel streams, it is afforded a lower data symbol (pulse representing one or more digital bits) rate per subcarrier to achieve the same bit rate as a single-carrier system[1], which would require a higher data symbol rate on its one carrier. A lower symbol rate means a greater symbol length, which yields improved robustness against inter-symbol interference (ISI). ISI is typically caused by multipath propagation, the phenomenon of a single wireless signal taking multiple paths (direct or reflective) and reaching the receiver at different times, because it causes time-delayed copies of a symbol to appear after the initial one, leading to the delayed copies interfering with subsequent symbols. With a longer symbol length, the delayed symbols interfere with a smaller percentage of the subsequent symbol periods, resulting in less performance degradation.

Clearly, FDM offers many advantages over SCM, but it also has the major disadvantage of reduced spectral efficiency. To avoid inter-carrier interference (ICI), unused spaces in the frequency spectrum between FDM subcarriers, called guard bands, must be allocated. Guard bands help prevent subcarriers from interfering with each other, but also create gaps in spectrum, making the channel underutilized. OFDM solves this spectral efficiency problem by allowing subcarriers to have overlapping spectra. By making the subcarriers orthogonal to each other, OFDM receivers are able to distinguish between each subcarrier's signal, despite their frequency overlap. This eliminates the need for inter-carrier guard bands (though guard bands are often still used between channels), which gives OFDM extremely efficient spectrum utilization capability. A frequency spectrum of orthogonal subcarriers is shown in Figure 2.1. One of the key characteristics of orthogonality is that the peak of each subcarrier's

[1]It is assumed that this SCM system has the same constellation mapping and channel bandwidth as the compared FDM system.

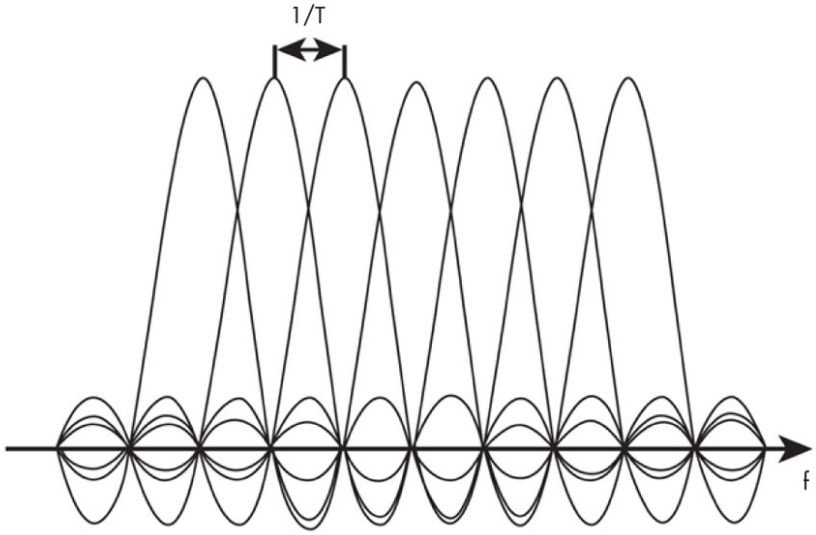

Figure 2.1: OFDM subcarrier spectrum. Retrieved from [3]

spectrum aligns with the nulls of all the other subcarriers' spectra. In essence, OFDM receivers calculate the spectrum values of received signals at these peaks, and therefore demodulation of each subcarrier's signal can be accomplished without ICI [5]. In the figure, each subcarrier's spectrum corresponds to a mathematical sinc function in frequency, which is the ideal case for OFDM systems. It can be seen that the peak of each sinc occurs at the same frequency as the zero crossings of all the other ones. Another essential OFDM characteristic is observable in the figure: the orthogonality requirement that the subcarrier frequencies are separated by 1/T, where T is the subcarrier symbol period.

Waveform Generation

A complex baseband time domain OFDM signal with a spectrum as in Figure 2.1 can be represented (before pulse shaping) as:

$$x(t) = \sum_{k=0}^{N-1} X_k \cdot e^{j2\pi k \Delta_f t} \tag{2.1}$$

where N is the number of subcarriers, X_k contains the I/Q modulation coefficients for each subcarrier from the bits-to-constellation symbol mapping, and Δ_f is the subcarrier frequency spacing, which is equivalent to the subcarrier symbol rate R_s (expressed as 1/T above). An amplitude scaling factor would also typically be included in practice. Eq. (2.1) intuitively holds considering that it describes a signal consisting of the sum of multiple equally spaced waveforms with the correct frequency spacing to achieve orthogonality.

Before the digital data to be transmitted in an OFDM system is parallelized, it begins as bits transferred serially through a constellation mapper, which maps bits to symbols specified by the constellation modulation scheme. These input symbols are transferred with a symbol rate R. Once the data is parallelized, the symbol rate (of each subcarrier) decreases, becoming $R_s = R/N$. As the digital domain is being discussed, define the discrete time $t = n\Delta t$, with n ranging from 0 to N - 1 in steps of 1. This ranges in time from 0 to subcarrier symbol period T. $\Delta t = 1/f_s$, where f_s is the sampling frequency. Thus:

$$T = N\Delta t = N \cdot \frac{1}{f_s} \tag{2.2}$$

Because $R_s = R/N$, the subcarrier symbol period T is equal to N times the input symbol period. Comparing this relationship with Eq. (2.2), it follows that f_s is equivalent to R, as 1/R equals the input symbol period. From here:

$$f_s = R = R_s \cdot N = \Delta_f \cdot N \tag{2.3}$$

since $R_s = \Delta_f$. Now:

$$\Delta t = \frac{1}{f_s} = \frac{1}{\Delta_f \cdot N} \tag{2.4}$$

Eq. (2.4) can be substituted into Eq. (2.1) with continuous time t being replaced with discrete time $n\Delta t$, simplifying to:

$$x_n = x(n\Delta t) = \sum_{k=0}^{N-1} X_k \cdot e^{j2\pi kn/N} \tag{2.5}$$

Eq. (2.5) is the inverse discrete Fourier transform (IDFT) without the 1/N amplitude scaling factor that the IDFT contains. Consequently, it is obvious that the IDFT can be used to

generate a time domain signal of orthogonal subcarriers for OFDM. The inverse fast Fourier transform (IFFT) is generally used to compute the IDFT, as it is an efficiently implemented equivalent. The number of IFFT points corresponds to the number of OFDM subcarriers generated. The 802.11n standard specifies 64 subcarriers for 20 MHz channels [17], so a 64-point IFFT could be used to create the baseband OFDM signal. This makes subcarrier spacing $\Delta_{\mathrm{f}} = \frac{20\ \mathrm{MHz}}{64} = 312.5\mathrm{kHz}$. Note that the above mathematical analysis containing Eqs. (2.2) to (2.5) is abbreviated from [28].

Comparison to Single-carrier Modulation

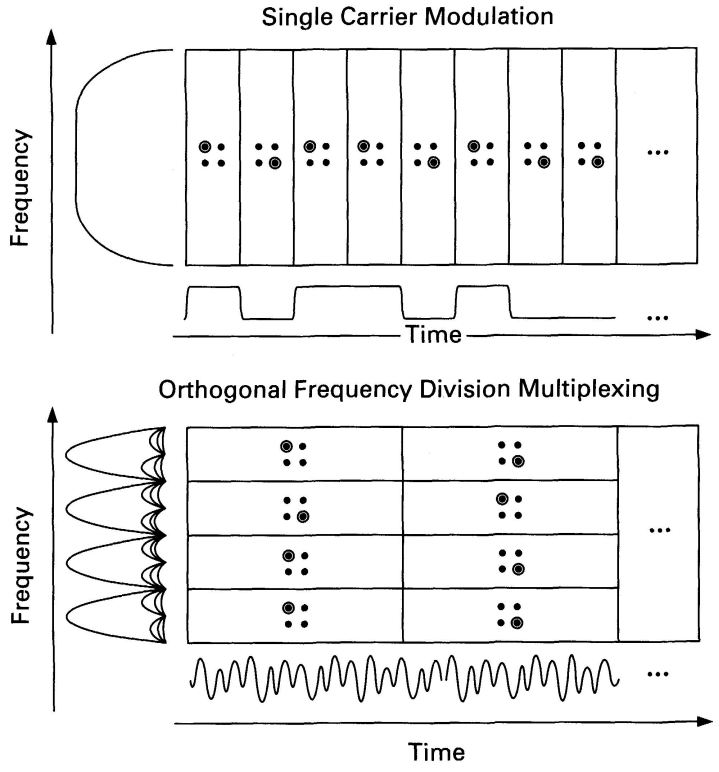

Figure 2.2: Single carrier modulation vs. OFDM. Retrieved from [4]

An example comparison between SCM and OFDM in the time and frequency domains is displayed in Figure 2.2. It also shows the constellation points of the data symbols being

selected for each. In this example, both the SCM and OFDM systems are using quadrature phase shift keying (QPSK) for the bits-to-symbol mapping. Since the SCM system has a single carrier, it only sends one QPSK symbol at a time. On the other hand, the OFDM system is using four subcarriers, each of which carries one QPSK symbol simultaneously. The OFDM waveform is the sum of the four subcarriers' signals, which explains why it appears to be more complicated then the simple SCM waveform. But in actuality, both systems are transmitting the exact same digital data, as the QPSK symbols being sent are exactly the same. Let us define the term *data symbol* to denote constellation-mapped symbols, in this example being the QPSK symbols. These may also be referred to as subcarrier symbols in this paper. We now define the term *OFDM symbol* to be the set of data symbols transmitted simultaneously in parallel by an OFDM transmitter. The OFDM symbol is the actual signal that is transmitted across the wireless channel. In this example, each OFDM symbol consists of four QPSK data symbols. In an SCM system, the signal transmitted through the channel consists of a single data symbol. Since it is only sending one data symbol at a time, to transfer the same amount of data, the SCM symbol period must be shorter than the OFDM symbol period (if both systems have identical channel bandwidth and constellation mapping), proportional to the number of subcarriers in the OFDM system. This explains why the OFDM symbol period is four times longer than the SCM symbol period in the figure.

Alternatively, an SCM system could maintain its data rate and keep a longer symbol duration by using a higher order constellation for its modulation, but this makes the transmitted symbols more vulnerable to noise and interference. This highlights another advantage of OFDM—its tendency to use lower-order constellations for its subcarriers than comparable SCM, yielding more resilient data symbols.

Guard Interval and Cyclic Prefix

To protect from ISI, guard time is added at the beginning of each OFDM symbol. This extra symbol time is not used for data, but is useful in reducing the negative effect of multipath signal interference. The data from each received symbol is meant to be sampled

after the guard time, so if any delayed copies of the preceding symbol arrive during the guard interval, effectively no ISI occurs. Implemented OFDM systems typically have guard times longer than the expected delay spread of the channel, making ISI highly unlikely [5]. When solely considering ISI reduction, it is plausible to transmit nothing during the guard interval. However, that could lead to ICI between subcarriers, because in some circumstances the received subcarrier signals would no longer be orthogonal. For example, if one subcarrier signal arrives later than others amid reception of an OFDM symbol, the empty guard time of the delayed subcarrier symbol could appear during the demodulation period of the OFDM symbol, causing the subcarrier signals to lose their orthogonality within the receive window.

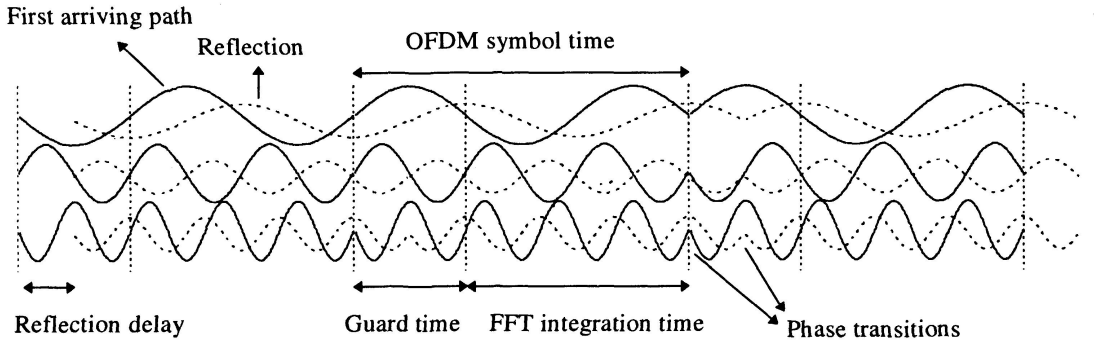

Figure 2.3: "Example of an OFDM signal with three subcarriers in a two-ray multipath channel. The dashed line represents a delayed multipath component." Retrieved from [5]

In order to avoid ICI, a cyclic prefix is added to the beginning of each OFDM symbol, covering the guard time. This cyclic prefix is a copy of the end segment of the OFDM symbol waveform, causing the signal to remain periodic and maintain orthogonality between subcarrier signals. Since the cyclic prefix contains part of the modulated data symbol, it also allows the data to be successfully sampled during any part of the symbol. Figure 2.3 illustrates a three subcarrier OFDM signal with each subcarrier signal depicted separately (instead of showing the actual OFDM signal which is the sum of the three) for reader clarity. Each data symbol is modulated with binary phase shift keying (BPSK), meaning that the

edges between symbols correspond to 180-degree phase changes. The dotted-line signals are delayed copies of the solid-line signals, which arrive at the receiver first. The merit of guard time can be seen in the figure. As the time delay of the duplicate signals is shorter than the guard time, their symbol transitions do not appear during the receive sampling time of the OFDM symbols, denoted in the figure as "FFT integration time" since a fast Fourier transform (FFT) is performed at the receiver. Thus, data corruption from ISI is prevented, and the subcarrier signals remain orthogonal during demodulation time. If a symbol transition from one of the delayed subcarrier signals did fall into the receive sample window, orthogonality would be destroyed [5].

Block Diagrams

Simplified block diagrams of an OFDM transmitter and receiver are shown in Figure 2.4. For the transmitter, shown on top, bits are mapped to data symbols with the modulator (based on the constellation chosen) and then converted from series to parallel. The IFFT generates time domain samples of the baseband OFDM modulated waveform, which are then converted from parallel to series and have the cyclic prefix added to assemble the full waveform of the OFDM symbol. Finally, the signal is converted from the digital domain to analog and sent to the upconverter, which brings its frequencies up to the passband for wireless transmission.

The OFDM receiver, displayed at the bottom of Figure 2.4, begins by downconverting the received wireless signal back to baseband and converting it back to the digital domain. It then has its cyclic prefix removed, is converted to parallel, and an FFT is performed on it. The FFT effectively reverses what the IFFT in the transmitter did, and so converts the time domain OFDM symbol back to its separate data symbols. Lastly, the parallel data symbols are serialized and sent to the rest of the receiver for demapping and decoding.

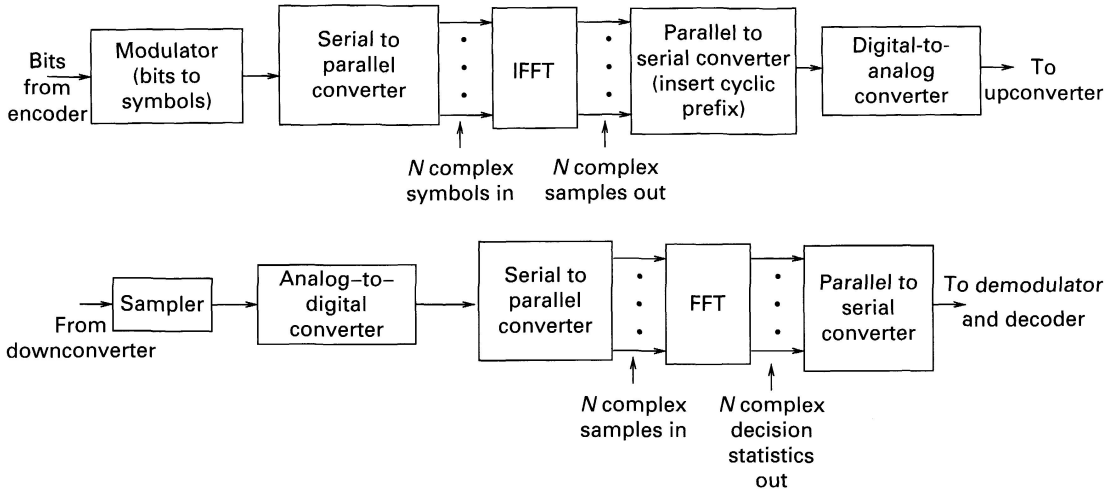

Figure 2.4: Simplified OFDM transmitter (top) and receiver (bottom) block diagrams. Adapted from [6]

End Remarks

OFDM systems do have some disadvantages compared to single-carrier systems. Because OFDM waveforms are the sum of multiple subcarrier waveforms, they often have a high peak-to-average power ratio. Power amplifiers in wireless transmitters normally have a limited linear power range, and distortion occurs if a signal exceeds it. This can lead to transmit power needing to be reduced or transmitters requiring more advanced power amplifiers [4]. OFDM is also more sensitive to carrier frequency offset and phase noise, phenomena arising from imperfect oscillators used in upconversion and downconversion, which disrupt the orthogonality of OFDM signals. Nevertheless, with all the advantages that OFDM provides, especially in noisy, frequency-selective channels, it is easy to see why it is such a popular choice in modern wireless communication systems.

2.1.2 MCS

IEEE 802.11n and the newer 802.11ac provide a set of standard PHY settings called the Modulation and Coding Scheme (MCS). Each setting is given a numbered MCS index, which corresponds to a specific modulation (constellation mapping) type and forward error correction code rate. The modulation type determines how many bits are transmitted per data symbol, with higher order modulations having more bits per symbol. More bits per symbol yields a higher data rate, but also decreases immunity to noise and interference, because constellation points[2] are closer together, making it more likely for the decoder to interpret received symbols incorrectly.

Forward error correction (FEC) coding is a method of reducing the probability of received bit error by including redundant bits within the data to be communicated. These redundant bits do not include any additional information to be transferred, but increase the robustness of the system to wireless channel effects. The code rate is the fraction of transmitted bits that are non-redundant, or that in other words provide new data to be transferred. It is typically represented as k/n, where for every k bits of new data, n - k redundant bits are sent, making n the total number of transmitted bits. For example, with a code rate of 3/4, for every three information bits, one redundant bit is included. A lower code rate increases the likelihood of successful bit decisions for received data but decreases the data rate, while a higher code rate does the opposite.

Table 2.1 contains the 802.11n MCS parameters and High Throughput (HT) data rates for 20 MHz channels with one spatial stream, meaning single input single output (SISO) antennas. As the MCS index increases, the data rate of the system increases, as well as the vulnerability to negative channel effects. This is accomplished with a combination of increasing modulation order and code rate. From MCS 0 to MCS 7, modulation increases

[2]In a digital communication system, constellation points are the set of amplitude and phase (I/Q) coordinates possible for a data symbol given a particular modulation type. Transmitters select the correct constellation point for each symbol-sized bit cluster to be sent, and receivers decode each symbol by selecting the constellation point that best matches the received symbol.

Table 2.1: IEEE 802.11n MCS parameters and HT data rates for 20 MHz channel with one spatial stream [17].

MCS index	**Modulation**	**Code Rate**	**Data Rate (Mbps)**	
			GI = 800 ns	GI = 400 ns
0	BPSK	1/2	6.5	7.2
1	QPSK	1/2	13.0	14.4
2	QPSK	3/4	19.5	21.7
3	16-QAM	1/2	26.0	28.9
4	16-QAM	3/4	39.0	43.3
5	64-QAM	2/3	52.0	57.8
6	64-QAM	3/4	58.5	65.0
7	64-QAM	5/6	65.0	72.2

from BPSK to QPSK to quadrature amplitude modulation (QAM). The code rate increases between MCSs in which the modulation scheme remains the same. The table provides two data rates for each MCS, one for each of the two guard interval settings. The short guard interval setting of 400 ns provides a higher data rate because more of the OFDM symbol is used for data transfer instead of guard time.

Implemented systems frequently use link adaptation to adjust the selected MCS in real time. Link adaptation is the practice of adjusting the modulation scheme and code rate over time in accordance with the wireless link quality (affected by noise, interference, channel fading, etc.) in attempt to achieve the best possible performance through the channel. For example, a high MCS index would not be very effective in a link with a low signal-to-noise ratio (SNR), so a link adaptation algorithm would select a lower MCS to improve system throughput. On the other hand, if the link had a high SNR and was set to a low MCS, link adaptation would set the MCS higher, because higher order modulation and code rates (and thus data rates) are more successful with higher SNR.

Many more PHY options exist besides the ones shown above. Data rates for each MCS index are higher when using 40 MHz channels. Higher MCS indices (and thus data rates) are also included in 802.11n but correspond to more spatial streams, requiring MIMO. 802.11ac has the same MCS 0 - 7 settings as 802.11n, but also provides two higher MCS settings for SISO, which use 256-QAM.

2.1.3 PHY Frame Structure

Digital data is communicated in Wi-Fi (and through the Internet) via packets, which are clusters of data. Not all of the information in a packet is part of the actual data that users intend to send or receive. Control data is also necessary in each packet for many networking and communication functions to be successful, such as network addresses that are required to determine where the packet will be sent to. Control data is essential, but is a form of network overhead because it utilizes the communication media without directly contributing to system throughput. Each network layer (application, transport, network, MAC, and PHY) contributes its own set of control data to each packet. At the MAC and PHY layers, packets are referred to as frames. This subsection focuses on PHY layer frames in 802.11n. As the PHY layer is closest to hardware, PHY frames contain all the control and user data from the other layers and consist of all the data that is physically transferred.

Figure 2.5 displays the three types of PHY frame structures in 802.11n: Legacy, Mixed Mode, and Greenfield. These are the three formats defined for 802.11n's Physical Layer Convergence Protocol (PLCP). In each structure type, the blocks preceding the labeled "Data" block will be referred to as the preamble. The "Data" block (hereinafter *payload*) comprises user data and control data from the higher network layers. It also includes a service field for data descrambling[3], tail bits (all zeros) for terminating the forward error correction decoder, and extra pad bits (all zeros) if necessary to ensure that enough bits are present to fill up the last symbol. Time durations are shown above each preamble segment,

[3]Bit scrambling is performed in the transmitter to randomize data and avoid long chains of zeros or ones, so the data must be descrambled at the receiver.

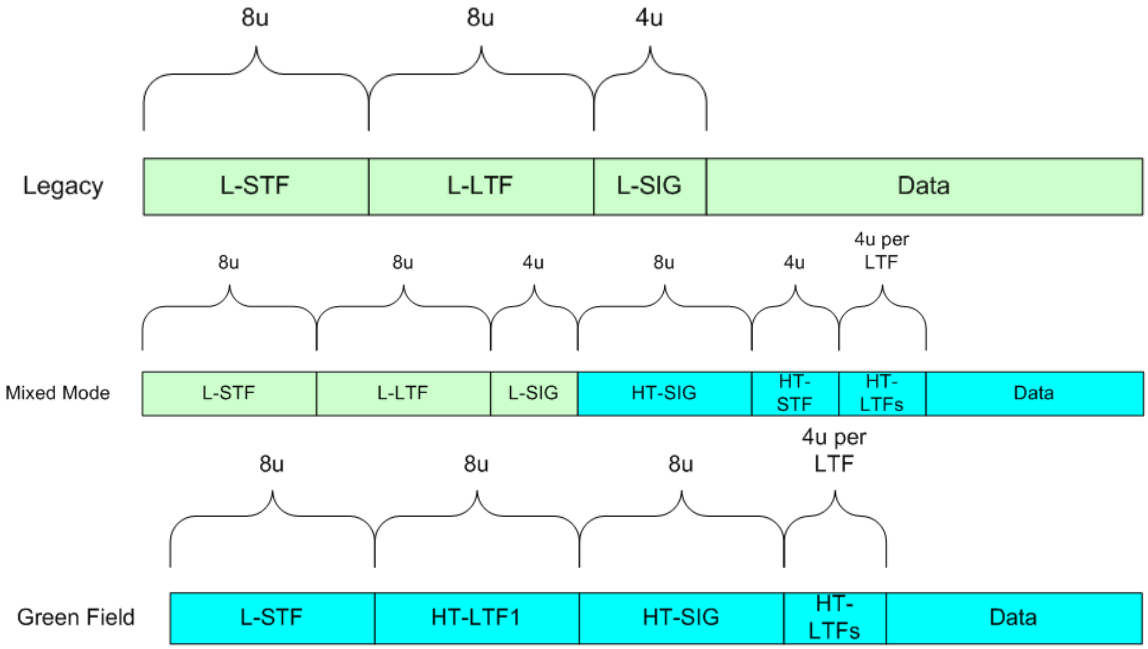

Figure 2.5: IEEE 802.11n PHY frame structures. Adapted from [8]

while payload duration is variable because it depends on user and control data from upper layers.

Legacy mode frames are identical to the OFDM frame format of the older 802.11a and 802.11g standards. As such, 802.11g and 802.11a devices can decode legacy mode frames, but lower legacy data rates are used. The preamble of legacy frames consists of the Legacy Short Training Field (L-STF), the Legacy Long Training Field (L-LTF), and the Legacy Signal field (L-SIG). The first preamble block, the L-STF, contains known subcarrier symbols that are used for frame finding, time synchronization, frequency offset measurement, and automatic gain control[4] (AGC) configuration. The next field, the L-LTF, contains known

[4]Automatic gain control (AGC) is an important function in wireless receivers to automatically adjust the receiver amplifier gain based on the amplitude of incoming signals. The goal is for the AGC to set the gain so that received signals are at a desirable amplitude range before being converted to the digital domain for processing and decoding.

subcarrier symbols that are used for channel estimation[5] and more precise time synchronization and frequency offset measurement [4]. The last legacy preamble field, the L-SIG, carries modulation, code rate, and frame length information for decoding the frame.

Mixed Mode frames also provide compatibility with legacy devices, but additionally include support for MIMO and HT data rates (802.11n HT SISO rates shown in Table 2.1). The first part of the mixed mode preamble is the same as an entire Legacy mode preamble, as described above. This allows the frame to maintain interoperability with legacy devices. The remaining portions of the mixed mode preamble are part of the HT preamble. Similar to the L-SIG in Legacy mode, the HT Signal field (HT-SIG) provides HT rate and frame duration information, and is also used by receivers to distinguish between Legacy mode and Mixed Mode frames [4]. The HT Short Training Field (HT-STF) is used to reestablish the AGC configuration with more accurate settings for MIMO. Lastly, the HT Long Training Fields (HT-LTFs) are for MIMO channel estimation with multiple spatial streams.

Greenfield mode frames are only compatible with 802.11n devices, not with older legacy devices. They provide the same MIMO capabilities and HT data rates that mixed mode frames do, but with less overhead and consequently higher throughput potential, since the legacy preamble is not included. The greenfield preamble has similar blocks and functions as the HT portion of the mixed mode preamble.

2.1.4 DCF

One of the core responsibilities of the MAC layer is to regulate which device has access to use the communications medium at any given time. In the case of Wi-Fi, the communications medium is the wireless channel. The primary MAC method to regulate channel access in 802.11 is the Distributed Coordination Function (DCF). Since all devices in a single Wi-Fi network use the same wireless channel, interference occurs if multiple devices transmit

[5]Channel estimation is a common technique performed in wireless receivers to measure the amplitude and phase fading effects of the channel over frequency and/or time. The channel estimate is generally used by an equalizer, which attempts to reverse the effects of the channel on received symbols to increase the likelihood of successful data decoding.

frames at the same time, which is called a collision. This leads to difficulties in successfully receiving frames. DCF provides mechanisms to help avoid these situations and promote efficient channel utilization. A flowchart outlining DCF operation is depicted in Figure 2.6.

Carrier Sensing

The basis of DCF is the Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) protocol. In the carrier sense function of CSMA/CA, Wi-Fi nodes (devices) perform sensing to determine if the channel is free, i.e., no other nodes are transmitting. To do so, they use a combination of *physical carrier sensing* and *virtual carrier sensing*. For the channel to be classified as idle, both physical and virtual carrier sensing must determine it as such. Frame transmission can only begin after both mechanisms specify the channel to be idle.

Physical carrier sensing is the determination of wireless channel status from physical inspection. One way it is performed is by listening for Wi-Fi frames being sent by other nodes. A node regards the channel as busy and will not transmit its own frames if it is detecting a frame in the channel from another node. Nodes can also take non-Wi-Fi sources into account by measuring received energy from the channel and comparing it to a threshold. If the threshold is exceeded, the node identifies the channel as busy and will not transmit. Both of these techniques are part of clear channel assessment (CCA), which uses physical carrier sensing for a specific period of time, known as the DCF Interframe Space (DIFS).

Virtual carrier sensing is another mechanism for nodes to ascertain the channel's status. The MAC header of each transmitted frame contains a field that specifies the time duration for which the channel is expected to be busy for that frame transmission and acknowledgment sequence, including frame fragments[6]. Any node that picks up a frame in the channel from any other node, even if the frame's data is not intended for it, can read the MAC header to determine this time duration. The receiving node then sets its Network Allocation Vector

[6]Frames can optionally be broken up into smaller frame fragments before being transmitted in the PHY layer. This can be useful for large data frames because, since they spend a long time in the channel, large frames are more vulnerable to changing channel conditions and collisions with other frames. Frames are compared to a fragmentation threshold to determine whether or not they become fragmented.

(NAV) to this time duration. Once it is set, the NAV begins to count down to zero, after which the frame transmission sequence should be finished.

Collision Avoidance and Acknowledgments

The collision avoidance part of CSMA/CA is executed in Wi-Fi with the random backoff procedure. Random backoff reduces the probability that nodes will begin transmitting at the same time after the channel is sensed to be idle, hence reducing the chance of collision. If a node has a new frame to transmit but senses the channel to be busy during CCA, it chooses a random number uniformly between zero and the Contention Window (CW), inclusive. The CW begins at implemented value CWmin. The chosen random number, the backoff count, corresponds to the number of slot times that the node has to wait for during idle channel periods before it can transmit. For 802.11a, 802.11g, 802.11n, and 802.11ac, the slot time is 9 μs [4]. The backoff counter only counts down in periods where the channel is sensed to be idle, meaning the NAV is zero and CCA has sensed the channel to be clear for a DIFS duration. If the channel is sensed busy (via NAV or CCA), the random backoff counter holds its value and does not resume counting down until the idle channel conditions are met again.

Once a node receives a standard data frame[7] or fragment, it waits a Short Interframe Space (SIFS) period then sends an acknowledgment (ACK) frame back to the sender node to let it know that the frame was received correctly by the destination node. If the sender node receives the ACK within the ACK timeout period, it waits a SIFS period before transmitting the next frame fragment, if any, and again waits for an ACK. This repeats until all frame fragments are successfully sent and acknowledged, after which CW is reset to CWmin and DCF starts over. The SIFS period is always shorter than the DIFS period to protect from other nodes interrupting the current node's frame transmission and ACK sequence, since nodes have to pass CCA for a DIFS interval before beginning frame transmission. In general,

[7]A data frame is a frame containing data that is the goal of the user to transfer. In contrast, control frames include various frame types that are used for different MAC functions to assist in successful communication of data frames, such as acknowledgments (ACKs).

SIFS is used to prioritize specific types of frame exchanges over others. It is chosen to be as short as possible but long enough to account for PHY latency and MAC level processing. The SIFS duration is 16 μs for 802.11a, 802.11g, 802.11n, and 802.11ac [4]. The DIFS duration is defined as follows:

$$DIFS = SIFS + 2 \cdot SlotTime \tag{2.6}$$

If, after transmitting a data frame or fragment, the sender node does not receive an ACK within the ACK timeout period, it regards the event as a "virtual collision". This node does not know if its frame was received incorrectly or if only the ACK was received incorrectly. In either case, it is likely that at some point, a collision occurred in the channel, which is more prone to occur in a WLAN with heavy traffic or many nodes. To reduce the chance of collision occurring again, the sender's CW value is doubled (unless it is at max value CWmax, which it will then remain at) and the random backoff procedure is run before retransmission is attempted. To understand the benefit of this, consider a network with nine nodes. If one node is transmitting and the other eight have new frames to transmit but sense the channel as busy during CCA, each of them draws a random backoff number. If the CW is set to a low number, say eight, it is more likely that more than one of the nodes draws the same number than if the CW was higher, say sixteen. If two or more nodes select the same random backoff number, they will begin transmitting at the same time, causing a collision. This is why doubling the CW reduces the probability of collision events. But the CW starts low because a high CW leads to longer backoff delay times (because of potential for higher-valued random draws), which can cause extra unnecessary idle channel time and thus decreased throughput in small networks.

RTS/CTS and Retry Counters

As described above, DCF provides efficient and fair channel access regulation. However, the above analysis assumed that each node in a WLAN is able to sense the presence of all the other nodes by detecting their frames. But in some situations, nodes are not able to do

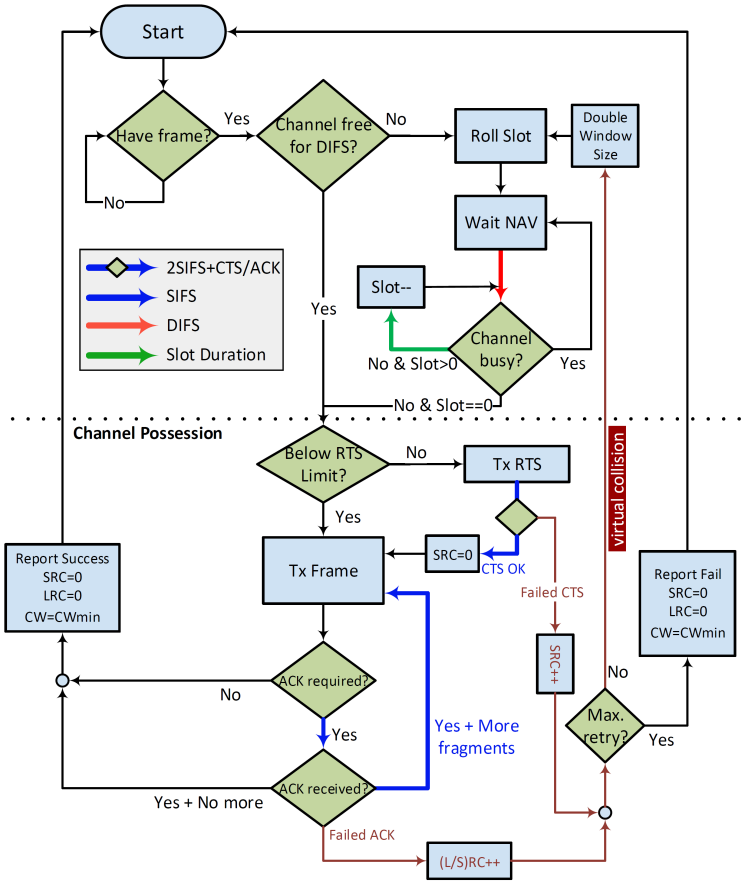

Figure 2.6: Distributed Coordination Function (DCF) flowchart. Retrieved from [9]

so, often due to distance. For example, this occurs when one node is able to receive frames from two other nodes but those two nodes are unable to detect each other. When one of these two nodes transmits, the other will not detect the channel as busy during CCA and will not receive the frame to set its NAV. Because of this, it may also begin transmitting before the first node is done with its transmission, causing a collision. This type of scenario is called the *hidden node problem*.

802.11 has the RTS/CTS option to address the hidden node problem. If this option is used, before sending a data frame, a node sends a short Request to Send (RTS) frame. After receiving the RTS frame, the receiving node waits a SIFS period then sends out a Clear to

Send (CTS) frame. If the sending node receives the CTS frame within the CTS timeout period, it begins the data frame transmission and ACK process, with each transmission separated by a SIFS duration. Just as with ACK exchanges, the short SIFS period gives RTS-CTS-data frame exchanges priority channel access. All nodes in range of the receiving node can receive its CTS frame—even nodes that are out of range of the sending node. Thus, the RTS/CTS process can let hidden nodes know about the sending node's transmission, helping reduce the chance of collision. If the receiving node is an access point, which is described in Subsection 2.1.5, all nodes in the network are in range and can receive the CTS. Both RTS and CTS frames contain, in their MAC header duration fields, the length of time that the RTS, CTS, data frame, and ACK exchange will take. This causes nodes in range of the sending node as well as potentially hidden nodes in range of the receiving node to set their NAV accordingly [29].

Since RTS frames are typically shorter than data frames, they are less likely to collide with other frames. Because of this, RTS/CTS can also be useful in non-hidden node scenarios that have a noisy or heavily loaded channel. A short RTS frame will have a higher chance of successful reception than a long data frame, and can be used to grant a node control of the channel so its data frames can be successfully received. When RTS/CTS is activated, an RTS threshold must be set. Data frames larger than the threshold utilize RTS and CTS frames, while smaller data frames just transmit normally. The RTS/CTS function can reduce throughput in situations that are not appropriate for it, since it uses extra channel time for control frames instead of data frames. Consequently, care should be taken when activating RTS/CTS and choosing the RTS threshold.

When data frame transmission fails (no ACK received), retransmission of the same frame is attempted after the random backoff period. Wi-Fi nodes keep track of the number of retries with the short retry count (SRC) and long retry count (LRC), which both start at zero. The SRC is incremented when transmission fails for frames smaller than the RTS threshold, and the LRC is incremented when transmission fails for frames larger than the RTS threshold [30]. If either the SRC or LRC reach their limit, the data frame is discarded

and no more retransmissions are attempted, while both SRC and LRC are reset back to zero and CW is set back to CWmin. SRC, LRC, and CW are also reset when a data frame and all its fragments are successfully ACKed.

2.1.5 WLAN Architecture

Wi-Fi networks have two possible topologies: ad-hoc mode and infrastructure mode. Ad-hoc mode is a "peer-to-peer" configuration where Wi-Fi devices communicate directly with each other. In the more common infrastructure mode, instead of Wi-Fi devices communicating directly with each other, they communicate through an access point (AP). Wireless routers are normally used as APs in home and office WLANs. Each user device in infrastructure mode is called a station (STA), or client. When a STA wants to send a packet to another STA in the WLAN, it sends the packet to the AP, which then sends it to the destination STA.

Ad-hoc networks are simple to set up and are useful for quick and easy data transfer between devices. They are also inexpensive, as no AP hardware is required. However, infrastructure mode has many advantages. It generally provides a wider range, since APs typically have stronger transmitters than STAs, and higher security. Infrastructure mode also allows wireless extension of wired networks via wired connection to an AP. In the common case of a wireless router serving as the AP, the Internet is connected to the AP via a wired connection and the AP can then share its Internet connection with all the STAs connected to it. In this scenario, all incoming and outgoing Internet traffic goes through the AP.

2.2 Radar

Radar (originally an acronym for Radio Detection and Ranging) is a wireless technology that uses electromagnetic (EM) waves to detect objects. Its applications include weather forecasting, aircraft and ship navigation, air traffic control, defense systems, surveillance, and many more. Depending on the type of system used, radar can determine the distance, direction, altitude, and/or velocity of objects in its path.

The basic principle of radar is that by transmitting an EM wave pulse in a direction and receiving it back at the source after it is reflected off of a potential object, properties of the object can be ascertained. For example, since it is known that EM pulses travel at the speed of light, radar systems can accurately calculate the distance to an object by measuring the amount of time it takes for the transmitted pulse to return after being reflected.

Since radar is solely treated as an interferer in this research and its simulations, it is covered in less detail in this chapter than Wi-Fi.

2.2.1 Basic Measurements

In this subsection, three basic radar measurements used to resolve characteristics of detected objects are discussed: time delay, Doppler shift, and directionality.

Time Delay

By measuring the time delay for a transmitted pulse to return to the radar source after reflecting off of an object, radar systems can calculate the distance to the object. Given time delay τ, the distance d can be calculated as [31]:

$$d = \frac{1}{2}\tau c \tag{2.7}$$

where c is the speed of radar signal propagation, which we approximate as the speed of light in a vacuum throughout the remainder of this section. In reality, the exact speed of propagation depends on the medium that the pulse travels through. The $\frac{1}{2}$ factor is present because the pulse travels distance d twice before being received—once to get to the object and another time to return to the radar source. Eq. (2.7) also assumes that the radar pulse propagates in a straight line.

Doppler Shift

A class of radar systems called Doppler radar can exploit the Doppler effect to determine the velocity of detected objects. The Doppler effect is the phenomenon where the frequency of

a wave emitted from a moving source is observed as being shifted from its emitted frequency. The same effect occurs when a wave reflects off of a moving object. Thus, after a radar pulse reflects off of a moving object, the radar receiver receives it at the Doppler-shifted frequency. The (shifted) frequency f_r of a received radar pulse is:

$$f_r = f_t \left(\frac{1 + \frac{v_o}{c}}{1 - \frac{v_o}{c}} \right) \tag{2.8}$$

where v_o is the radial velocity[8] of the detected object and f_t is the emitted frequency of the transmitted radar pulse. The Doppler frequency f_d, the frequency shift between the transmitted and received signals, is then:

$$f_d = f_r - f_t = f_t \left(\frac{1 + \frac{v_o}{c}}{1 - \frac{v_o}{c}} \right) - f_t = f_t \left(\frac{c + v_o}{c - v_o} - 1 \right) = 2 f_t \frac{v_o}{c - v_o} \tag{2.9}$$

after substituting in Eq. (2.8) for f_r. Since v_o is much less than c for the majority of radar applications, $c - v_o$ is approximately c. Using this, Eq. (2.9) becomes:

$$f_d \approx 2 f_t \frac{v_o}{c} \tag{2.10}$$

Consider the voltage signal of a transmitted radar pulse, where A_t, f_t, and θ_t are its amplitude, frequency, and phase shift respectively:

$$v_t(t) = A_t \cdot cos(2\pi f_t t + \theta_t) \tag{2.11}$$

Then, the voltage signal of the same pulse being received after time delay τ is:

$$v_r(t) = A_r \cdot cos(2\pi f_t (t + \tau) + \theta_r) = A_r \cdot cos \left(2\pi f_t \left(t + \frac{2d(t)}{c} \right) + \theta_r \right) \tag{2.12}$$

after substituting in Eq. (2.7) for τ, where A_r and θ_r are its amplitude and phase shift respectively. The instantaneous frequency of a sinusoid is calculated by taking the derivative of its argument, i.e. $\phi(t)$ for $\cos(\phi(t))$, with respect to time. From this, the frequency of the received radar pulse is [32]:

$$f_r = \frac{1}{2\pi} \frac{d}{dt} \phi(t) = \frac{1}{2\pi} \frac{d}{dt} \left(2\pi f_t \left(t + \frac{2d(t)}{c} \right) + \theta_r \right) = f_t + \frac{2 f_t}{c} \frac{d}{dt} d(t) = f_t + \frac{2 f_t v_o}{c} \tag{2.13}$$

[8]The radial velocity is the component of the target object's velocity that is pointed directly toward or away from the radar source antenna. If the object is traveling completely perpendicular to the radar beam, its radial velocity is zero and no Doppler effect will be observed at the radar receiver.

since $\frac{\mathrm{d}}{\mathrm{dt}}\mathrm{d(t)}$ equals the radial velocity $\mathrm{v_o}$ of the detected object. The derivative equates to zero if the object is stationary, making received frequency $\mathrm{f_r}$ equal to transmitted frequency $\mathrm{f_t}$, which makes sense since no Doppler shift occurs with a stationary object. The last term of Eq. (2.13) matches up with Eq. (2.10), showing that $\mathrm{f_d} = \mathrm{f_r} - \mathrm{f_t}$ indeed.

These relationships provide some insight into how Doppler radar systems can determine the velocity of objects. One possible implementation for doing so is using a frequency mixer to mix a reference signal at the transmit frequency with the received signal. Low-pass filtering the mixer output would yield a signal at the Doppler frequency (from one of the subtracted frequency components of the mixer output signal). This Doppler frequency signal could then be used to calculate the detected object's velocity.

Directionality

Radar systems generally use highly directional antennas that change in direction over time, usually by rotating. When a pulse is transmitted from a highly directional antenna, the vast majority of the EM energy is radiated in the direction of the antenna's main lobe. In a simple sense, radar systems can exploit this antenna directionality to determine the direction of detected objects relative to the radar source.

2.2.2 Parameters

Radar systems have numerous parameters that can be tailored to performance goals. The basic radar signal-to-noise ratio equation [33]:

$$SNR = \frac{P_t G_t G_r \lambda^2 \sigma}{(4\pi)^3 d^4 k T_0 F B} \tag{2.14}$$

includes many of them. This equation provides the SNR of a received radar pulse $\mathrm{P_r}$, assuming free space propagation loss and thermal noise, as a function of adjustable system parameters $\mathrm{P_t}$, $\mathrm{G_t}$, $\mathrm{G_r}$, λ, F, and B. It essentially determines radar system performance, as the ability of a radar system to receive returning pulses correctly depends on the SNR. $\mathrm{P_t}$ is the total peak transmit power, $\mathrm{G_t}$ is the transmit antenna gain, $\mathrm{G_r}$ is the receive antenna

gain, λ is the wavelength of the radar pulse carrier wave (which is inversely proportional to its frequency), F is the receiver noise factor, which is the measure of SNR degradation that the receiver causes, and B is the receiver bandwidth, which contributes linearly to the amount of noise seen at the receiver. The remaining variables are the distance to the target object d, its radar cross section σ, Boltzmann's constant k (1.38 x $10^{-23} \frac{\mathrm{W \cdot s}}{\mathrm{K}}$), and the noise reference temperature T_0 (290 K). The radar cross section is the area that a perfectly isotropic reflector would need to yield the same return power as the actual target object [31].

As previously addressed, radar operates by transmitting EM pulses and receiving them after they are reflected off of objects. In order to successfully receive a reflected pulse, the radar transceiver must not be transmitting when the pulse returns (except in continuous wave radar systems, which are not discussed in this thesis). This is because if it is attempting to transmit and receive at the same time, it will effectively interfere with itself since the receiver will pick up some of the signal being transmitted. For this reason, radar systems are configured to have a *listening period*, a time period after a pulse transmission and before the next one in which they do not transmit. Returning pulses can only be successfully received if they reach the receiver during the listening period. The listening period should be set long enough to allow for a pulse to reach and return from the maximum desired range of the radar system. Because of this, the listening period tends to be much longer than the *pulse duration*.

Additional quiet periods (non-transmission periods) may also be necessary depending on the type of radar system. These can include a "recovery time" for the transmitter to ramp down or a "dead time" for control mechanisms or to guard from interference effects due to late-arriving return pulses. The combination of all these quiet periods, including the listening period, for a given system is termed the *rest time*, or *rest period*. Now, the pulse repetition interval (PRI) can be defined as:

$$PRI = \tau_p + \tau_r \tag{2.15}$$

where τ_p is the pulse duration and τ_r is the rest time. Figure 2.7 displays a visual representation of PRI, pulse duration, and rest time. The pulse repetition frequency (PRF) is

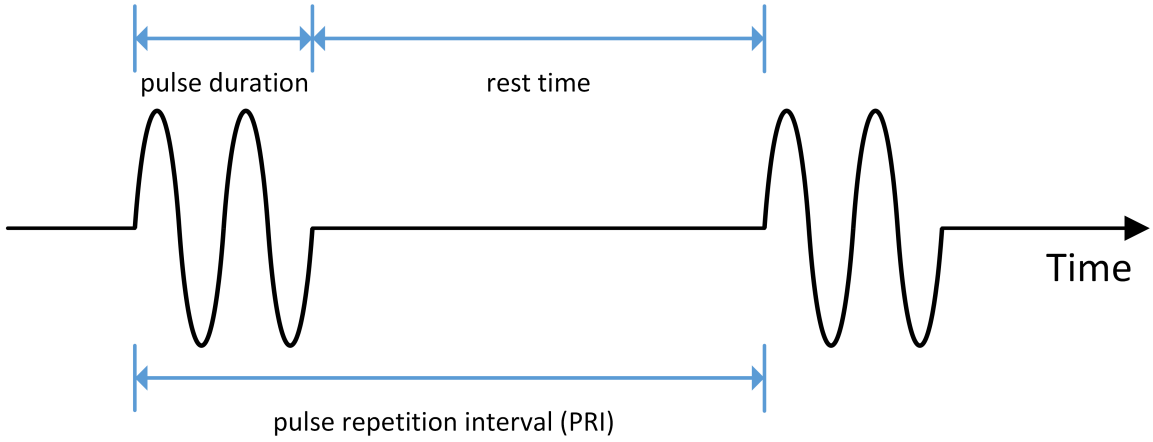

Figure 2.7: Radar timing periods.

then:

$$PRF = \frac{1}{PRI} = \frac{1}{\tau_p + \tau_r} \tag{2.16}$$

which is the number of pulses transmitted per second. Systems with high PRF have shorter ranges than those with low PRF (assuming the same pulse duration) because their listening periods are shorter. However, low PRF systems have more frequent pulse transmissions which tends to make object detection easier. These are trade-offs to take into consideration when choosing the PRF for a radar system.

The *scan rate* is the rate of antenna rotation for a rotating radar system, often measured in revolutions per minute or degrees per second. A higher scan rate should provide more frequent updates in each direction surrounding the radar source. However, with higher scan rate, fewer radar pulses will hit desired objects per rotation, making detection probability lower. Along with scan rate, a higher PRF would also increase the number of pulses that hit an object on each rotation. Therefore, systems with high PRF can typically afford a higher scan rate than those with low PRF [34].

A rotating radar can be implemented by rotating its antenna a fixed number of degrees before stopping to transmit a number of pulses (with rest periods in between adjacent transmissions). Then it rotates again, transmits, and repeats. For this discretized rotation, the

radar rotation period can be calculated as:

$$T_R = (\tau_p + \tau_r) \cdot \frac{360}{degrees\ per\ rotation} \cdot \#\ TX\ pulses\ per\ angle \tag{2.17}$$

The radar scan rate in revolutions per minute (rpm) is then simply:

$$scan\ rate\ (rpm) = \frac{60}{T_R} \tag{2.18}$$

2.2.3 Pulse Compression

The *range resolution* is the minimum difference in range (distance from the radar source) between two target objects at the same bearing (the same radar pulse hits both targets) at which a radar system can resolve the targets separately, and is given by:

$$\Delta R \geq \frac{c}{2B} \tag{2.19}$$

where B is the radar signal bandwidth. If two objects have a range separation smaller than ΔR, the radar system will detect them as one object. The bandwidth[9] of a *simple pulse*, a baseband rectangular-shaped pulse (before upconversion by the wireless carrier wave), is typically defined as $\mathrm{B} = 1/\tau_\mathrm{p}$, where τ_p is the pulse duration. In this case, Eq. (2.19) becomes:

$$\Delta R \geq \frac{c\tau_p}{2} \tag{2.20}$$

showing that range resolution performance is inversely proportional to pulse duration (since lower range resolution is better) in a radar system transmitting simple pulses.

In a radar system, received SNR is proportional to the transmitted energy per pulse. Since radar receiver bandwidth is matched to the transmit signal bandwidth, for a simple pulse system we can substitute in $\mathrm{B} = 1/\tau_\mathrm{p}$ into Eq. (2.14), yielding:

$$SNR = \frac{P_t \tau_p G_t G_r \lambda^2 \sigma}{(4\pi)^3 d^4 k T_0 F} \tag{2.21}$$

[9]The bandwidth of a signal pulse is inversely proportional to the pulse duration. This can be observed through the Fourier transform.

The energy of the simple pulse is the product of the peak transmit power and the pulse duration, or $E = P_t \tau_p$, which is part of the numerator of Eq. (2.21), exhibiting that increasing energy does increase the received SNR. Eq. (2.21) also points out that received SNR is proportional to pulse duration in a simple pulse radar system.

It has been shown that a simple pulse radar system's range resolution performance is inversely proportional to its pulse duration, while its received SNR is directly proportional to its pulse duration. Consequently, through pulse duration, simple pulse systems have an unwanted coupling between received SNR and range resolution. Pulse duration can be increased to increase received SNR, but this reduces range resolution performance, and vice versa.

Pulse compression is a popular technique used to decouple the negative relationship between pulse energy (and consequently received SNR) and range resolution that exists in simple pulse radar systems. By using frequency or phase modulated waveforms, signal bandwidth can be increased while keeping the pulse duration constant [33]. The converse is also true—pulse duration can be increased while maintaining the same signal bandwidth.

Matched filters are commonly used in radar receivers to achieve optimal SNR. When a *pulse compression waveform* (frequency or phase modulated) is passed through a matched filter, the resulting output waveform is compressed in time, hence the name pulse compression. This time compression is proportional to the signal bandwidth of the received pulse. Range resolution is enhanced in this way because as the compressed pulses are shortened in duration, they are less likely to overlap in time and thus easier to distinguish. Figure 2.8 displays a transmitted linear frequency modulated (FM) waveform pulse (also known as a chirp waveform pulse) in part (a) and the resulting "de-chirped" waveform in part (c) after being passed through a matched filter. Linear FM pulses are popular pulse compression waveforms that have linearly increasing frequency over time and a bandwidth of Δf, the total change in frequency throughout the pulse duration. The figure makes it clear that the compressed pulse duration of approximately $\frac{1}{\Delta f}$ is shorter than the transmitted pulse duration.

Consider a pulse signal bandwidth B for a desired range resolution, per Eq. (2.19). The

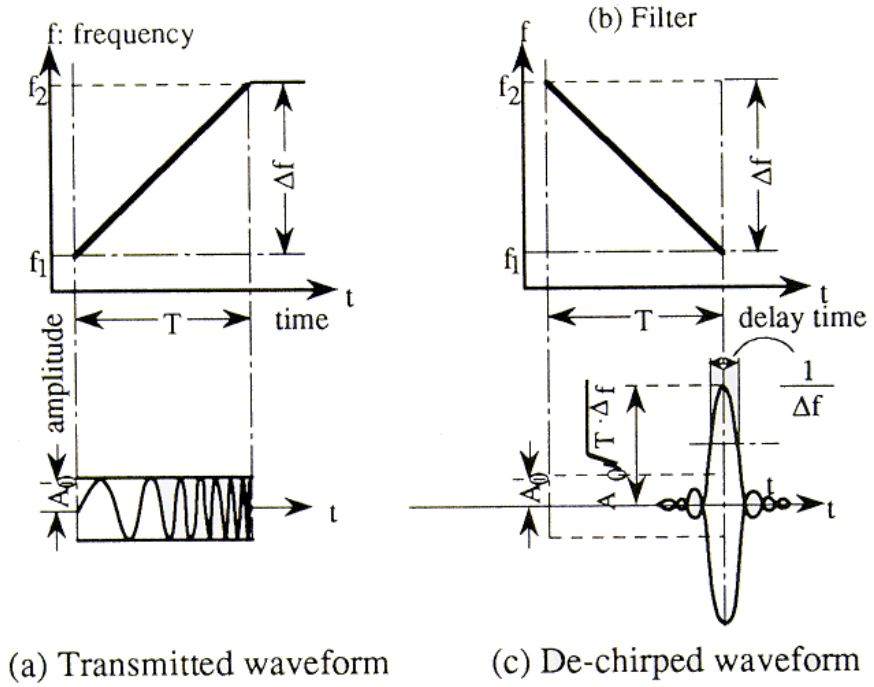

Figure 2.8: Matched filter compression of a linear FM pulse. Retrieved from [10]

pulse duration of a simple pulse system would then be $\tau = 1/\mathrm{B}$. By using a pulse compression waveform instead, pulse duration can be increased to T while maintaining the same signal bandwidth B of the original simple pulse. This increases the transmitted energy and thus received SNR while using the same peak transmit power, which is often limited in radar systems. Since B stays constant, range resolution remains the same. Hence, by using the pulse compression waveform, received SNR is increased without degrading range resolution. The energy and received SNR are increased by a factor of [35]:

$$PCR = \frac{T}{\tau} = TB \tag{2.22}$$

which is termed the pulse compression ratio (PCR). In addition to being the duration ($<$ T) of the simple pulse waveform with equivalent range resolution to the described pulse compression waveform, τ is the approximate duration of the return pulse compression waveform (transmitted with duration T) after being compressed by the receiver's matched filter. This

can be observed in the compressed "de-chirped" waveform of Figure 2.8, whose pulse duration is approximately $\frac{1}{\Delta f}$, which equals τ since the linear FM pulse bandwidth $\Delta f = B = \frac{1}{\tau}$.

2.3 Summary

This chapter explained many key concepts of Wi-Fi networks and radar systems. The Wi-Fi portion (Section 2.1) discussed OFDM, MCS, and PHY frame structure, all part of the Wi-Fi PHY layer. The DCF of the Wi-Fi MAC layer was also described, along with WLAN Architecture for Wi-Fi. In the radar portion (Section 2.2), basic radar measurements and parameters were covered, as well as pulse compression. These principles are referenced in the remainder of this thesis to describe the simulators used in this research and their results, as well as the proposed mitigation schemes.

Chapter 3
LINK SIMULATOR

When a radar pulse interferes with a Wi-Fi signal at its receiver, it distorts the received Wi-Fi signal's amplitude and phase, making it more difficult to decode successfully. This link simulator models this behavior and its effect on Wi-Fi communication performance. It models, at the signal level, a full baseband Wi-Fi transmit and receive chain, a wireless channel, and an interfering radar pulse generator.

This link sim is joint work with Hossein Safavi from University of Washington's Fundamentals of Networking Laboratory. Work is still in progress on it as additions are being made, but its essential features are complete.

3.1 Functional Overview

This section provides a brief overview of the link sim's basic functionality and configuration. Subsequent sections go into more detail about some specific components.

This link sim, written in MATLAB, models a full baseband SISO OFDM Wi-Fi transmitter and receiver as well as a radar waveform generator. The Wi-Fi and radar signals are kept at baseband throughout the simulation—carrier upconversion and downconversion are not modeled. A block diagram of the link sim is shown in Figure 3.1. The six adjacent blocks at the bottom are part of the Wi-Fi receiver model, while the five adjacent blocks above them are part of the Wi-Fi transmitter model. The radar pulse generator is the top right block, and the channel models are labeled between the Wi-Fi transmitter and receiver as well as the radar generator and Wi-Fi receiver.

The simulation starts at the leftmost block of the Wi-Fi transmitter, where random data bits for the data frame are generated and FEC coded using convolutional code based

on the code rate of the specified MCS. Next, the bits are interleaved and mapped to data symbols based on the constellation mapping of the MCS, and pilot symbols are also added. Following this is the preamble generator, after which a 64-point IFFT of the data and preamble symbols is done to OFDM modulate the frame. Digital to Analog Conversion (DAC) is then performed to generate the time domain Wi-Fi frame waveform.

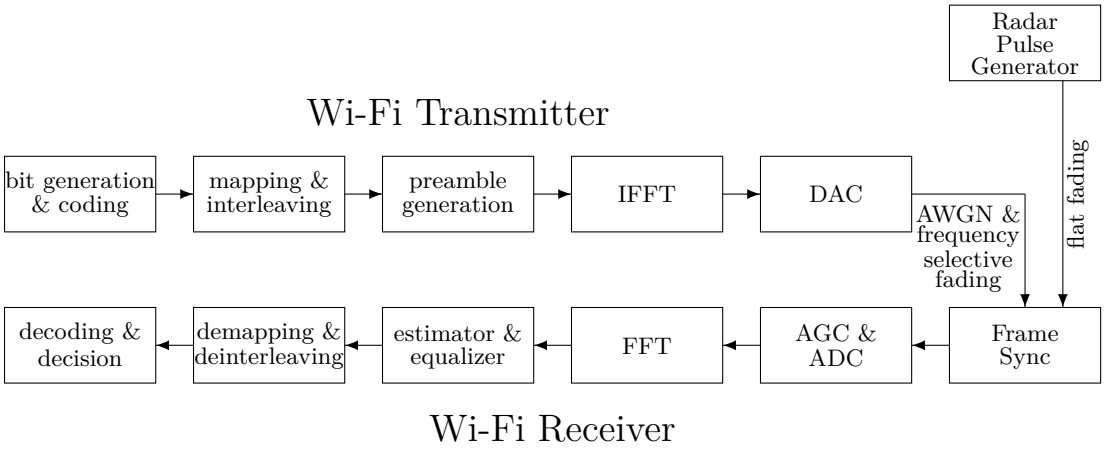

Figure 3.1: Block diagram of link simulator.

The Wi-Fi channel bandwidth configured in the link sim is 20 MHz, which yields a sampling rate of 20 MHz. From the 64-point IFFT and the 20 MHz channel bandwidth, the subcarrier spacing is $\frac{20 \text{ MHz}}{64} = 312.5$ kHz, as there are 64 subcarriers. Of the 64 subcarriers, 52 of them are actually transmitted over. Subcarrier 0 (in the center) is unused as it is a DC null. Subcarriers -26 to -1 and 1 to 26 are occupied (as in legacy mode), with pilots at -21, -7, 7, and 21. The remaining outside subcarriers are left alone as guard bands—6 on one side and 5 on the other. The OFDM guard interval is set to the default of 800 ns for 802.11n, and the OFDM symbol duration is 4 μs.

After the Wi-Fi frame waveform is generated, a channel model is applied to it, with options for AWGN and frequency selective fading. Before the Wi-Fi frame gets received, the radar pulse generator transmits a radar pulse waveform with specified parameters and flat fading is applied. Now both the Wi-Fi frame and radar pulse waveforms are at the

Wi-Fi receiver and are added together. An example waveform of a Wi-Fi frame with radar interference is shown in Figure 3.2.

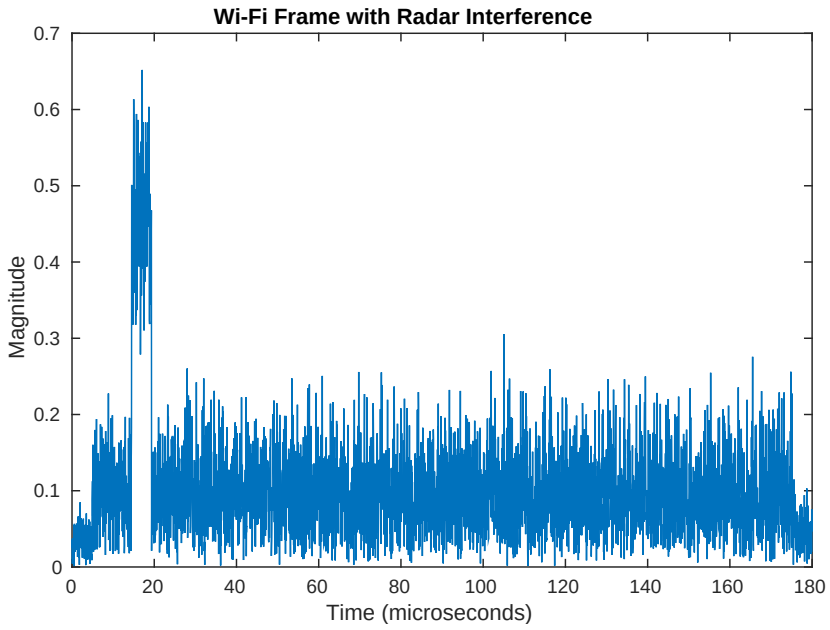

Figure 3.2: Waveform of Wi-Fi frame with interfering 5μs radar pulse at 20 dB INR.

The first thing that the Wi-Fi receiver must do when receiving the frame (with radar interference) is attempt to time synchronize to the frame so that it can be correctly interpreted. After this, the Automatic Gain Control (AGC) is applied and sets the receiver gain based on the first 4 μs of the frame (first half of the L-STF). The Analog to Digital Converter (ADC) follows, bringing the frame back to the digital domain. A 64-point FFT is then done for OFDM demodulation. Next is the channel estimator and equalizer, which attempt to remove some unwanted channel distortions from the data symbols. Constellation de-mapping and de-interleaving occurs, then finally the decoding and decision process to determine the identity of each received data bit. Options for both hard and soft Viterbi decoding are available. A frame error occurs if one data bit fails after receiver bit decisions are made, causing the cyclic redundancy check (CRC) to fail and the frame to be dropped.

3.2 Channel Model

3.2.1 Flat Fading

Flat channel fading is used for the radar pulse waveform. Flat fading is the simplest type of channel model, as it assumes that all frequency components of a signal are attenuated by the same amount. This was chosen to be used for radar because since radar pulses often travel extremely long distances across varying terrain, so many different channel effects can occur that an accurate implementable model may not be feasible without immense complexity. And with the wide array of radar applications, different radar types could experience differing degrees of fading. Additionally, the high directionality of radar antennas makes radar signals less susceptible to multipath effects than the majority of wireless signals.

3.2.2 AWGN

To model channel noise effects, white Gaussian noise is added to the Wi-Fi waveform before it is received. The impact that this AWGN will have on successful frame reception strongly depends on the SNR of the received waveform and the MCS, along with transmitter and receiver chain configurations. In the link sim, the SNR is an input while the Wi-Fi transmit power is fixed. The noise variance can be easily calculated from the SNR and transmit power.

If a data symbol $\mathbf{s}$ is transmitted through the channel, after AWGN the received symbol is simply:

$$\mathbf{y} = \mathbf{s} + \mathbf{n} \tag{3.1}$$

where $\mathbf{n}$ is the Gaussian noise vector with variance σ_{N}^2. $\mathbf{y}$, $\mathbf{s}$, and $\mathbf{n}$ are all complex valued, as I/Q constellation points are represented in the complex plane.

3.2.3 Frequency Selective Fading

The transfer function of a frequency selective fading channel varies with frequency—signals propagating through the channel at different frequencies will experience different magnitudes of fading. A common cause of frequency selectivity in a wireless channel is multipath propagation. Frequency selective fading is common in real-world scenarios, especially in enclosed indoor environments where there are many walls and ceilings that provide various paths for signals to take to get to the receiver.

In 2004, the 802.11 High Throughput Task Group (TGn) provided a set of frequency selective channel models for Wi-Fi network modeling and analysis. These models are based on the Rician fading model, which is statistical model for channels with one dominant Line of Sight (LOS) propagation path and multiple Non Line of Sight (NLOS) paths.

Table 3.1: TGn channel model parameters [18].

Model	**K (dB) (LOS)**	**NLOS RMS delay spread (ns)**	**Number of clusters**
A	0	0	1
B	0	15	2
C	0	30	2
D	3	50	3
E	6	100	4
F	6	150	6

Table 3.1 contains the TGn channel models and some of their important parameters. K is the Rician K-factor, which is the ratio of signal power received from the LOS path to the power from the NLOS paths. A set of scattered paths with different delay times are specified for each model, and the table lists the root mean square (RMS) of the delay spread of these NLOS paths. Also included is the number of clusters, where a cluster is a group

of adjacent-delay propagation paths with very similar angular spreads, arrival angles, and departure angles.

TGn channel model B is implemented in the link sim. It was chosen because it is the middle tier model for a weak LOS scenario (as the K-factor is 0 dB), so as to provide moderate intensity selective fading. A weak LOS component is typical for many home and office scenarios, where there is often no direct path from node to node that is not through walls. TGn lists channel model B to be useful for office scenarios [18].

3.3 Wi-Fi Model

3.3.1 Interleaving

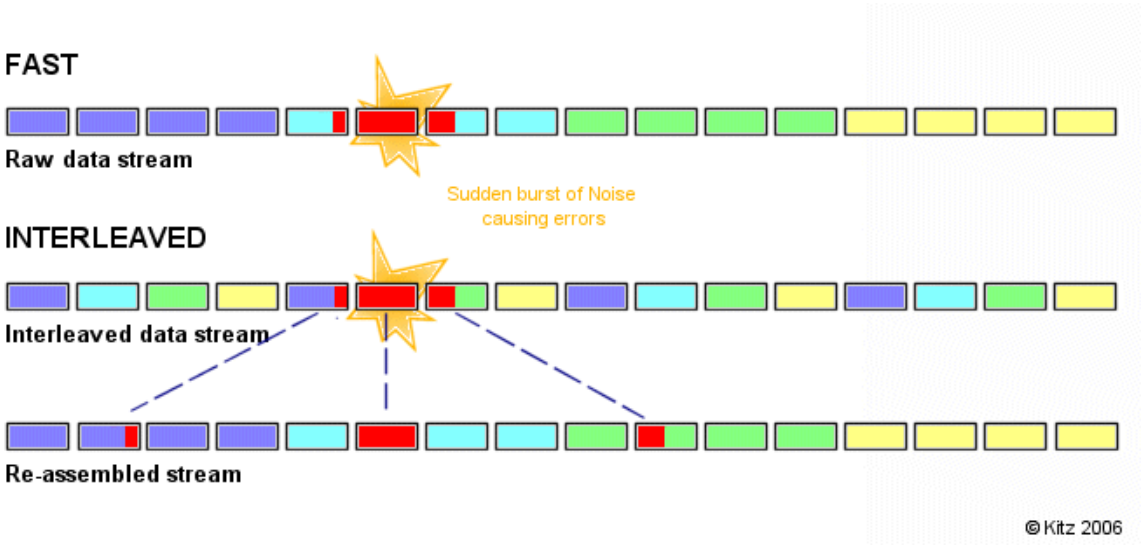

Figure 3.3: Example of interleaving and its usefulness in mitigating the impact of interference. Retrieved from [11]

Interleaving is the practice of switching the order of adjacent bits in some fashion to provide added robustness against receive errors. It is typically done over time or frequency. Figure 3.3 shows a comparison of interleaved data to normal data in sequence. The bit stream entitled "FAST" has the data in its original sequence, with bits of the same color

corresponding to the same data block. The interleaved stream on the other hand has its sequence switched so that adjacent bits are from separate data blocks. After arriving at its destination the interleaved bit stream must eventually be re-assembled back to its original order to be used for its purpose. When a burst of noise interferes with the streams, the received data is affected differently for both. All three bits interfered with in the un-interleaved stream are part of the same data block, making it likely that its reception will fail. But when the interleaved stream gets hit with the same noise burst, three bits still experience interference, but these bits are from different data blocks. Therefore, none of the data blocks from the re-assembled stream have as high a likelihood of failure as the blue data block in the un-interleaved stream.

The example explained above can apply to both time and frequency interleaving. As time interleaving, each bit would be in a time sequence from left to right. The original stream is labeled "FAST" because it suffers less delay than the interleaved stream, since the entire interleaved stream must be received before any of the data blocks can be re-assembled. This is one of main tradeoffs of interleaving.

If the Figure 3.3 example is interpreted as frequency interleaving, each bit would be on a different subcarrier. Wi-Fi only uses frequency interleaving. By spreading the bits of an FEC codeword onto separate subcarriers, robustness against interference is added because in many situations not all the subcarriers will be hit with interference at the same time.

3.3.2 Preamble

The 802.11n HT Data rates of Table 2.1 are used in this link sim, which yield either mixed mode or green field mode PHY frame structure. Green field mode is not commonly used so mixed mode makes more sense to simulate. But since the link sim model is solely SISO and in mixed mode the HT portion of the preamble (see Figure 2.5) is used for MIMO capabilities, they are not needed in our simulation and are not used. So by removing the HT preamble portions, the preamble is essentially just a legacy preamble–with L-STF, L-LTF, and L-SIG.

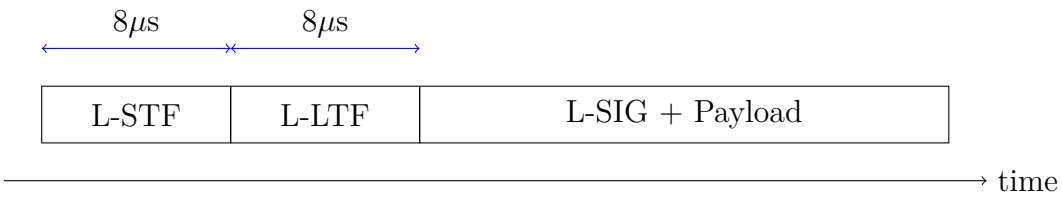

Figure 3.4: Wi-Fi link sim PHY frame structure.

The L-STF is the first 8 μs of the frame. In the link sim, the first half of it is used for AGC setting and the entirety of it is used for frame synchronization. The L-LTF takes up the next 8 μs and is used for channel estimation in the link sim. The L-SIG's functionality is not currently simulated in the link sim, but its 4 μs duration is taken into account, making the simulated preamble 20 μs in duration. Regarding radar interference location, the L-SIG portion is treated as the payload. Figure 3.4 shows the simulated PHY frame structure.

The preamble always has the same constellation mapping and code rate no matter what MCS is set for the Wi-Fi frame. This is why only one set of frame synchronization results is necessary in Section 3.5.

3.3.3 Frame Synchronization

Frame time synchronization (hereinafter *frame sync*) is an extremely important part of a Wi-Fi receiver. If a receiver fails to correctly synchronize with a frame, it has no hope to recover the frame's data because it will not be interpreting it correctly. A two-step frame sync method which the author deems *delayed autocorrelation* is used in the link sim to model frame finding. This method, detailed in [36], involves correlating the L-STF with delayed copies of itself, first using a coarse timing synchronization followed by a finer one. The autocorrelation helps mitigate channel effects, since the copies are all affected by the channel in the same way.

Frames are always dropped when failing to be synchronized, so when generating the link sim results of Section 3.5, the frame sync simulations using the delayed autocorrelation

method were kept separate from the normal frame error rate simulations. This was done so that the frame error rate could be calculated accurately, without depending on frame synchronization success. The normal frame error rate results were generated using ideal frame sync so that the receivers would always sync successfully. As this delayed autocorrelation frame sync method primarily uses the L-STF, radar pulses that interfere with the L-STF are considered in the frame sync error results.

3.3.4 Channel Estimation

Two wireless channel estimator types are currently available in the link sim. The estimator determines an estimate of the channel response over all the OFDM subcarriers. The equalizer then applies it to received data symbols in attempt to reverse the fading effects of the channel and improve correct bit decision probability.

Ideal Estimator

The link sim has the option of being configured with an ideal channel estimator. The ideal estimator perfectly estimates the channel transfer function. It is impossible to have an ideal estimator in actual practice, but it is useful in simulation for obtaining an upper bound on estimator performance.

MMSE Estimator

A minimum mean square error (MMSE) estimator is implemented in the link sim. The MMSE provides an estimate [7]:

$$\hat{\mathbf{H}} = \mathbf{W}\tilde{\mathbf{H}} \tag{3.2}$$

where $\hat{\mathbf{H}}$ is the MMSE channel estimate, $\tilde{\mathbf{H}}$ is the raw channel estimate (received L-LTF symbols divided by corresponding known transmitted L-LTF symbols, see Equations 6.7 and 6.11 in [7]), and $\mathbf{W}$ is a weighting matrix. Now, with actual channel response $\mathbf{H}$, the

error betweeen the MMSE estimate and the real channel response is [7]:

$$\mathbf{e} = \mathbf{H} - \hat{\mathbf{H}} \tag{3.3}$$

which means the mean squared error (MSE) is simply [7]:

$$MSE = E[||\mathbf{e}||^2] = E[||\mathbf{H} - \hat{\mathbf{H}}||^2] \tag{3.4}$$

The MMSE estimator uses the known (in practice it will have an estimate of these parameters) noise variance and channel power delay profile to adjust weighting matrix $\mathbf{W}$ in attempt to minimize the mean squared error of Eq. (3.4). See Equations 6.12 - 6.14 of [7] for details.

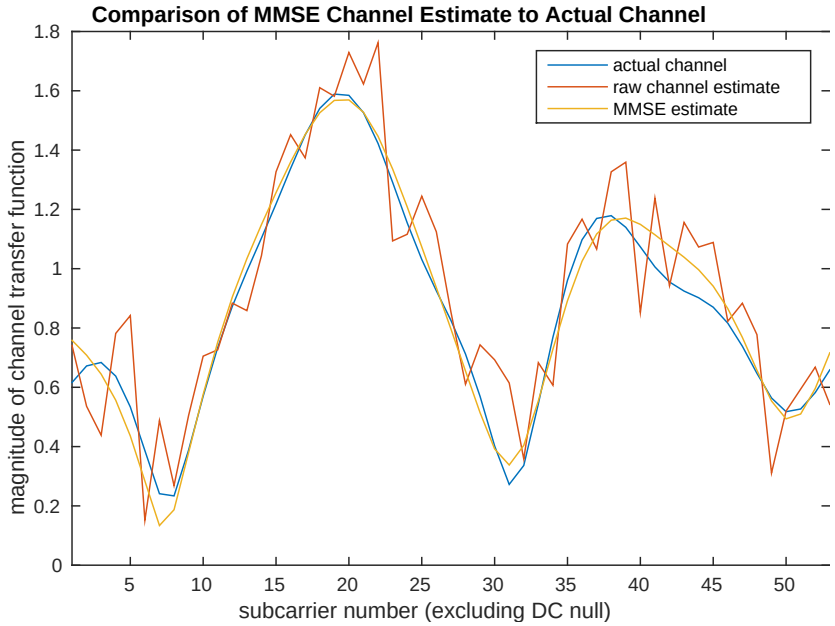

Figure 3.5: Comparison of actual channel transfer function, raw channel estimate, and MMSE estimate for 8 dB SNR with channel model B.

A comparison of the MMSE estimate in the link sim to the raw channel estimate (directly from subcarriers) and the actual channel transfer function is shown in Figure 3.5. The frequency selective channel response is a realization of TGn channel model B. The raw estimate's jagged edges, peaks, and troughs come from AWGN. In most locations, the

MMSE estimator smooths down the raw estimate to very close to the actual channel transfer function. When determining its estimate for a given subcarrier, the MMSE estimator actually chooses its weighting based on all the subcarriers, though with higher weighting for subcarriers closer in frequency. This helps produce the smoothing effect seen in the figure.

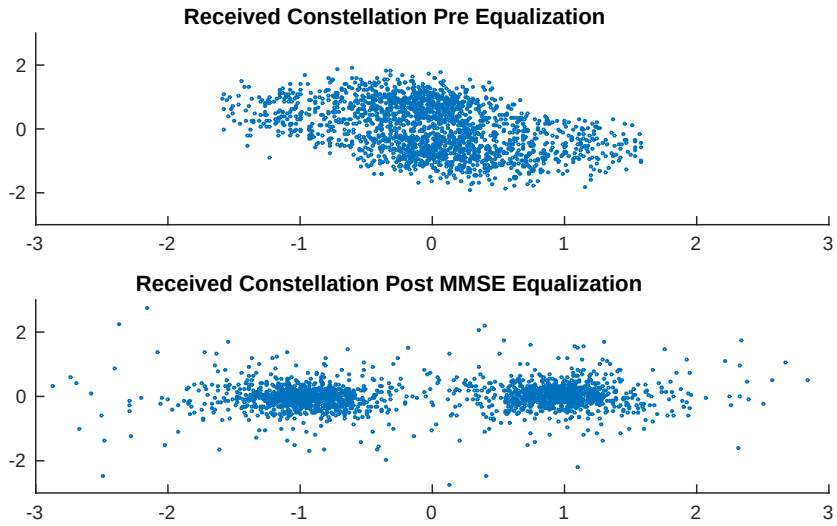

Figure 3.6: Data symbol constellation diagram before and after MMSE equalization for 8 dB SNR and MCS 0 (BPSK) with channel model B. Same simulation instance as that of Figure 3.5.

Figure 3.6 displays the received data symbol constellation points before and after MMSE equalization for the same simulation instance as Figure 3.5. MCS 0 was configured, which uses BPSK, and it is clearly shown in the figure that after equalization the received symbols take on a shape much more indicative of BPSK and that points were shifted out of the center zone. This should increase the successful decision probability.

The MMSE estimator in the link sim takes its raw measurements from the known symbols of the L-LTF. Thus, when an interfering radar hits the L-LTF, the channel estimate is perturbed.

3.4 Radar Waveform Model

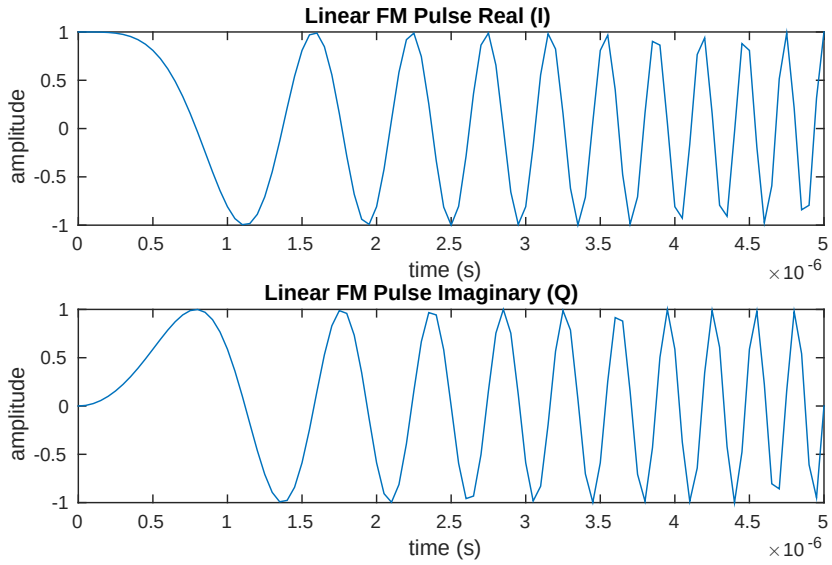

Figure 3.7: Time plot of 0–4 MHz, 5 μs linear FM waveform pulse with unit amplitude.

Linear FM waveforms are the pulse types used in the link sim radar model. Briefly discussed in Subsection 2.2.3, linear FM waveforms, also known as "chirp" signals, are very commonly used in radar systems for pulse compression. The frequency of a linear FM waveform is modulated linearly over time. Defining f_0 as the starting frequency (frequency at the start time of the pulse) and f_1 as the end frequency, the chirp rate k is:

$$k = \frac{f_1 - f_0}{\tau_r} \tag{3.5}$$

where τ_p is the radar pulse duration.

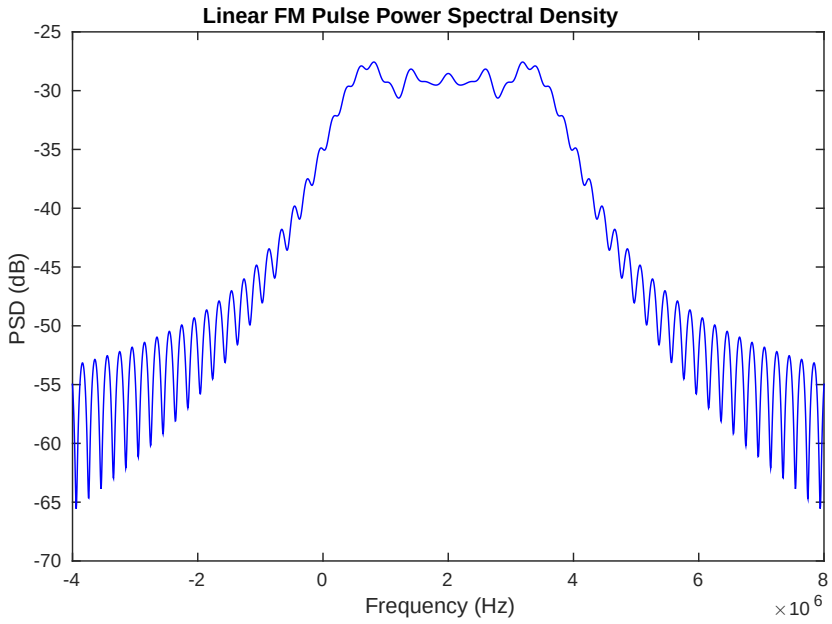

Figure 3.8: PSD plot of 0–4 MHz, 5 μs linear FM waveform pulse with unit amplitude.

A complex chirp pulse can then be expressed as:

$$y(t) = Ae^{j[2\pi(f_0 t + \frac{k}{2}t^2)+\phi]} \tag{3.6}$$

where A is the magnitude and ϕ is the phase offset. For the link simulation results, the linear FM waveform had a 5 μs pulse duration and a chirp range (frequency range) of 0–4 MHz. This waveform and its PSD are depicted in Figures 3.7 and 3.8. In the link sim, the radar power is not explicitly set by the user. Instead, the received Interference-to-Noise Ratio (INR) is configurable and is defined with respect to the Gaussian noise at the Wi-Fi receiver (which is used in the Wi-Fi SNR calculations). Adjusting the INR effectively adjusts the radar power. For each Wi-Fi frame simulation up to one radar waveform can be generated.

3.5 Results and Analysis

Before getting into radar interference scenarios and results, it was helpful to run Wi-Fi only simulations to validate the Wi-Fi model itself. Figure 3.9 contains the frame error rate

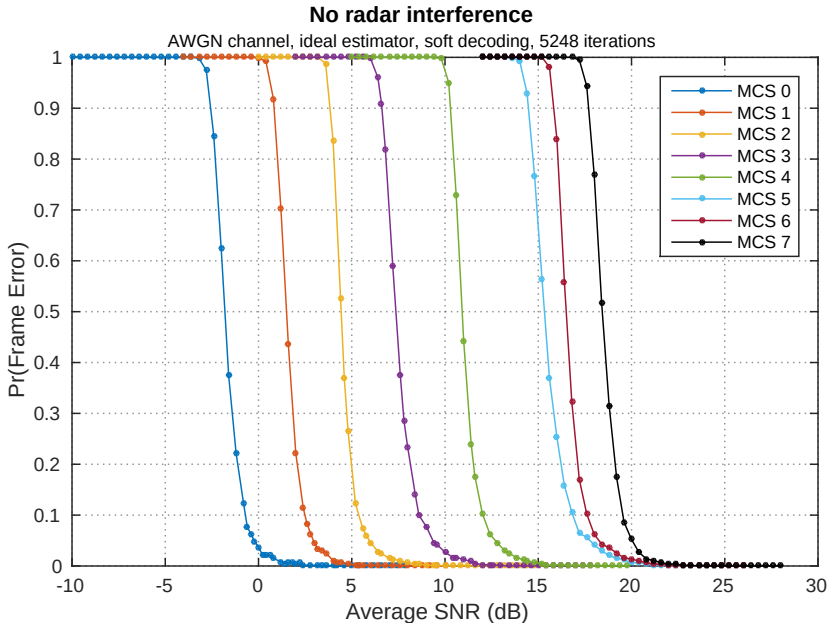

Figure 3.9: Frame error rate vs. SNR for MCS 0–7 with no radar interference using AWGN only channel, ideal estimator, soft decoding, and 180 μs frame duration.

(FER) versus average received SNR (hereinafter referred to as simply "SNR") for each SISO MCS using an AWGN only channel, ideal estimator, and soft decoding. The figure also says "5248 iterations" meaning that many iterations were averaged over for each data point. These curves make intuitive sense, as each subsequent increasing MCS performs noticeably worse than the previous. Though it is interesting that MCS 5 and MCS 6 are closer in FER performance than any other adjacent MCS pair.

The rest of the FER link sim results are only using MCS 0 or MCS 7, as they are the highest and lowest SISO MCS settings and are the ones simulated in the system simulation (Chapter 4), which these link sim results will be mapped to. As mentioned in 3.3.3, the frame sync results are separated from the general FER results. The frame sync error rate results use delayed autocorrelation frame sync and the FER results use ideal frame sync as they assume that frame sync successfully occurred already. All simulated data frames use 40 payload OFDM symbols which, along with the preamble, yield 180 μs frame duration. Soft

decoding is used in all the simulations, because it provides much better performance than hard decoding in frequency selective channels.

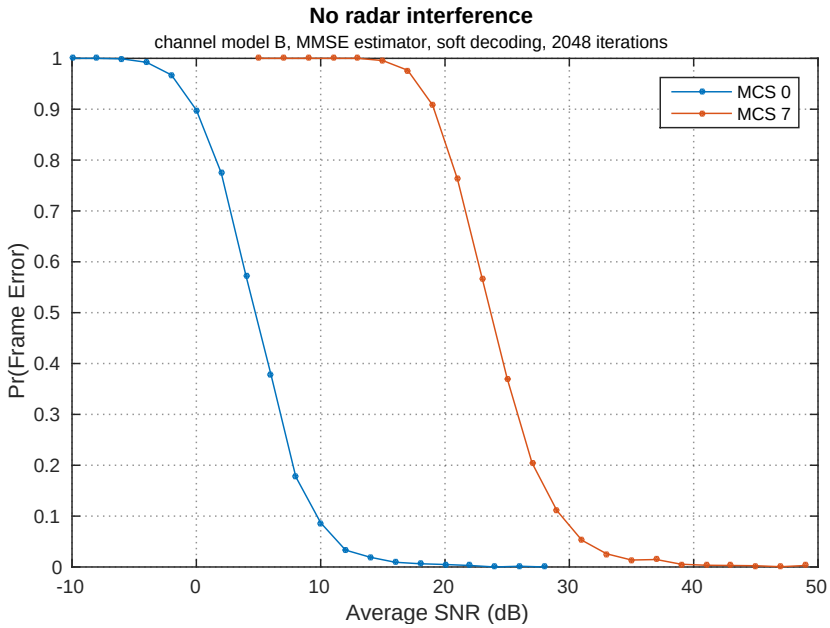

Figure 3.10: Frame error rate vs. SNR for MCS 0 and MCS 7 with no radar interference using channel model B, MMSE estimator, soft decoding. interleaving, and 180 μs frame length.

Figure 3.10 depicts the FER versus SNR curves for MCS 0 and MCS 7 in channel model B. *Compared to the AWGN results, the performance is significantly worse with channel model B as the MCS 0 and MCS 7 zero FER crossings are almost 20 dB higher.* This disparity is primarily because the frequency selective fading of channel model B is much more difficult to compensate for at the Wi-FI receiver than AWGN. These results are also using the MMSE estimator instead of the ideal channel estimator, which reduces performance. The results of Figure 3.10 are in the vicinity of those in Figure 5.13 of [4], which helps validate these link sim results. The results from this thesis have slightly better performance, which is expected because the simulations in the cited figure have PHY impairments that are not modeled in this work.

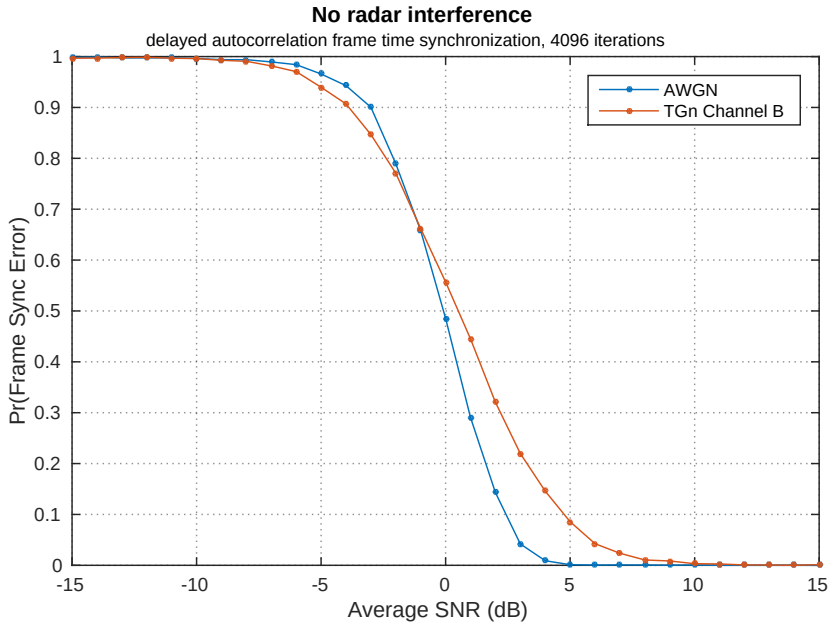

Figure 3.11: Frame synchronization error rate vs. SNR with no radar interference using delayed autocorrelation frame synchronization method. Plots shown for AWGN and channel model B.

The delayed autocorrelation frame sync results for no radar interference are in Figure 3.11 for both AWGN and channel model B. It is clear that *this frame sync technique is very resilient to channel effects, as the probability of frame sync error for channel model B performs better than MCS 0 data reception*, achieving zero percent error at 10 dB instead of the 20 dB required for MCS 0 in Figure 3.10.

The remainder of the link sim results in this section are configured with the settings in Table 3.2. The next plot, Figure 3.12, is a three dimensional error surface showing the frame sync error rate as a function of SNR and radar INR. It is important to recognize that in the scenarios generating this curve, the radar pulses only interfered with the L-STF, which is the portion of the preamble used for frame sync. Thus, frame sync performance suffers with increasing radar INR. For 40 dB INR, the frame sync error rate is high even halfway up the SNR range, while without radar the error rate is virtually 0 above 10 dB SNR.

Table 3.2: Settings used for all link simulation 3D error surface results (when parameters are applicable).

Radar		**Wi-Fi**	
Waveform	Linear FM	Interleaving	Yes
Chirp Range	0 - 4 MHz	Channel Model	B
Pulse Duration	5 μs	Estimator	MMSE
Phase	Random	Decoding	Soft
Frequency Offset	None	Frame Length	180 μs

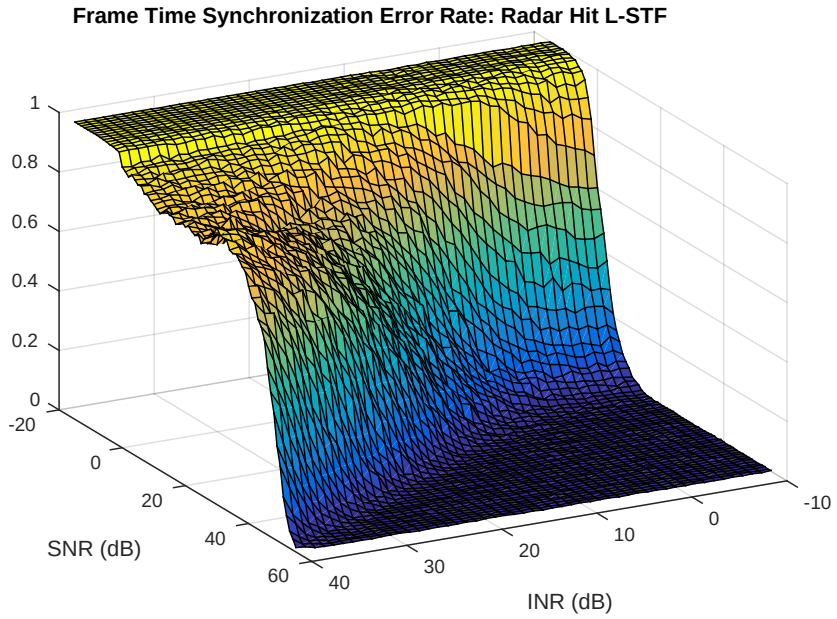

Figure 3.12: Frame synchronization error rate vs. INR and SNR (1024 iterations) for radar interfering with L-STF using delayed autocorrelation frame synchronization method and channel model B with radar settings in Table 3.2.

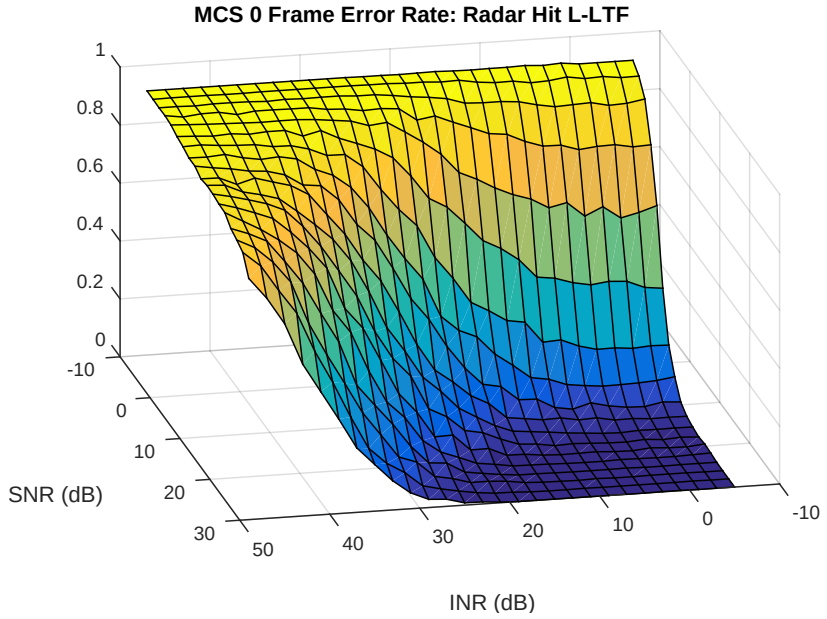

Figure 3.13: MCS 0 frame error rate vs. INR and SNR (1024 iterations) for radar interfering with L-LTF. Configured with settings in Table 3.2.

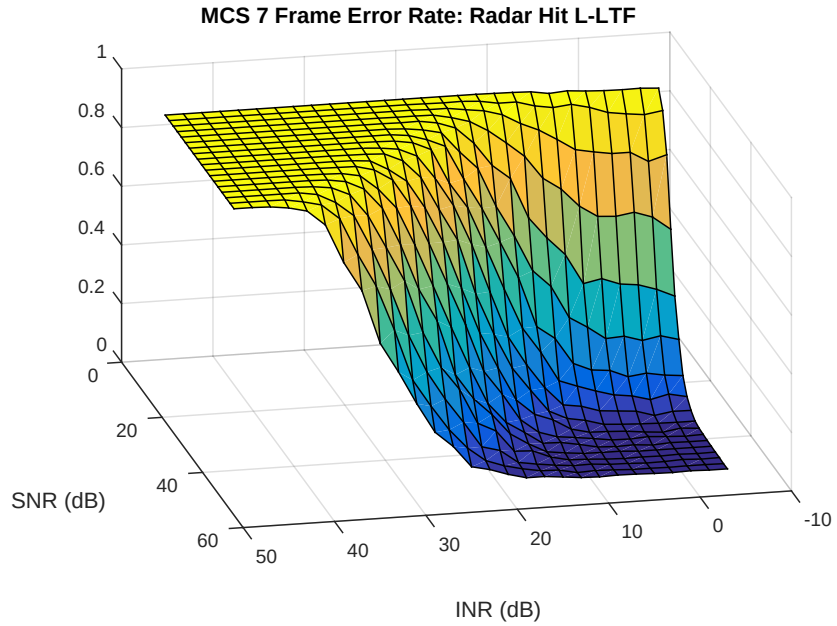

Figure 3.14: MCS 7 frame error rate vs. INR and SNR (1024 iterations) for radar interfering with L-LTF. Configured with settings in Table 3.2.

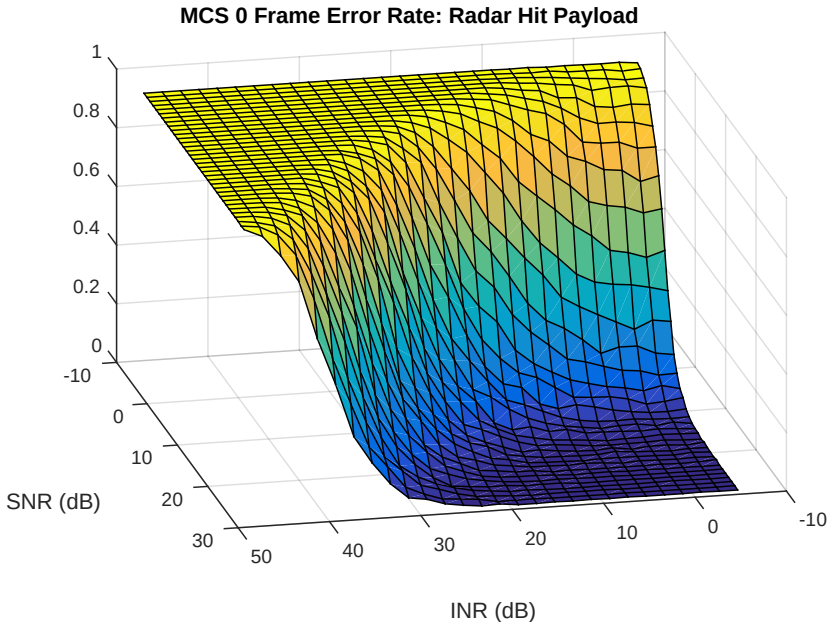

Figure 3.15: MCS 0 frame error rate vs. INR and SNR (1024 iterations) for radar interfering with payload. Configured with settings in Table 3.2.

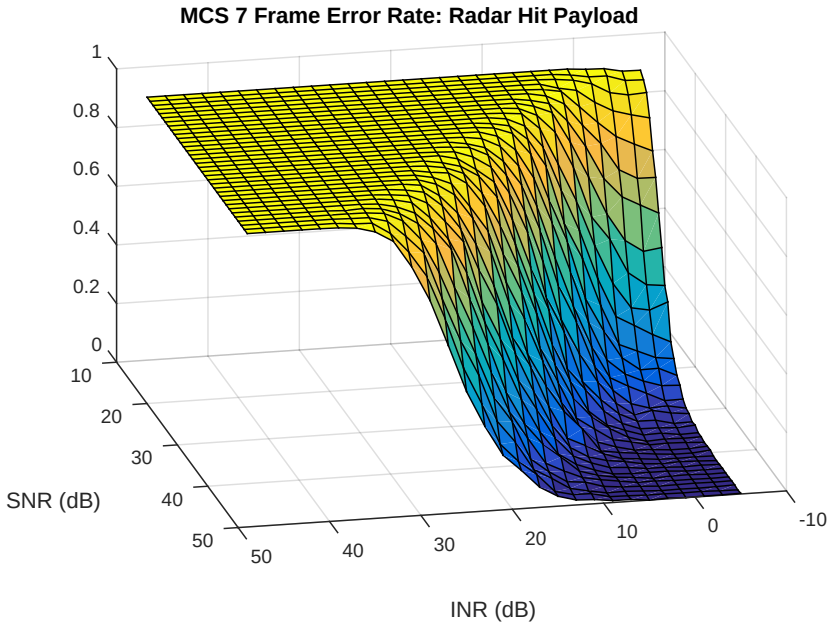

Figure 3.16: MCS 7 frame error rate vs. INR and SNR (1024 iterations) for radar interfering with payload. Configured with settings in Table 3.2.

The remaining results are three dimensional error surfaces showing FER as a function of SNR and INR for MCS 0 and MCS 7 and for two different radar hit locations: L-LTF and payload. As expected, *MCS 0 is more resilient to radar interference than MCS 7 due to its lower modulation order (BPSK vs. 64-QAM) and code rate (1/2 vs. 5/6)*, as the FERs in the MCS 7 surfaces increase more quickly with INR than those in the corresponding MCS 0 surfaces. But *in all situations, high enough radar INR can even cause frames with extremely high SNR to fail reception the majority of the time.*

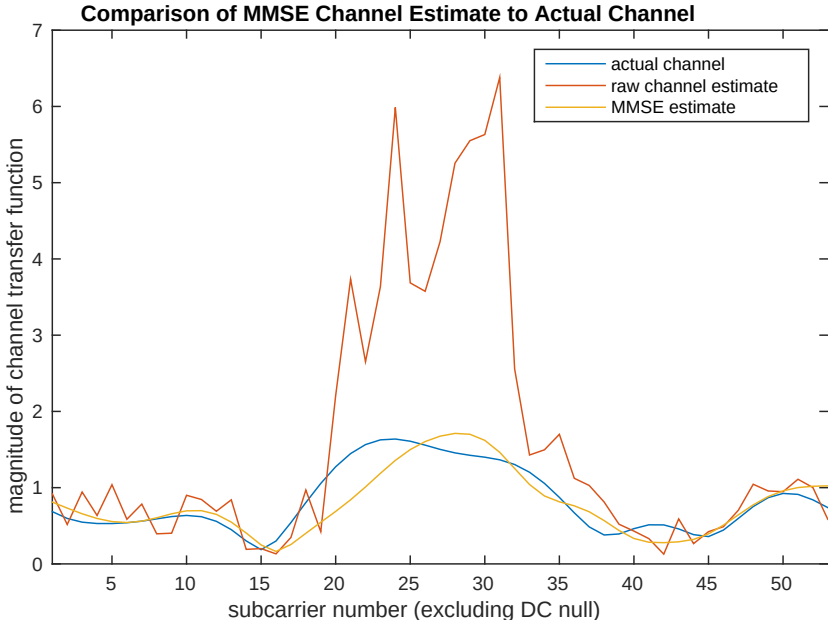

Figure 3.17: Comparison of actual channel transfer function, raw channel estimate, and MMSE estimate for 8 dB SNR with channel model B where radar interferer hits L-LTF with 20 dB INR. MMSE estimate is still relatively accurate even with massive interference. See Subsection 3.3.4 for an overview of the MMSE estimator and Section 6.2 of [7] for further details.

It can be observed *in the error surfaces above that the situation where the radar pulse hits the L-LTF performs better than that of radar-payload interference*, as the hit payload surfaces reach an FER of 1 faster. Recall that the L-LTF is the portion of the frame used for channel estimation. This result surprised the author at first, because his intuition was that

radar hitting L-LTF would be worse due to yielding a poor channel estimate. The channel estimate is applied to all symbols in the frame, so if a terrible channel estimate is taken, the frame is very likely to fail. However, *the MMSE estimator is actually quite resilient to radar interference.*

Figure 3.17 displays the MMSE channel estimate for 8 dB SNR and 20 dB INR with the radar pulse hitting the L-LTF. It is strongly affecting the raw channel estimate which strays very far from the actual channel transfer function. But the MMSE estimator still maintains a relatively close match with the true channel. One reason behind this is the fact that for each subcarrier's estimate, the MMSE provides weighting from all subcarriers, not just adjacent ones, and the radar pulse is not interfering with all the subcarriers. In addition, as the L-LTF is 8 μs long and the OFDM symbols are 4 μs, L-LTF uses two OFDM symbols and it uses that whole period to make its estimates. The 5 μs may only interfere with one of the OFDM symbols, still leaving a full one undisturbed to contribute to a good MMSE estimate. Lastly, the simulated MMSE estimator is assumed to know the exact noise variance and channel power delay profile, when in reality it would have an estimate. This grants it optimal performance, whereas in practice it would not be perfectly optimal. The L-LTF-interfering radar pulse would surely cause a worse MMSE channel estimate if it was more wideband such that it interfered with all the subcarriers, or if its duration was long enough to interfere throughout the entire two OFDM symbol duration of the L-LTF.

Chapter 4
SYSTEM SIMULATION

The system simulations performed in this research provide system level performance results of Wi-Fi networks with radar interference. The primary metric that the author is interested in is Wi-Fi network throughput (the rate of successfully communicated useful data bits per unit time), which is measured in numerous simulation scenarios. The Wi-Fi module of the ns-3 network simulator is used to run the simulations, and modifications to it were made to model radar interference. The error rate results from Chapter 3 are mapped to the system simulation to provide detailed Wi-Fi frame reception probabilities with radar interference.

The layout of this chapter is presented as follows. First discussed is an assumption made regarding radar beam coverage, along with the path loss model used in the simulations. Next is a short overview of the ns-3 Wi-Fi module, followed by a description of the modifications made to it. Some proposed mitigation schemes are then presented. The next sections describe the parameters and scenarios used in the simulations, while the final section contains simulation results and analysis.

4.1 Radar Beam Coverage

A key assumption made in the system simulator is that a radar pulse covers the entire simulated network. In this section, the percentage of radar beam coverage over a Wi-Fi network's area is analytically expressed. This is used to justify the stated assumption.

Upon transmission, an antenna radiates power into its surrounding space. The magnitude of the propagated signal in each direction around the antenna depends on characteristics of the antenna. Isotropic, or omnidirectional, antennas propagate signals with the same

magnitude in all directions around the transmitter. True isotropic antennas do not exist, but in certain applications such as in Wi-Fi transmitters and receivers, isotropic-like antennas are used in attempt to provide coverage in all directions. On the other hand, directional antennas propagate the majority of a signal's energy in a specific direction. Radar systems typically use highly directional antennas in order to determine the direction of detected objects as the antenna rotates. The following discussion is limited to the horizontal (azimuthal) plane, from which all the mentioned parameters will be referenced, because only the horizontal plane is modeled in the Wi-Fi-radar system simulator used in this research.

An important characteristic of a directional antenna is its beam width. The horizontal beam width is the azimuth angle range (in degrees) that is within -3 dB of the maximum transmit antenna gain. We define the horizontal beam radius at any given time as half of the width of the maximum magnitude portion of the signal, which is the portion of the signal transmitted from within the beam width. Any object within the beam diameter (two times the beam radius) of the signal will be encapsulated by the maximum magnitude beam if it comes into contact.

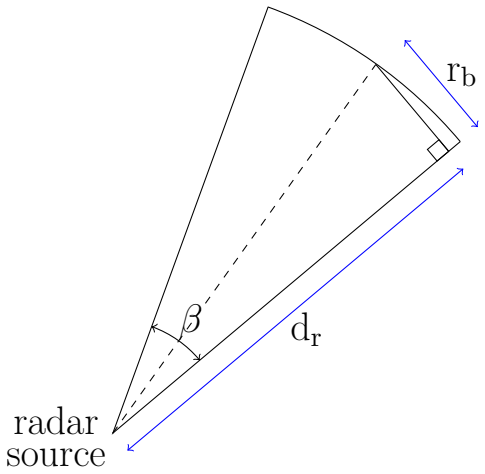

Figure 4.1: Relationship between horizontal beam radius r_b and propagation distance d_r of a radar pulse for a given horizontal beam width β. Top-down perspective.

The horizontal beam radius (and thus diameter) increases the farther the signal propa-

gates, which is shown in Figure 4.1. This results in the signal energy being spread across a wider azimuthal radius. Define the horizontal beam slope as the ratio of horizontal beam radius increase to propagation distance increase. From the figure, the beam slope can be written as:

$$S_b\ (horizontal) = \frac{r_b}{d_r} = tan\left(\frac{\beta}{2}\right) \tag{4.1}$$

where r_b is the approximate beam radius, d_r is the distance from the transmit source, and β is the horizontal beam width. Using this, the horizontal beam diameter is:

$$beam\ diameter\ (horizontal) = 2r_b = 2d_r \cdot S_b = 2d_r \cdot tan\left(\frac{\beta}{2}\right) \tag{4.2}$$

which provides the relationship between beam diameter, propagation distance, and beam width.

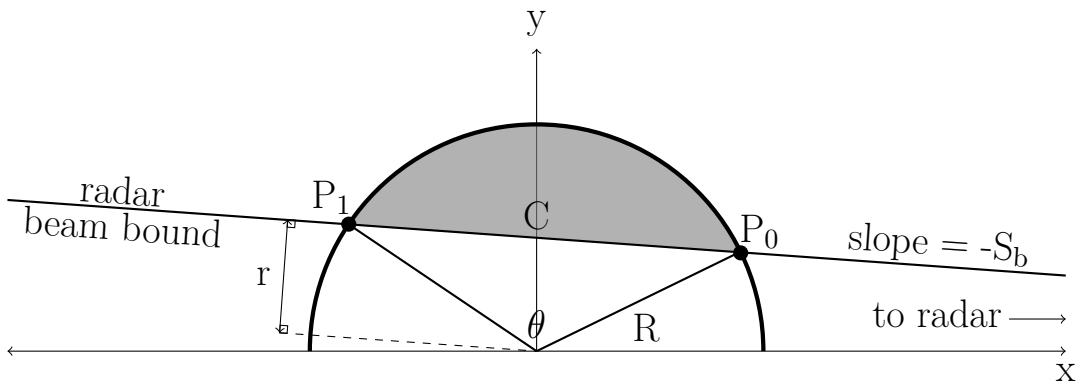

Figure 4.2: Radar beam bounding line intersecting with circular WLAN area with radius R. P_0 and P_1 are the intersection points. Gray shaded area is not covered by the radar beam.

Consider a WLAN with a circular area centered around origin point (0,0) in an xy Cartesian coordinate system. If it has a radius R, the equation of the upper half of the area circle is:

$$y = \sqrt{R^2 - x^2} \tag{4.3}$$

Now consider a radar source at some distance d_r from the center of the circular WLAN area. We will assume the radar source is aligned with the x-axis of our system for convenience,

but the results of this analysis will hold regardless of the orientation of the radar source with respect to the WLAN. With the radar source at coordinate (d_r, 0), from the beam slope of Eq. (4.1), the equation of the upper half radar beam bounding line is:

$$y = -S_b x + d_r S_b \tag{4.4}$$

Plugging a value of x into Eq. (4.4) gives the horizontal beam radius at distance x away from the WLAN center. There also exists a lower beam bounding line which is symmetric across the x-axis. The bounding lines are the edges of the maximum magnitude portion of the radar signal (transmitted from within the beam width). Since both the radar beam bounding lines and the WLAN circular area are symmetric, this analysis is performed on only the positive y field (quadrants I and II) but can be applied to the lower y field as well.

Assuming the radar source is farther away from the WLAN center than the WLAN radius, the upper radar beam bounding line will become a secant line and intersect the WLAN circle at two points, unless the radar source is far enough away that the upper beam bounding line has a higher y value within the WLAN radius x range. The intersection points can be calculated by equating Eqs. (4.3) and (4.4), which after rearrangement yields:

$$(S_b^2 + 1)x^2 - 2S_b^2 d_r x + (d_r^2 S_b^2 - R^2) = 0 \tag{4.5}$$

Solving Eq. (4.5) for x gives the x coordinates of the two intersection points. To calculate the y coordinates of the intersection points, these x values can simply be plugged into Eq. 4.3.

Figure 4.2 depicts the described scenario. The semicircle bounds half of the WLAN area, the line intersecting the semicircle is the upper radar beam bounding line, and P_0 and P_1 are the two intersection points. To determine the percent coverage, the *coverage area* needs to be calculated, which is the area of the WLAN that is covered by the radar beam. In the figure, the half coverage area (since it only shows half the WLAN circle area) is the unshaded part of the semicircle. Therefore, the half coverage area can be calculated as half the WLAN area minus the area of the gray shaded region. This gray section is a type of region commonly referred to as a circular segment.

In Figure 4.2, the radar beam bounding line forms a chord between the intersection points. The chord length C is simply the distance between P_0 and P_1. The height of the triangular region r is calculated as [37]:

$$r = \frac{1}{2}\sqrt{4R^2 - C^2} \tag{4.6}$$

The central angle θ is then [37]:

$$\theta = 2 \cdot arctan\left(\frac{C}{2r}\right) \tag{4.7}$$

The area of the shaded segment is the area of the sector wedge (triangular region and segment combined) minus the area of the triangular region, which is [37]:

$$A_{seg} = \frac{1}{2}R^2(\theta - sin\theta) \tag{4.8}$$

The half coverage area is the difference between the area of the semicircle and A_{seg}. Since the WLAN circle and radar bounding beam is symmetric across the x-axis, the full coverage area is:

$$A_c = A_{circle} - 2A_{seg} = \pi R^2 - R^2(\theta - sin\theta) \tag{4.9}$$

Finally, the *coverage percentage*, the percentage of the WLAN network area covered by the radar beam is:

$$coverage\ percentage = \frac{A_c}{A_{circle}} \times 100 = \frac{\pi R^2 - R^2(\theta - sin\theta)}{\pi R^2} \times 100 \tag{4.10}$$

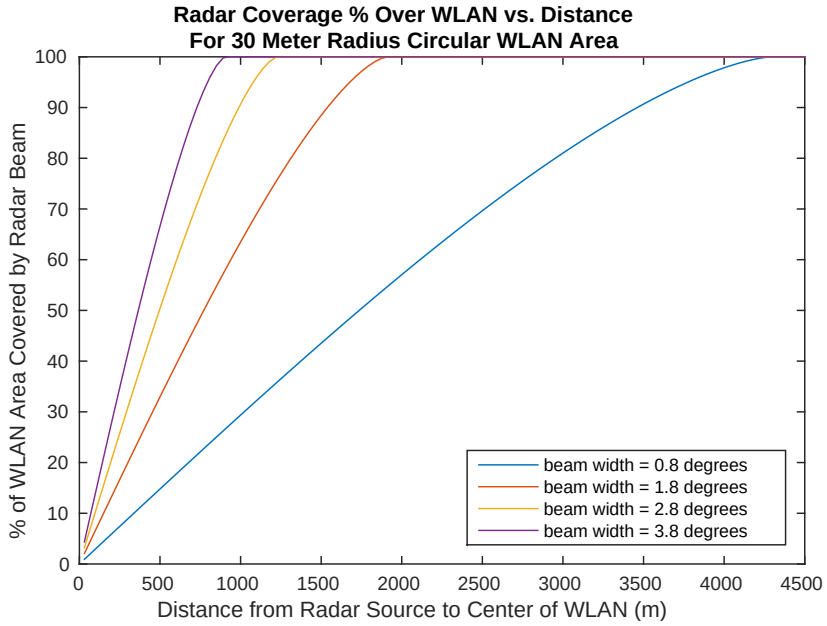

Figure 4.3: Percentage of WLAN area covered by radar beam vs. distance. Fixed 30 meter radius horizontal circular WLAN area and varying horizontal radar beam width.

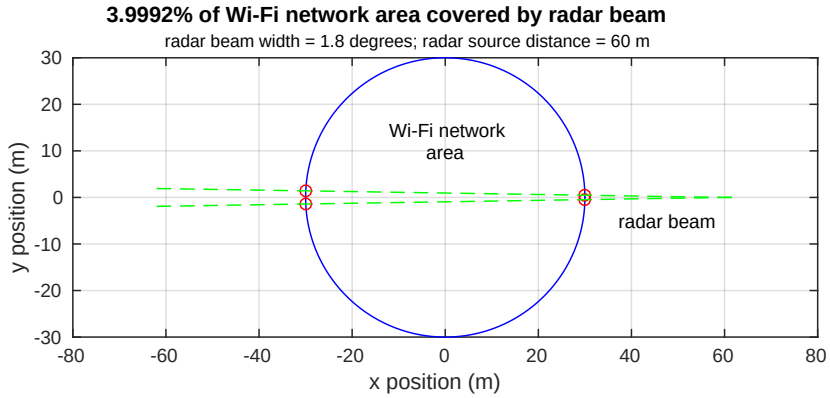

Figure 4.4: Radar beam coverage over 30 m radius Wi-Fi network 60 meters away.

Figure 4.3 shows, for various radar beam widths, the coverage percentage as a function of distance (d_r) to the WLAN center for a 30 meter radius horizontal circular WLAN area. As expected, a higher beam width yields a lower distance required for maximum area coverage. Also, the coverage percentage per distance change increases faster for larger beam

Table 4.1: Horizontal radar beam width vs. minimum distance from WLAN center for full area coverage of 30 meter radius horizontal circular WLAN area by radar beam.

beam width (degrees)	**min. distance for full coverage (m)**
0.8	4298
1.8	1910
2.8	1228
3.8	905
4.8	716
5.8	593
6.8	505
7.8	441
8.8	390
9.8	350
10.8	318

widths. This can also be observed in Table 4.1, which contains radar beam widths and their corresponding minimum distances from the WLAN center for the radar beam to fully cover the 30 meter radius area.

In the system simulations run for this thesis, the horizontal radar beam width is 1.8 degrees. Figure 4.4 depicts the radar beam coverage over a 30 meter radius circular Wi-Fi network for a radar source only 60 meters away. As it is so close, the radar beam only covers approximately 4 percent of the area. From Table 4.1, a radar with a 1.8 degree beam width needs to be at least 1.91 km away from the center of the network to fully cover it with its beam. In the simulations run, the radar source to Wi-Fi network distance is always set to be greater than 1.91 km and the Wi-Fi network radius is always less than 30 meters to ensure that the entire Wi-Fi network can be covered by the interfering radar signals. As a result,

it is assumed that interfering radar signals hit all nodes in the Wi-Fi network at the same time with the same power (as the distance between Wi-Fi nodes is much smaller than the total distance that the radar signal propagates).

The preceding argument actually applies to both horizontal and vertical coverage of a radar beam over a Wi-Fi network, based on horizontal and vertical beam width respectively. But since vertical radar rotation is not modeled in this system simulator, it is assumed that the radar antenna has a vertical beam width and orientation such that its beam can always fully cover the Wi-Fi network in the vertical plane.

4.2 Path Loss

The wireless propagation loss model used in the system simulation stems from the basic free space path loss equation, in which the path loss is proportional to the square of the propagating signal's wavelength and inversely proportional to the square of the distance traveled. Since in reality the signal loss due to distance depends on the environment that the signal propagates through, the distance term can be generalized to have an exponent of n, where n would equal 2 to yield the free space path loss equation described above. Then, received power from one way propagation (from transmitter to receiver, not round trip back to the transmit source) can be expressed as:

$$P_r = P_t G_t G_r \left(\frac{\lambda}{4\pi}\right)^2 \frac{1}{d^n} \tag{4.11}$$

where P_t is the transmit power, G_t is the transmitting antenna gain, G_r is the receiving antenna gain, λ is the signal wavelength, d is the distance between the transmit and receive antennas (the distance that the signal propagates), and n is the path loss exponent.

Path loss exponents for different propagation environments are given in Table 4.2. In the system simulations performed in this research, it is assumed that the Wi-Fi network nodes are in buildings and do not have complete line-of-sight paths (1.6 to 1.8) to each other, but do not have completely obstructed paths (4 to 6) either. So the path loss exponent for Wi-Fi signals is set to 3. These buildings are assumed to be in urban areas, which interfering

Table 4.2: Path loss exponents for various propagation environments [19].

Environment	Path Loss Exponent, n
Free Space	2
Urban Area	2.7 to 3.5
Shadowed Urban Area	3 to 5
In Building Line-of-Sight	1.6 to 1.8
Obstructed In Building	4 to 6
Obstructed In Factories	2 to 3

radar signals must propagate through to get to the Wi-Fi networks. In the simulations we set the a radar a sufficient distance away such that its main beam covers the WLAN (from Section 4.1), so it is reasonable to assume that the radar source is outside of the urban area. Assuming the radar signal propagates through free space (2) for 3/4 of the distance to the WLAN and experiences moderate shadowing (4) due to buildings in the urban area for 1/4 of the distance, the average path loss exponent comes out to 2.5, which we use for radar signals.

4.3 ns-3 Wi-Fi Module

The ns-3 Wi-Fi module models the MAC and PHY layers for IEEE 802.11 [12]. The MAC model is broken up into a lower MAC and an upper MAC. The lower MAC model contains a DCF implementation, as described in Subsection 2.1.4. The upper MAC model includes scheduling of higher level functions such as beaconing and probing.

The Wi-Fi PHY model takes care of the physical transmission and reception of the Wi-Fi frames. At any given time, a node's PHY is in one particular state: either Transmit, Receive, Idle, CCA Busy, or Sleep. An incoming frame can only be received if the PHY is in the Idle state or in the CCA Busy state. If this criteria is met, the power of the frame signal is compared to the frame detection threshold. If it is higher than the threshold, the

PHY transitions into the Receive state and attempts to successfully receive the frame. If the frame signal power is lower than the detection threshold, it is dropped and the PHY remains in its state.

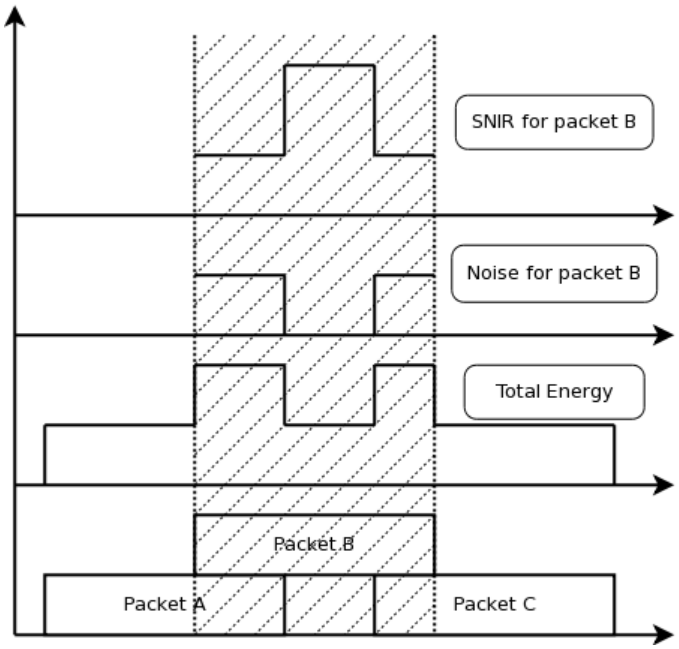

Figure 4.5: Example SNIR vs. time in ns-3 Wi-Fi PHY model. Retrieved from [12]

When a node is in the Receive state and attempting to receive a frame, the probability of successful reception is determined based on the Signal-to-Noise-Plus-Interference Ratio (SNIR). Thermal noise is the noise source modeled at the PHY layer, with the effective noise power given by:

$$P_n = kT_0BF \tag{4.12}$$

where k is Boltzmann's constant, T_0 is the noise reference temperature (290 K), B is the receiver bandwidth, and F is the receiver noise factor, which is the measure of SNR degradation that the receiver causes. Interference power from other frames in the channel also contributes to the SNIR. At any given time, the SNIR is calculated by dividing the power of the frame signal attempting to be received by the sum of the thermal noise power and

the inter-frame interference power. This is illustrated in Figure 4.5, where frame and packet are interchangeable terms in this context. Packet B is the frame that the PHY is trying to receive in this example.

The Wi-Fi module provides error rate models (similar to the FER results of Chapter 3) that map the SNIR, modulation and coding scheme, and number of bits to a receive success rate value. For each time segment of the frame being received that has a different SNIR than the previous segment, of which there are three in the example of Figure 4.5, a success rate value is determined from the error rate model. The packet success rate (PSR) for the whole frame is the product of all these time segment success rates. The frame error rate is then:

$$FER = 1 - PSR \tag{4.13}$$

Once the FER is determined, a random number between 0 and 1 is drawn from a uniform distribution. If this value is higher than the FER, the frame is received successfully. Otherwise, frame reception fails. Thus, it is much more likely for a frame with a low FER to be received successfully compared to one with a high FER.

For this work, some modifications and additions were made to the ns-3 Wi-Fi PHY model to model radar interference and its effects on Wi-Fi. ns-3.22 is the version of ns-3 that is used in this research.

4.4 Radar Interference Additions

Discussed in this section are the major changes and additions to the ns-3 Wi-Fi PHY model that enable radar interference modeling.

4.4.1 Radar Structure

The amount of radar interference power that a Wi-Fi network experiences changes over time due to radar rotation as well as radar pulse transmission time and rest time. A custom data structure is used to keep track of this, which contains radar signal power at the Wi-Fi node receivers (which is assumed to be the same for all nodes in the network based on the

assumptions of Section 4.1) over time after radar transmit antenna gain, path loss[1], and Wi-Fi receive antenna gain. To hold this information, the structure stores the time point and value of each power transition (positive or negative) for one full rotation (scan) cycle. These data points depend on a variety of parameters, including the radar's transmit power, distance to the Wi-Fi network, frequency, antenna radiation pattern, pulse duration, rest time, scan rate, as well as the Wi-Fi receive antenna gain.

In this model it is assumed that all the radar parameters, importantly scan rate and PRF, remain constant throughout the simulations. This is why data from only one full radar rotation cycle needs to be stored, because subsequent cycles will have the same data points, just shifted in time by the rotation period times the cycle number (where the first cycle is cycle 0, the second is cycle 1, etc.).

Table 4.3: Radar settings used for radar structure example.

Degrees per Rotation	1 degree
Pulses per Angle	1
Pulse Duration	1 μs
Rest Period	1999 μs
Distance from Radar to Wi-Fi	10 km
Path Loss Exponent	2.5
TX Radar Power	30 kW
Wi-Fi RX Gain	0 dB
Frequency	3.5 GHz
Antenna Pattern	in Figure 4.9

Figure 4.6 shows a simple example depiction of the radar structure for the settings in Table 4.3. It displays the received radar power as a function of time for one rotation cycle.

[1]The radar path loss is calculated using the distance from the radar source to the center of the Wi-Fi network—it doesn't account for the distance to each individual Wi-Fi node.

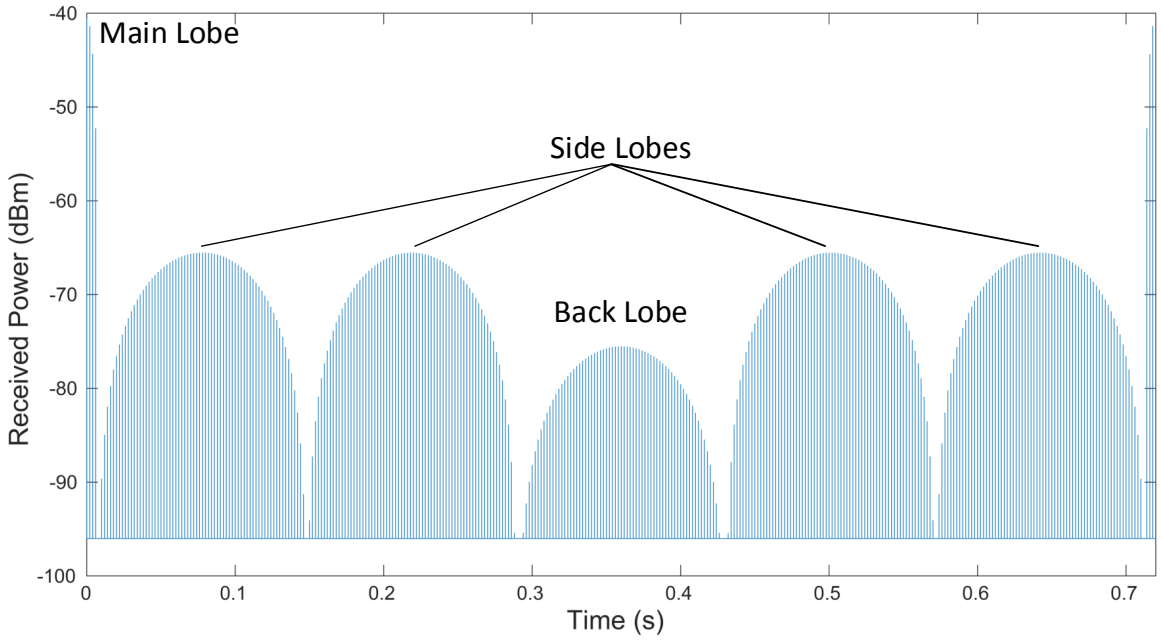

Figure 4.6: Depiction of received radar power at WLAN corresponding to implemented radar structure using example settings in Table 4.3 with -96 dBm noise floor. The radar has a 500 Hz PRF and 83.33 rpm scan rate. Data points are at each power transition (positive or negative). Data points are stored for one full rotation period, which is 720 ms in this example.

It shows a series of received radar pulses at different power values separated by rest periods of no received power. The data points (time and power change) that are stored correspond to each power transition, so for each pulse there are two data points—one at the time of the start of the pulse with a positive power value, and one at the time of the end of the pulse with a negative power value. The time difference between each pulse's end time and start time is the radar pulse duration, and the time difference between the start time of a pulse and the end time of the previous pulse is the radar rest time.

The varying power is due to the radar antenna rotation. The most interference power occurs when the radar antenna's main lobe (which has the highest gain) is pointed precisely in the direction of the Wi-Fi network. This is the highest peak received power in the cycle,

which is stored as the first data point in the radar structure. As the radar antenna rotates, the amount of power it radiates in the direction of the Wi-Fi network (and that the Wi-Fi nodes receive) changes due to the antenna gain. The lower peaks in Figure 4.6 correspond to the side lobes and back lobe of the radar antenna.

The discretized radar rotation described in Subsection 2.2.2 is used to generate the radar structure. From Eqs. (2.17) and (2.18), the scan rate is determined from the number of degrees per rotation, number of TX pulses per angle, pulse duration, and rest time. Recall from Eq. (2.16) that the pulse duration and rest time also determine the PRF. Using these equations with the settings in Table 4.3, the radar's PRF comes out to 500 Hz and the scan rate is 83.33 rpm (720 ms rotation period).

4.4.2 Radar Detection

When a Wi-Fi frame is being received, the simulation framework must be able to determine if a radar pulse interferes with it. Furthermore, if a radar pulse does interfere with the frame, the specific part of the frame that the radar pulse interferes with should be identified so that interference effects can be more accurately modeled by referencing the link sim error rate results from Chapter 3. Functionality was added to the simulator to allow for this capability.

When a receive frame event is occurring, the simulator knows the start and end time of the frame. Using these times, it checks the radar structure (described in Subsection 4.4.1) to see if any radar pulses are at the receiver in the frame time period. If not, no radar interference occurs and the simulator attempts to receive the frame with error probability solely based on its SNIR (from thermal noise and interference from other Wi-Fi frames). If a radar pulse does overlap with the frame over time, the time period overlap is identified and the interfering radar signal power is retrieved from the radar structure.

Once the interfering time period is resolved, the portion of the Wi-Fi frame that the radar pulse interferes with can be determined. Recall from the link sim discussion in Chapter 3 that Wi-Fi legacy mode frames (explained in Subsection 2.1.3) are being simulated in this

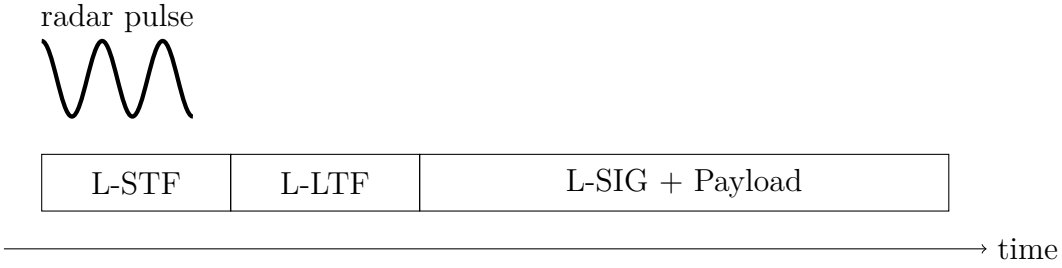

Figure 4.7: Radar pulse interfering with L-STF (frame synchronization portion) of Wi-Fi frame at receiver.

work (but with the SISO MCS and HT rates of Table 2.1). Legacy mode frames consist of four main portions: the L-STF, the L-LTF, the L-SIG, and the payload. Interference effects on the L-SIG were not modeled in the link sim, so the L-SIG and payload are grouped together as one sector for the system simulation interference modeling, and will be referred to as simply the payload for the remainder of this section. As shown in Figure 2.5, the L-STF, the part of the frame used for frame finding and time synchronization, is the first 8 μs of the frame. The L-LTF, the portion used for channel estimation, follows as the next 8 μs. The remaining part of the frame is then the payload (and L-SIG). From the knowledge of the time periods of each frame portion and the interfering radar pulse, the simulator identifies which of the three frame portions that the radar interference falls into. If the radar interference time overlaps with more than one frame portion, the one that it interferes with for the longest time is selected as the radar interference portion. Figure 4.7 illustrates a radar pulse interfering with the L-STF of the Wi-Fi frame being received.

4.4.3 Link Sim Error Rate Mapping

The link sim (of Chapter 3) frame synchronization error rate and frame error rate results are mapped for the system simulation. From the SNIR, radar INR, and radar hit location (portion of the Wi-Fi frame that the radar interferes with), the link sim error curves can be used to look up the error probabilities (using SNIR as the SNR input to the link sim

mapping). The frame sync modeled in the link sim uses the L-STF portion of the Wi-Fi frame, so if a radar hits the L-STF the frame sync error rate is looked up using both SNR and INR. To achieve the mapping, link sim result data is output into text files that the system simulator parses, storing the data in lookup tables.

The results used for the system simulations in this work come from link simulations with the following parameters: channel model B, MMSE estimator, soft decoding, and interleaving. Table 4.4 lists the link sim results that are mapped to the system simulator. The system simulations in this work were performed using only MCS 0 and MCS 7, which is why only results for these MCSs are listed.

Table 4.4: Link sim error rate result figures (from Chapter 3) that are mapped to the system simulator.

Figure Number	**Notes**
3.11	Frame sync for no radar hitting L-STF; Channel B curve only
3.12	Frame sync for radar hitting L-STF
3.10	Frame error for no radar hitting L-LTF or payload
3.13 & 3.14	Frame error for radar hitting L-LTF for MCS 0 & 7
3.15 & 3.15	Frame error for radar hitting payload for MCS 0 & 7

Bilinear interpolation is used to obtain error rate values in between SNR and INR data points. If SNR and/or INR values for error rate lookup do not align exactly with the SNR and/or INR values of the link sim data points, this technique determines the error rate value by approximating the data in between data points as changing linearly.

4.4.4 CCA and Frame Reception

The clear channel assessment (CCA) and frame receive logic in ns-3 was altered to address the requirements of this work. The flowchart in Figure 4.8 describes the new implementation. First, for any CCA or RX event to occur, the PHY state must be Idle or CCA Busy. Nodes

in Sleep, Transmit, or Receive states are already busy. Although the CCA Busy state means the node regards the channel as busy and will not transmit, it can still perform CCA to determine it to be busy longer, or even attempt to RX an incoming frame (though this is more likely to fail if the channel was already busy).

If the PHY is Idle or CCA Busy it monitors channel energy for incoming signals. 802.11 has a CCA-ED threshold of -62 dBm for 20 MHz channels (which we are simulating). If the aggregate channel energy including Wi-Fi and non-Wi-Fi sources is above -62 dBm, the Wi-Fi node must detect the channel as busy and defer transmission. Thus, interfering radar signals need to be included for CCA. To enable this in the simulator, an RX event is scheduled at each Wi-Fi node for each radar pulse in the simulation.

When an Wi-Fi or radar signal arrives at a Wi-Fi receiver, the receiver should only do anything if the signal power is above the RX sensitivity threshold. This threshold is typically close to the noise floor, which is the thermal noise power (Eq. (4.12)) plus the RX noise figure. For the 20 MHz channel and 5 dB noise figure simulated in this work, the noise floor is approximately -96 dBm. Although the RX sensitivity threshold is not actually implemented in the simulator, it has no effect in the simulations run for this work because all the RX Wi-Fi signal powers are well above -96 dBm at any node in each scenario.

After sensing a signal the node attempts to frame synchronize to it. Obviously, if it radar it will automatically fail because it is not a Wi-Fi frame. If it is a Wi-Fi signal, the probability of frame sync error is looked up from the error rate tables (of Subsection 4.4.3) using the RX SNR as well as the radar INR if any radar pulse interferes with the L-STF. A probabilistic coin toss is then done to determine whether or not frame sync was successful. If frame sync is not successful (due to the signal being a radar pulse or failing the coin toss), the aggregate energy in the channel is compared to the -62 dBm CCA-ED threshold and PHY state is set to CCA Busy if above it.

If frame sync is successful the node attempts to receive the whole frame. The SNIR is

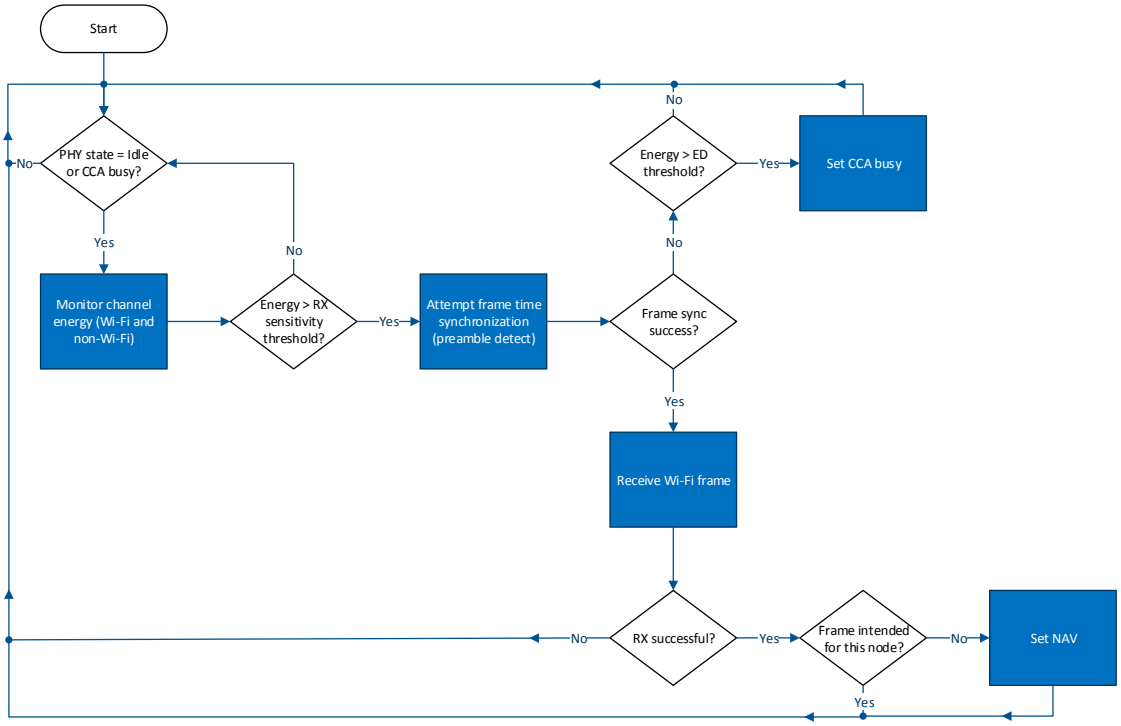

Figure 4.8: Wi-Fi clear channel assessment and reception flowchart.

calculated as:

$$SNIR\ (dB) = 10 \cdot log_{10}\left(\frac{P_s}{P_n + P_{i_{avg}}}\right) \tag{4.14}$$

where P_s is the power of the frame signal being received, P_n is the noise power as calculated from Eq. (4.12), and $P_{i_{avg}}$ is the average[2] inter-frame interference power during the receive

[2]The default ns-3 error rate model has bit length granularity, which enables it to calculate the error rate based on the SNIR of multiple separate time segments as depicted in Figure 4.5. The link sim results mapped from Chapter 3 do not have bit length resolution as they are frame error rates. This is why we must determine one SNIR to use for the error rate lookup. The average inter-frame interference power makes a fine assumption for the SNIR and INR calculations because in all the simulations run in this thesis, the scenarios are configured to have no hidden node problem. Thus, collisions can only occur when two nodes pick the same backoff counter and begin transmitting at nearly the same time. Since all data frames in the simulations are the same length, the colliding frames almost 100% overlap, making the average inter-frame interference power close to an optimal representation. The only other possible collision occurrence can happen if radar interference causes frame sync to fail so that a node fails to carrier sense a frame. It is then possible for that node to transmit in the middle of another node's transmission, but this event is rare.

frame duration, which is calculated as a weighted average over the duration of the RX frame. The radar INR is also calculated using the same average inter-frame interference power:

$$INR\ (dB) = 10 \cdot log_{10}\left(\frac{P_r}{P_n + P_{i_{avg}}}\right) \tag{4.15}$$

where P_r is the radar signal power at the Wi-Fi receiver.

After the SINR and INR (if applicable) are calculated, the frame error rate is looked up from the tables using the SINR, INR, and radar hit location (L-LTF or payload). Once retrieved, another probabilistic coin toss is done to determine frame RX success or failure. If it is successful and not intended for the receiving node, the NAV is set to activate virtual carrier sensing (explained in Subsection 2.1.4). A frame must pass both the frame sync and data RX coin tosses to be successfully received.

4.5 Simulation Parameters

Figure 4.9 displays the radar antenna radiation pattern used in the system simulations, generated from [13]. It provides the power gain [3] (in dBi) of the antenna as a function of horizontal azimuth angle (in degrees). The direction that the antenna is pointed in corresponds to 0 degrees. The antenna radiates the vast majority of its energy in the direction it is facing, as it contains the highest peak, which is termed the main lobe. Some energy is radiated in all directions, however. The four equal gain side peaks seen in the figure are called side lobes, while the small peak around 180 degrees is the back lobe. The (horizontal) beam width of the antenna is defined as the angle difference between the two points at half the peak gain of the main lobe. Half power corresponds to -3 dB, so these half peak gain points are each -3 dB below the peak main lobe point (0 degrees), one on each side. There is also a vertical beam width, but it is not specified in this simulation because vertical radar rotation is not modeled.

The beam width of the antenna is 1.8 degrees and the peak gain is 28 dBi, while the peak

[3] All antenna gains mentioned in this thesis are given in dBi, which is referenced to an ideal isotropic antenna. Ideal isotropic antennas radiate with unity gain in all directions.

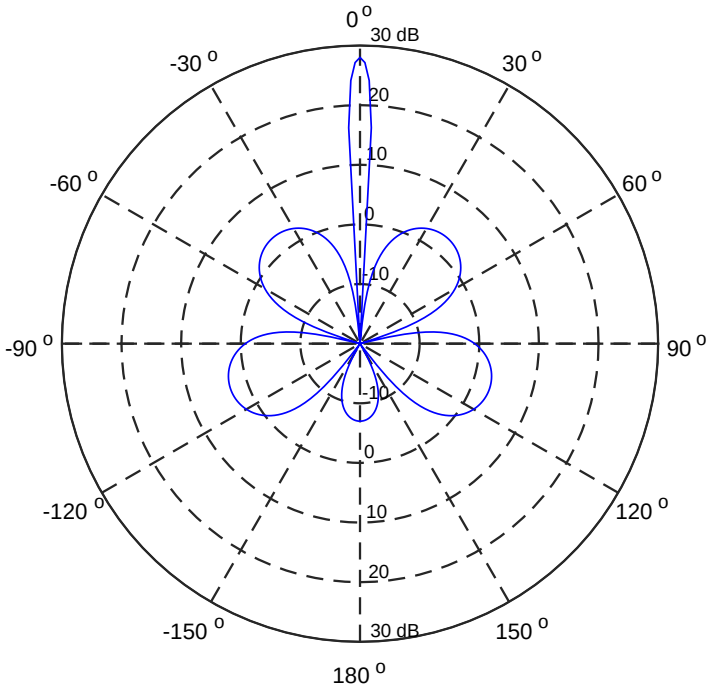

Figure 4.9: Radar antenna gain radiation pattern with beam width of 1.8 degrees, peak gain of 28 dBi, side lobe loss of -25 dB, and back lobe loss of -35 dB. Generated from [13] and plotted using [14]

transmit power used in the system simulations is 30 kW, all inspired by the United States navigational radars described in [38]. The side lobe loss is -25 dB and the back lobe loss is -35 dB, which are the gain reductions (from the main lobe peak) of the side lobe peaks and back lobe peak respectively. The gain data points from this antenna used in simulation are taken at 1 degree intervals. The radar settings used for all system simulations run in this work are listed in Table 4.5. Since the degrees per rotation is set to 0.5 degrees, half of the antenna gain data point interval, linear interpolation is used to calculate the gains between the data points.

Table 4.6 contains the Wi-Fi settings used in all the system simulations performed. Frame aggregation and fragmentation are not active. RTS/CTS is also inactive, as it is unnecessary since none of the simulated scenarios (detailed in Section 4.6) suffer from the hidden node

Table 4.5: Radar settings used for all system simulation results.

Peak Antenna Gain	28 dBi
Side Lobe Loss	-25 dB
Back Lobe Loss	-35 dB
Beam Width	1.8 degrees
Peak Transmit Power	30 kW
Path Loss Exponent	2.5
Pulse Duration	5 μs
Frequency	3.5 GHz
Degrees per Rotation	0.5 degrees

problem.

User Datagram Protocol (UDP) traffic is used for both upstream (packets from STA to AP) and downstream (packets from AP to STA) traffic, which is connectionless at the transport layer. A Poisson traffic model is implemented to manage the packet arrival rates. Packet arrivals follow a Poisson distribution, where the delay time between single packet arrivals is random according to an exponential distribution. The traffic rates used are listed in Section 4.6.

The number of data bits in each Wi-Fi frame depends on PHY parameters and is calculated in these simulations as:

$$N_{data\,bits/frame} = N_{sub} \cdot N_{OFDM\,syms} \cdot N_{bits/data\,sym} \cdot R_{code} - N_{SF\ bits} - N_{t\ bits} - N_{p\ bits} \quad (4.16)$$

where N_{sub} is the number of OFDM subcarriers used for data (excluding pilots)—48 in this case (from link sim), $N_{OFDM\,syms}$ is the number of OFDM symbols used for payload bits (not including the OFDM symbols of the preamble)—40 in this case (from link sim), $N_{bits/data\ sym}$ and R_{code} are the number of bits per data symbol and code rate respectively as per the selected MCS, $N_{SF\ bits}$ is the number of service field bits—16 in this case, $N_{t\ bits}$ is the

Table 4.6: Wi-Fi settings used for all system simulation results.

Standard	802.11n 5 GHz	OFDM Guard Interval	800 ns
Topology Mode	Infrastructure	LDPC	Not used
Frequency	3.5 GHz	CCA-ED Threshold	-62 dBm
Channel Bandwidth	20 MHz	Frame Fragmentation	None
Antenna Type	Isotropic	Frame Aggregation	None
Antenna Configuration	SISO	RTS/CTS	Inactive
TX Antenna Gain	0 dBi	Traffic Type	UDP
RX Antenna Gain	0 dBi	Traffic Model	Poisson
RX Noise Figure	5 dB	Num. Payload OFDM Syms.	40
Path Loss Exponent	3	Data Frame Duration	180 μs

number of tail bits—6 in this case, and $N_{p\ bits}$ is the number of pad bits—chosen to be 2 to make the control bits an even octet (16 service field bits + 6 tail bits + 2 pad bits = 24 control bits, which is divisible by 8). Eq. (4.16) was applied to calculate the number of data bits per frame required in each simulated scenario using the 40 OFDM symbol payload (and consequential 180 μs frame duration) that is configured for all the system simulations.

The mapped link sim error rate results were generated with the settings in Table 3.2 and the autocorrelation frame synchronization method described in 3.3.3. The remaining parameters vary throughout the different simulations performed and are listed in the remaining sections. Table 4.7 contains the ns-3 random variable seed, run, and stream values used in the simulations.

4.6 Scenarios

The selected simulation scenarios are discussed in this section. All STA node positions in each scenario are within a 30 meter circular radius of the single AP node in the center, which meets the beam coverage assumptions of Section 4.1, as the radar source in every performed

Table 4.7: Seed, run, and stream value assignments used for ns-3 random variables in all system simulation results.

Method	**Argument**
RngSeedManager::SetSeed	1
RngSeedManager::SetRun	2
InternetStackHelper::AssignStreams	4
OnOffHelper::AssignStreams [for AP]	200
OnOffHelper::AssignStreams [for STAs]	317
WifiHelper::AssignStreams	512

simulation is always at least 2 kilometers away from the Wi-Fi network. As previously mentioned, only the horizontal plane (X,Y) is modeled in this system simulation.

For each scenario the minimum Wi-Fi received SNR without interference was calculated by accounting for the transmit power and the two STA nodes farthest away from each other. This minimum receive signal power is higher than the 10 dB minimum required for a 100 percent frame sync success rate (from Figure 3.11, Channel B), ensuring that all nodes are capable of detecting frames from any of the other nodes. Thus, no hidden node problem exists in any of the scenarios.

The total traffic arrival rate averages were selected such that the channel should be close to, but not completely saturated with frames. The total upstream rate averages were chosen to be approximately one third of the total downstream rate averages to reflect how downstream traffic (from downloads, web browsing, etc.) is often more prevalent.

For each scenario number, version A uses MCS 0 and version B uses MCS 7, both with transmit power and packet data size (based on Eq. (4.16) for 40 payload OFDM symbols) adjusted accordingly. The other settings remain the same among versions A and B.

4.6.1 Small Home WLAN

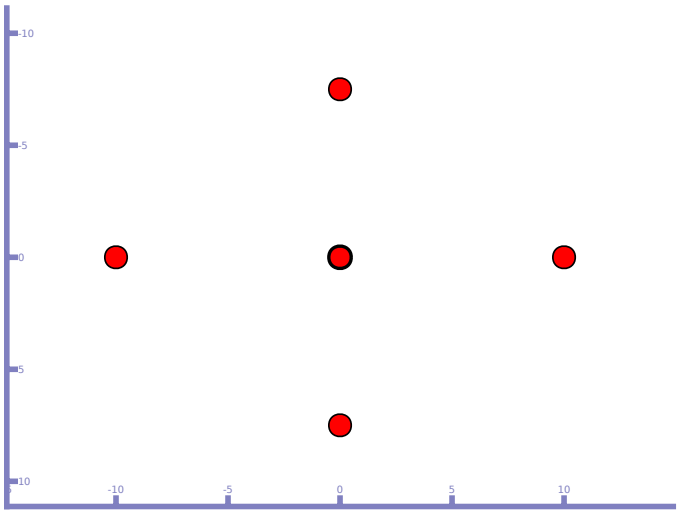

Figure 4.10: Node positions (in meters) for network scenario 1, small home. AP node is in center. Closest STA to AP distance is 7.5 m while farthest is 10 m. Depicted using [15]

Scenario 1 models a small home WLAN with 1 AP and 4 STAs. Its node positions can be seen in Figure 4.10, where the AP is in the center of the network. Two of the STAs are 10 meters away from the AP, while the other two are 7.5 meters away. Using a rectangular border aligned with the STA nodes, the area of the network is calculated as 300 meters squared. Transmit powers were chosen such that the received SNR range would yield frame error rates of approximately 0 (from Figure 3.10) for both MCS 0 and MCS 7 scenarios without inter-frame or radar interference. Table 4.8 contains parameters and properties of scenarios 1A and 1B.

Table 4.8: Network scenario 1 parameters and properties.

Parameter	**Scenario 1A**	**Scenario 1B**
Area	300 m^2	
Description	Small home	
Number of STA Nodes	4	
Number of AP Nodes	1	
Min STA to AP Distance	7.5 m	
Max STA to AP Distance	10 m	
Max STA to STA Distance	20 m	
Upstream Traffic Rate Mean	174 packets/sec per STA	
Upstream Traffic Rate Bound	17.4 packets/sec per STA	
Downstream Traffic Rate Mean	2083 packets/sec total	
Downstream Traffic Rate Bound	208.3 packets/sec total	
Wi-Fi MCS Mode	MCS 0	MCS 7
Packet Data Size	117 bytes	1197 bytes
Transmit Power (STA & AP)	2 dBm	22 dBm
Min Non-Interfered RX SNR	24.64 dB	44.64 dB
Max Non-Interfered RX SNR	28.39 dB	48.39 dB
Min Non-Interfered SNR at Any Node	15.61 dB	35.61 dB

4.6.2 Small Office Building WLAN

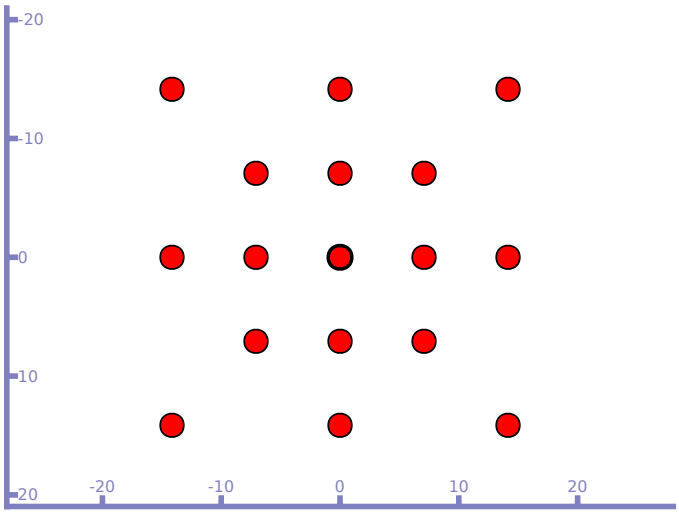

Figure 4.11: Node positions (in meters) for network scenario 2, small office building. AP node is in center. Closest STA to AP distance is 7.07 m while farthest is 20 m. Depicted using [15]

A small office building with 1 AP and 16 STAs is modeled in scenario 2. Figure 4.11 depicts the node positions with the AP in the center. One way the node positions can be interpreted is as a set of 9 cubicles in the center (with the middle one housing the AP) and 8 surrounding offices. The WLAN area is 800 meters squared when rectangularly bordering the outer nodes. The transmit powers were selected to make the received SNRs render a maximum frame error rate of approximately 0.005 and minimum frame error rate of approximately 0 (from Figure 3.10) for both MCS 0 and MCS 7 scenarios in absence of inter-frame or radar interference. Properties and settings of scenarios 2A and 2B are found in Table 4.9.

Table 4.9: Network scenario 2 parameters and properties.

Parameter	**Scenario 2A**	**Scenario 2B**
Area	800 m^2	
Description	Small office building	
Number of STA Nodes	16	
Number of AP Nodes	1	
Min STA to AP Distance	7.07 m	
Max STA to AP Distance	20 m	
Max STA to STA Distance	40 m	
Upstream Traffic Rate Mean	43.4 packets/sec per STA	
Upstream Traffic Rate Bound	4.34 packets/sec per STA	
Downstream Traffic Rate Mean	2083 packets/sec total	
Downstream Traffic Rate Bound	208.3 packets/sec total	
Wi-Fi MCS Mode	MCS 0	MCS 7
Packet Data Size	117 bytes	1197 bytes
Transmit Power (STA & AP)	5.5 dBm	25.4 dBm
Min Non-Interfered RX SNR	19.11 dB	39.01 dB
Max Non-Interfered RX SNR	32.65 dB	52.55 dB
Min Non-Interfered SNR at Any Node	10.08 dB	29.98 dB

4.7 Results and Analysis

Select system simulation results are shown and analyzed in this section. For the full set of collected results, see Appendix A. The radar and Wi-Fi settings in Tables 4.5 and 4.6 respectively are used in all of the simulations.

The distances from the radar source to the center of the Wi-Fi network chosen for simulation are 50, 20, 10, 5, and 2 km. These values are all above the required 1.91 km to achieve full radar coverage of the Wi-Fi network for a 1.8 degree radar beam width and 30

meter circular network area according to Table 4.1. Since all Wi-Fi nodes in both simulation scenarios are within a 30 meter radius, the assumptions of Section 4.1 hold.

Table 4.10: Radar simulation variable settings.

		Option			
		1	**2**	**3**	**4**
Input	rest period	195 μs		995 μs	
	pulses per angle	15	30	3	6
Outcome	PRF	5 kHz		1 kHz	
	scan rate	27.78 rpm	13.89 rpm	27.78 rpm	13.89 rpm

Table 4.10 contains the simulation variables and their four configuration options in these simulations. Setting the rest period in turn sets the PRF, from Eqs. (2.15) and (2.16). With the configured 5 μs radar pulse duration, the PRF comes out to either 5 kHz or 1 kHz from the two rest period options.

The upstream, downstream, and total average Wi-Fi throughput for scenarios 1 and 2 with no radar interference are listed in Table 4.11. These are used as benchmarks to compare with the throughput for radar interference scenarios. As expected, the upstream throughputs are around one third of the downstream throughputs.

Two types of plots were generated as results for each scenario configuration with radar interference. For the first, percentage of no radar average throughput versus radar distance, each simulation was run for one radar rotation period with the throughput averaged over the single period. This plot type compares the average single radar rotation upstream, downstream, and total throughput to the corresponding total throughput in Table 4.11 (row 3). The second plot type contains total throughput (upstream and downstream) versus time, where each simulation was run for 15 seconds (simulation time, not actual time), with 100 milliseconds between each data point and each throughput measurement averaged over the previous 100 milliseconds.

Table 4.11: Average throughput over 15 seconds for scenarios 1 and 2 with no radar interference. Total throughputs are comparison benchmarks for each corresponding radar interference configuration.

	Scenario 1		Scenario 2	
	MCS 0	**MCS 7**	**MCS 0**	**MCS 7**
Upstream	0.642 Mbps	6.36 Mbps	0.646 Mbps	6.55 Mbps
Downstream	1.93 Mbps	18.1 Mbps	1.94 Mbps	19.4 Mbps
Total	2.57 Mbps	24.5 Mbps	2.59 Mbps	26.0 Mbps

In all system simulations run, the throughput measurements begin (time = 0 in the time plots) after 9 seconds of Wi-Fi simulation time to allow the system throughput to stabilize. The radar interferer begins transmitting one full rotation period before measurements begin, so that the center of its main antenna lobe (0 degrees in Figure 4.9) is directed toward the Wi-Fi network at measurement start.

4.7.1 PRF Impact

As a radar interferer gets closer in distance to a Wi-Fi network, its signals experience less path loss before reaching the WLAN, causing the Wi-Fi nodes to experience higher INR and thus reduced communication performance.

Figure 4.12 displays the percentage of total average Wi-Fi throughput of one radar rotation period with respect to the no radar interference case versus radar to Wi-Fi distance (hereinafter *distance plot*) for scenario 1A with a 1 kHz PRF and 13.89 rpm scan rate radar. It is shown that the radar interferer does not begin to noticeably degrade the total average throughput until it is between 10 and 5 km from the WLAN.

Figure 4.13 shows the distance plot for the same scenario and configuration besides an increase in radar PRF from 1 kHz to 5 kHz. The total average throughput begins to noticeably drop between 20 and 10 km, and falls off much sharper than the 1 kHz PRF configuration.

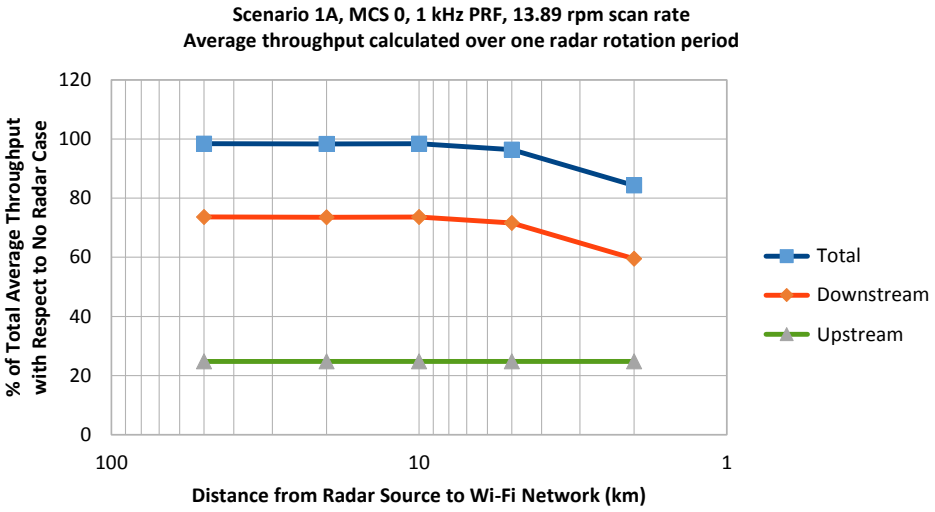

Figure 4.12: Percentage of total average throughput relative to no radar interference case for scenario 1A option 4.

At 2 km radar distance, the 5 kHz configuration throughput is close to 20 percent of the total no radar throughput, while the 1 kHz one is still around 80 percent.

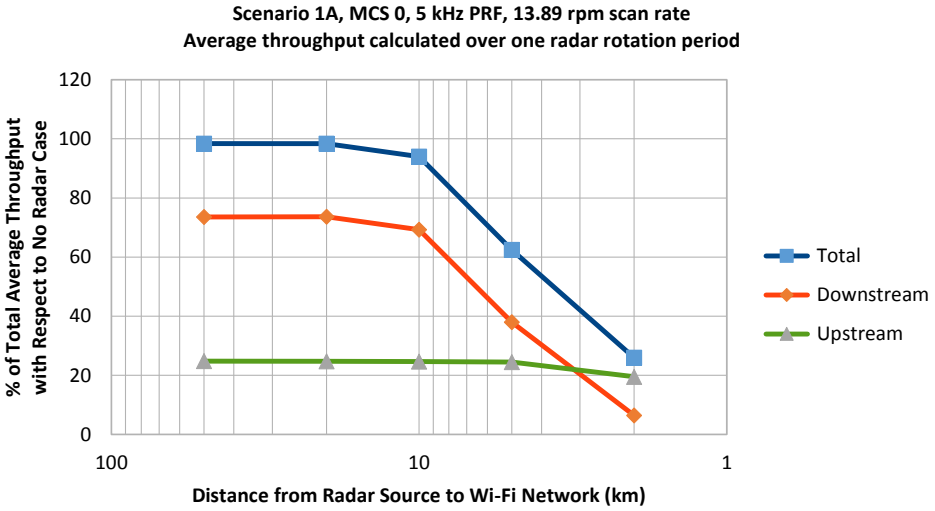

Figure 4.13: Percentage of total average throughput relative to no radar interference case for scenario 1A option 2.

The reason why the 1 kHz PRF configuration performs much better than the 5 kHz one is that its radar pulse transmissions are much less frequent. At 1 kHz PRF, the rest period is 995 μs while the pulse duration is only 5 μs. This means that the pulse makes up only 0.5 percent of the PRI. Since the radar pulses are only occupying the channel 0.5 percent of the time, it is relatively uncommon for a radar pulse to interfere with a Wi-Fi frame and cause communication failure. On the other hand, at 5 kHz PRF, the rest period is 195 μs (still with 5 μs pulse duration), causing radar pulse transmissions to be five times more frequent and comprise 2.5 percent of the PRI. As the PRI is 200 μs and the Wi-Fi data frame duration is 180 μs, most data frames are interfered with by a radar pulse. This causes many more frames to be dropped due to radar interference than the 1 kHz PRF case, yielding more throughput degradation.

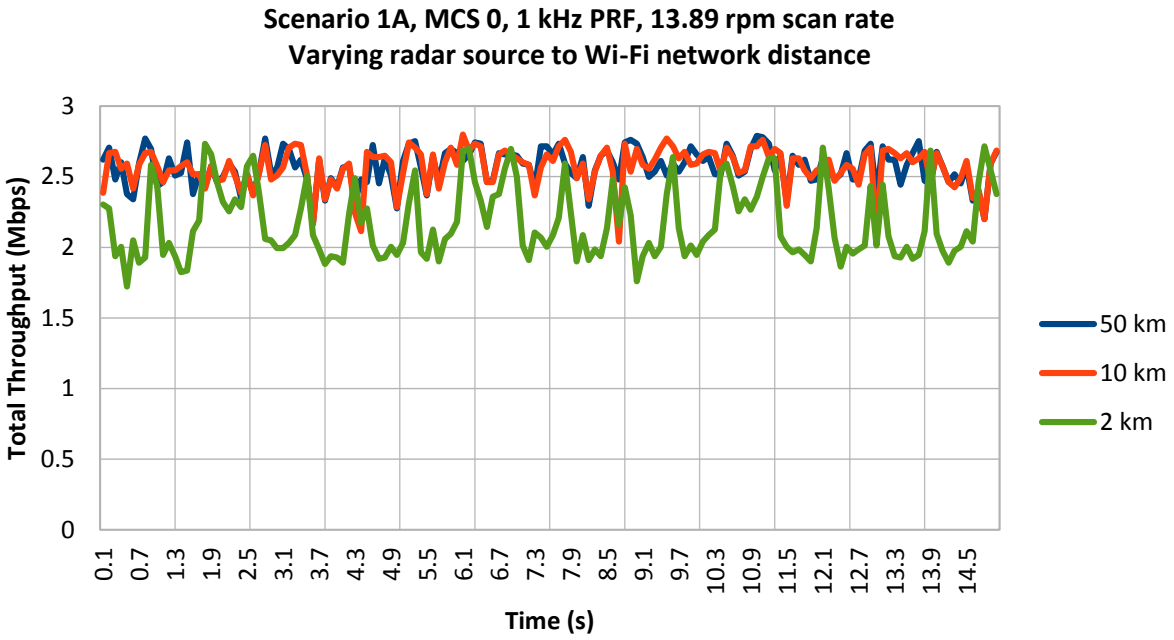

Figure 4.14: Total throughput over time for scenario 1A option 4. Each curve has one data point every 100 ms with each throughput measurement averaged over the previous 100 ms.

Figures 4.14 and 4.15 depict the total throughput (upstream and downstream) over time

(hereinafter *time plot*) for the same two configurations for three different radar distances. It can be seen that *even at 2 km radar distance, the total throughput does not greatly drop over time for the 1 kHz PRF scenario, since radar pulse transmissions are not very frequent. But for 5 kHz PRF, the performance impact of the rotating radar interferer is clearly visible.*

4.7.2 Throughput Dips Over Time

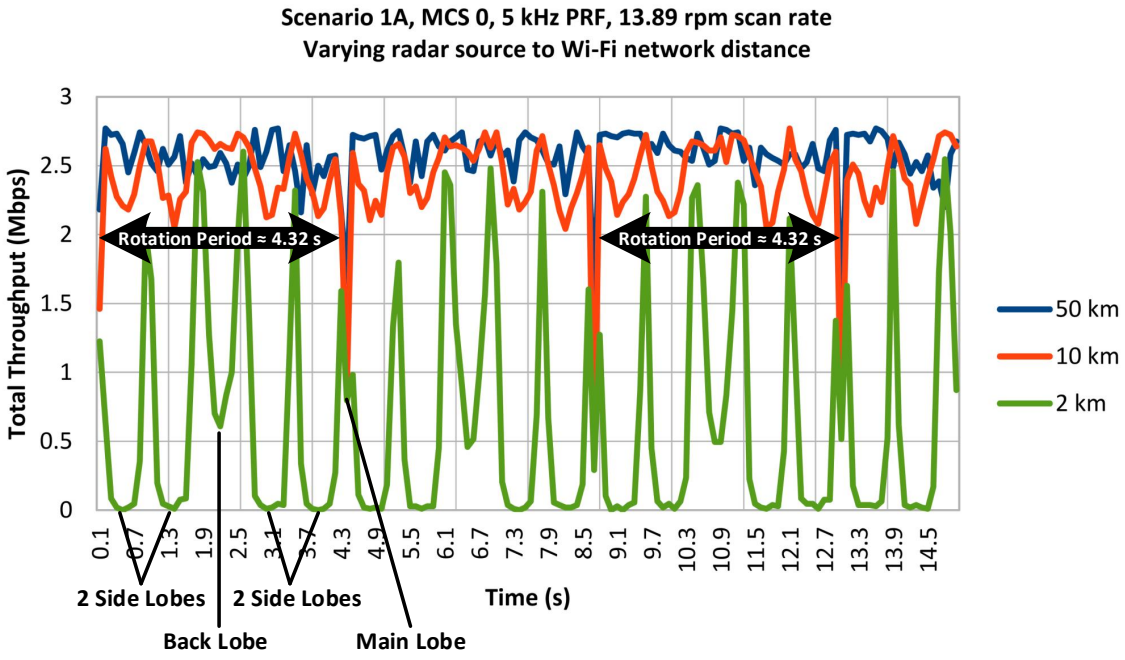

Figure 4.15: Total throughput over time for scenario 1A option 2. Each curve has one data point every 100 ms with each throughput measurement averaged over the previous 100 ms. Contains annotations relating the throughput curves to the radar scan rate and antenna pattern.

In the 50 and 10 km curves of the 5 kHz PRF time plot of Figure 4.15, it is shown that *at farther radar distances the primary throughput dip corresponds to the main lobe of the radar antenna—when the radar is pointing directly at the Wi-Fi network.* The main lobe has the highest gain, and thus the WLAN experiences the most interference when the main

lobe is oriented toward it. These throughput troughs are periodic, with a repetition interval corresponding to the radar rotation period of approximately 4.32 seconds (60 / 13.89 rpm). As the radar main lobe has a 1.8 degree beam width, only during a 22 ms period (4.32 s * (1.8 deg / 360 deg)) can the WLAN be interfered with by radar pulses from the main beam. For this reason, *the main lobe throughput dips only take up a single data point, since the time interval between throughput measurements is 100 ms, longer than the "dwell time" of the main lobe.*

At the farther distances, interfering radar pulses transmitted from the side and back antenna lobes are not strong enough to make a heavy impact on throughput. But, as shown for 2 km in Figure 4.15, *at closer distances the side lobes and back lobe contribute heavily to the throughput degradation over time and significantly decrease average throughput.* A repeating throughput dip pattern can be observed over time. Starting at 0 seconds (main lobe oriented at WLAN at time = 0 s) in the 2 km curve, the first two throughput dips correspond to the first two side lobes as the radar rotates, both with a nearly equal effect. The third dip corresponds to the back lobe and is smaller in intensity, since it has a lower gain than the side lobes (-7 dB instead of 3 dB). The third and fourth dips correspond to the last two side lobes before the radar completes its rotation. A main lobe dip can be seen following this, but it is less significant because the main lobe spans many less degrees than the side and back lobes so its signals interfere with the WLAN for a shorter time. *The throughput dip pattern of two side lobes, back lobe, two side lobes, and main lobe periodically repeats with the radar rotation period as the repetition time interval.* The strong impact of the side and back lobes at close distances is why the total average throughput for 2 km is only about 25 percent of the no radar case, while it is near 100 percent at 50 km (from Figure 4.13) .

4.7.3 MCS Impact

Figures 4.16 and 4.17 are the distance plots for scenario 1 with a 5 kHz PRF and 27.78 rpm radar for MCS 0 and MCS 7 respectively. *Compared to the MCS 0 scenario, the MCS 7*

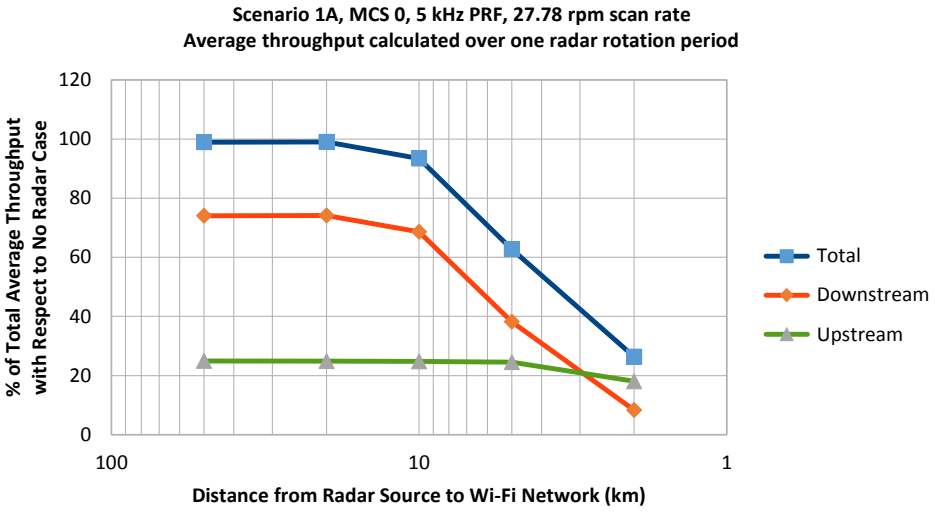

Figure 4.16: Percentage of total average throughput relative to no radar interference case for scenario 1A option 1.

scenario experiences considerably more percent throughput degradation from the radar interferer and begins to show noticeable drops in average throughput at a farther radar distance.

The total, downstream, and upstream percent average throughput curves all drop more sharply with distance for MCS 7. For MCS 7 the total throughput begins to heavily decline between 50 and 20 km, while for MCS 0 it doesn't noticeably fall until between 20 and 10 km, and even then the decline is much less steep. The total throughput of the MCS 7 scenario drops all the way down to approximately 10 percent at 2 km (with nearly 0 percent downstream throughput), whereas for MCS 0 it is around 25 percent.

MCS 0 having better resistance to radar interference than MCS 7 makes intuitive sense, because MCS 7 has a higher modulation order (64-QAM vs. BPSK) and code rate (5/6 vs. 1/2). This is also reflected in the radar interference results of Chapter 3.

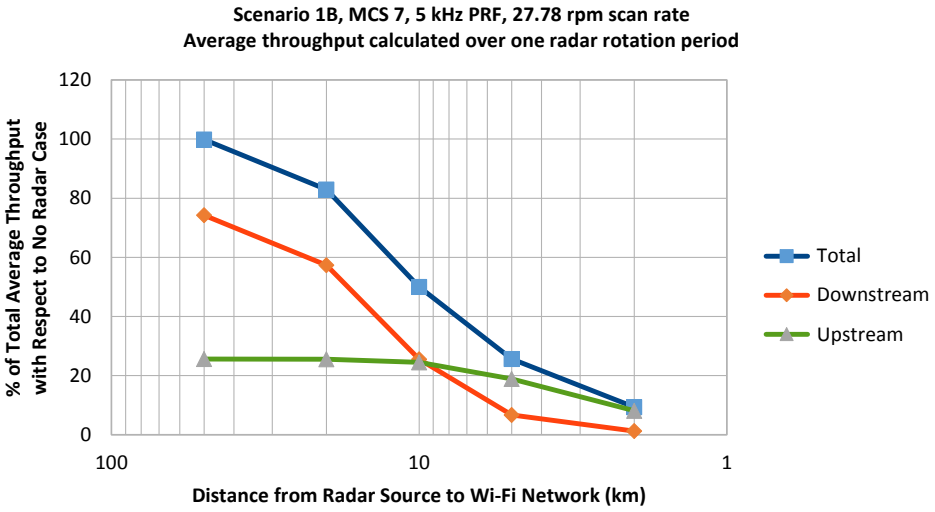

Figure 4.17: Percentage of total average throughput relative to no radar interference case for scenario 1B option 1.

4.7.4 Scenario 1 and 2 Comparison

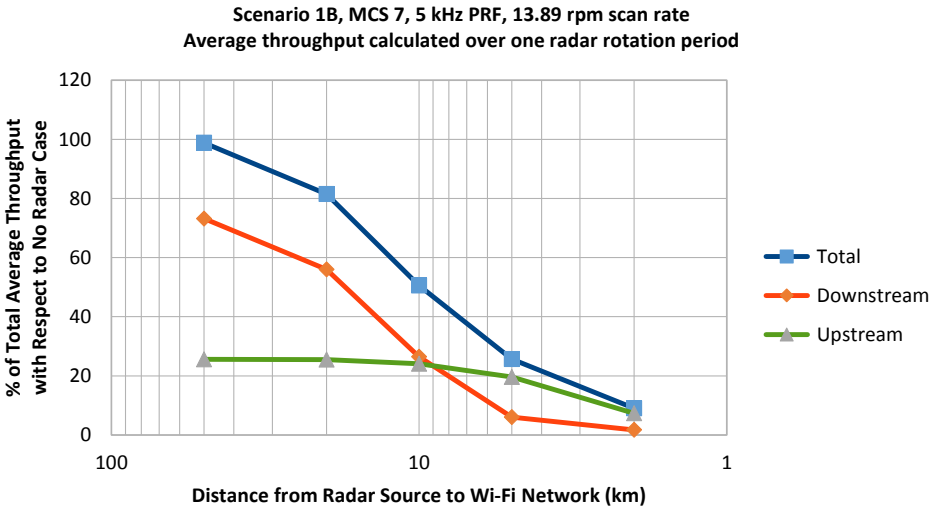

Figure 4.18: Percentage of total average throughput relative to no radar interference case for scenario 1B option 2.

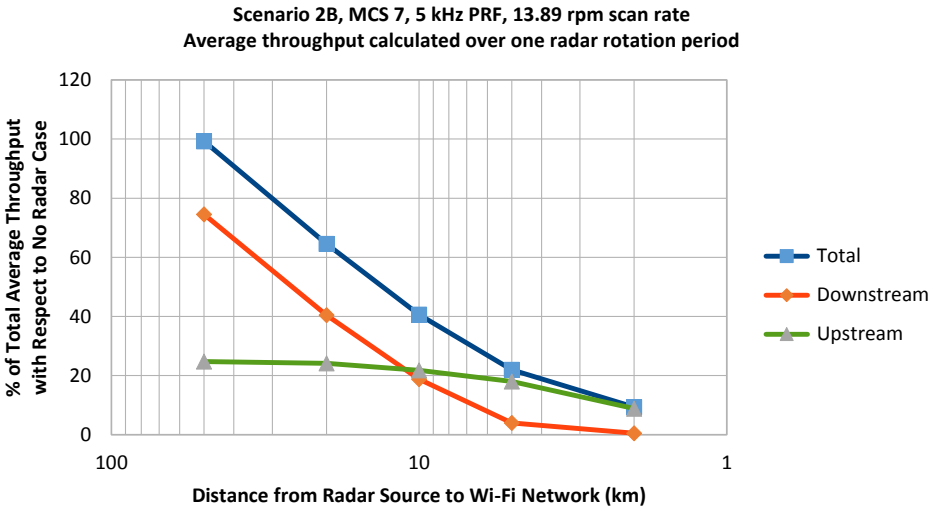

Figure 4.19: Percentage of total average throughput relative to no radar interference case for scenario 2B option 2.

Distance plots are given in Figures 4.18 and 4.19 with 5 kHz PRF and 13.89 rpm scan rate for scenarios 1B and 2B respectively. The average throughput in scenario 2B declines faster than in scenario 1B because 2B has farther away STA nodes from its AP, with a minimum non-interfered RX SNR of 39.01 dB while the minimum for 1B is 44.64 dB (from Tables 4.8 and 4.9). A lower SNR is more susceptible to frame errors from radar interference (as shown in the Chapter 3 radar interference results), and thus scenario 2B experiences sharper throughput degradation.

The higher number of STA nodes in scenario 2 over scenario 1, 16 instead of 4, does not seem to cause a noticeable impact on throughput performance. The total upstream (approximately 700 packets per second) and downstream (2083 packets per second) traffic arrival rate mean is the same for both scenarios, meaning although scenario 2 has more STAs, each STA is less likely to have a packet to transmit than the STAs in scenario 1. If instead both scenarios were configured to have the same "per STA" upstream arrival rate, scenario 2's total upstream throughput would be higher and more frame collisions would likely occur.

4.7.5 Scan Rate Impact

Figure 4.20 is the distance plot for the same scenario 2B and 5 kHz PRF as Figure 4.19 but with double the scan rate, at 27.78 rpm instead of 13.89 rpm. Between the two, *the average throughput performance over a radar rotation period is nearly identical for the varying radar scan rate.*

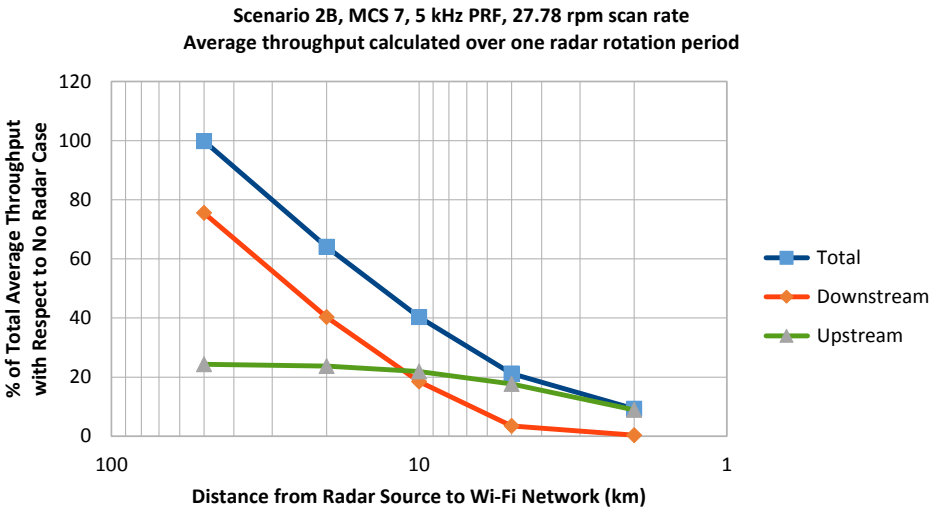

Figure 4.20: Percentage of total average throughput relative to no radar interference case for scenario 2B option 1.

However, Figures 4.21 and 4.22 show the time plots for the same to scenarios, for 13.89 and 27.78 rpm respectively. Differences between these two plots are visible. The throughput dips corresponding to the side lobes and back lobe are visibly about twice as long for 13.89 rpm compared to 27.78 rpm. This makes sense because the 13.89 rpm radar spends twice as long in each antenna direction than the 27.78 rpm one. The same phenomenon occurs with the main lobe troughs at 50 km radar distance. But since the throughput measurement interval is longer than the main lobe dwell time in both scan rate cases, it appears as a lower throughput dip in the 13.89 rpm scenario. This is because the 22 ms main lobe duration (duration of time that the 1.8 degree beam width main beam is facing the Wi-Fi network)

for 13.89 rpm is twice as long as the 11 ms for 27.78 rpm, causing the throughput data point over the 100 ms measurement interval to be lower for the former.

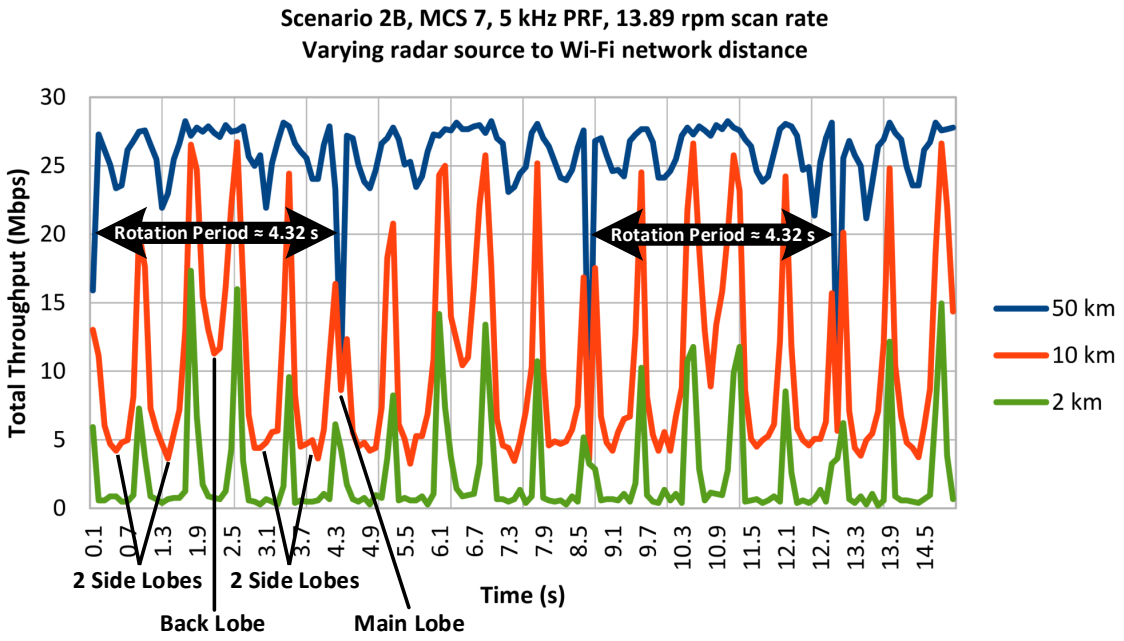

Figure 4.21: Total throughput over time for scenario 2B option 2. Each curve has one data point every 100 ms with each throughput measurement averaged over the previous 100 ms. Contains annotations relating the throughput curves to the radar scan rate and antenna pattern.

Since scan rate does not seem to significantly impact overall average throughput, it may not have much correlation with user-experienced performance of non time-sensitive applications such as email and file downloads. But for time-sensitive applications such as video chat and VoWLAN, user experience can suffer dramatically from throughput dips in time. *Time-sensitive applications will likely perform worse with lower radar scan rate because throughput dips will last longer, causing longer interruptions to the application. But a higher scan rate radar will cause more frequent shorter interruptions, which can also be detrimental to performance.*

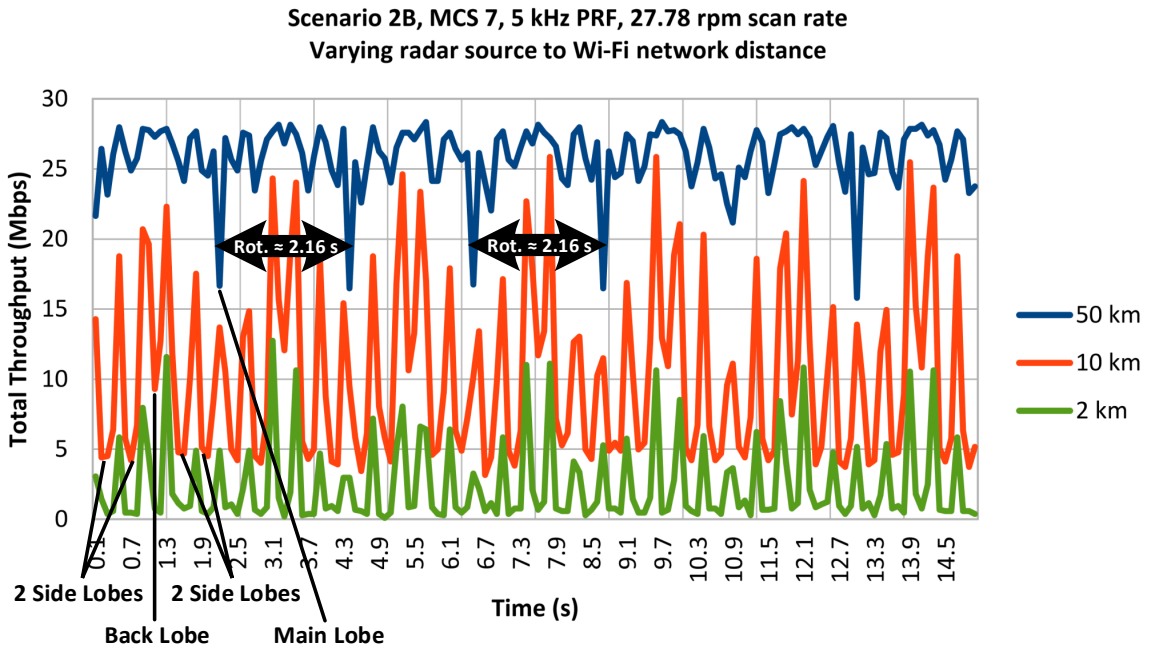

Figure 4.22: Total throughput over time for scenario 2B option 1. Each curve has one data point every 100 ms with each throughput measurement averaged over the previous 100 ms. Contains annotations relating the throughput curves to the radar scan rate and antenna pattern. In the figure, "Rot." stands for rotation period.

4.7.6 Upstream Resistance to Throughput Degradation

In all simulated scenarios and configurations, upstream throughput did not degrade much as a function of radar distance to WLAN. There are two main reasons for this: the lower upstream arrival rate and total upstream throughput's "resistance" to random backoff.

In all scenarios, the average total upstream arrival rate was configured to be close to a third of that of the downstream. Hence, fewer upstream data frames are being transmitted than downstream ones in these simulations. This makes it less likely for radar pulses to interfere with upstream data frames—contributing to the apparent throughput preservation.

Additionally, *in WLANs with multiple STAs and a single AP, total upstream traffic is more resilient to random backoff than total downstream traffic.* The "total" distinction is

important here, as it denotes the sum throughput of all STA to AP flows. If a radar pulse interferes with a downstream data frame or its ACK and causes a missed ACK, the AP must double its CW and enter random backoff. The same occurs with upstream traffic when a STA has a missed ACK. However, since only one STA actively transmits at a time (barring rare collision events), when a radar pulse hits the data frame and causes it to fail, only the transmitting STA doubles its CW and enters backoff. Other STA nodes are not forced into the increased backoff and can still transmit if they have a packet ready. Therefore, having multiple STA nodes helps to maintain total upstream throughput.

Consider the situation where a radar's main lobe is directed toward a WLAN and each peak interfering radar pulse is causing a Wi-Fi frame to fail. It is likely that a pulse will interfere with at least one downstream frame or its ACK, forcing the AP into extended backoff. It is less likely that the pulse will interfere with the upstream transmission of every single STA in the network, as main lobe dwell time is typically not very long since radar beam widths are extremely small. So by the time the main beam passes the WLAN and less interference is occurring, at least one STA that wasn't interfered with by the main beam will likely be ready to transmit and be successful.

All downstream data transmission is halted if the single AP is in random backoff, which explains the heavy downstream throughput degradation from radar interference seen in the simulation results.

4.7.7 Overview

A summary of the observed effects of system variables on average Wi-Fi throughput is contained in Table 4.12. The radar PRF and distance to the WLAN had the most impact on average throughput performance. The Wi-FI MCS index is also relevant, as higher MCSs experienced more percent average throughput reduction than lower ones. Radar scan rate did not appear to have a significant effect on average throughput, but it did shape the throughput-time relationship, as lower scan rates yielded longer throughput dips than higher ones.

Table 4.12: Observed impact of changing system variables on average Wi-Fi throughput.

Parameter	**Direction of change proportional to avg. throughput degradation**	**Impact on % avg. throughput performance**
Radar PRF	increasing	high
Radar scan rate	no clear relation	very low
Radar to Wi-Fi distance	decreasing	high
Wi-Fi MCS index	increasing	medium

Chapter 5
MITIGATION SCHEMES

The link and system simulation results of chapters 3 and 4 have shown that radar interference often makes a large negative impact on Wi-Fi network performance. In response, several mitigation schemes are proposed in this chapter. They are attempts to reduce the Wi-Fi performance degradation incurred from radar interference.

Some of the proposed mitigation schemes assume that Wi-Fi nodes are able to detect radar pulses. This paper does not cover radar detection methods, but the energy detection technique detailed in [39] and the bit error pattern identification explored in [24] are possible. A mitigation scheme may also assume that Wi-Fi nodes can estimate some properties of an interfering radar system, such as pulse duration, PRF, scan rate, and signal power. Any dependencies on these estimations are mentioned in the mitigation scheme descriptions below.

5.1 Contention Window Preservation

Recall from the Wi-Fi MAC layer DCF (Subsection 2.1.4) that after a node transmits a data frame, if a corresponding ACK frame is not received within the ACK timeout period (or if it transmits a RTS frame and the CTS frame is not received within the CTS timeout period), the CW value is doubled and the random backoff procedure is run with a random number from 0 to CW. This mechanism is present in DCF to reduce the probability that nodes draw the same random backoff number and thus transmit at the same time and collide. If an ACK (CTS) is missed, it is likely that a collision occurred with another node's frame, so doubling the CW will make a collision less likely for the next transmission. However, doubling the CW also makes it more likely for nodes to have to wait longer before transmitting, since larger backoff numbers can be drawn. So if the CW is unnecessarily high, less than optimal

throughput is likely.

A missed ACK (CTS) will occur if an interfering radar pulse causes a data frame (RTS) or its ACK (CTS) to be received incorrectly. If this takes place, the CW will be doubled as per the DCF. But radar interference has no correlation with channel contention from other Wi-Fi nodes. Hence, doubling the CW does not do anything to address the communication failure that was caused by radar interference but still leads the node with the missed ACK (CTS) to be more likely to have to wait longer before transmitting again. This is an unnecessary phenomenon that can reduce system throughput. To address this issue, the author has crafted the *Contention Window Preservation* mitigation scheme, which stops the DCF from doubling the CW in select circumstances.

Algorithm 1 Contention Window Preservation

```
1: if missed ACK or missed CTS then
2:     if radar pulse was detected while RXing during expected ACK or CTS time and RX
       power was greater than the threshold then
3:         maintain CW value
4:     else
5:         double CW value
6:     end if
7:     commence random backoff
8: end if
```

The Contention Window (CW) Preservation scheme assumes that Wi-Fi nodes can identify when an interfering radar pulse hits and estimate the radar signal power. If a missed ACK (CTS) event occurs, the node checks if frame reception was attempted (the node was trying to receive a in incoming frame) during the expected ACK (CTS) time interval. The start of this interval should be one SIFS duration after the end transmission time of the data frame (RTS) corresponding to the missed ACK (CTS), since data frames (RTSs) and ACKs (CTSs) are separated by a SIFS period. The end of the expected ACK (CTS) time

interval should be the interval start time plus the ACK (CTS) frame duration. If frame reception was attempted within this interval and an interfering radar pulse above a specified power threshold was also detected during the interval, the CW value is retained instead of being doubled. The random backoff procedure is then run. Evidently, the CW Preservation scheme requires Wi-Fi nodes to record and keep track of frame receive attempts within expected ACK (CTS) time intervals during which radar interference is detected. The CW retention mechanism would also be helpful if it could be triggered from a radar pulse above the power threshold being detected during data frame (RTS) transmission, but it is not included in the proposed algorithm because simultaneous transmission and reception/detection capability is uncommon.

The author envisions the aforementioned radar power threshold to be either manually configurable or automatically updated from empirical measurements. This threshold can affect the performance gain yielded from the CW Preservation scheme. If the threshold is set too low, it could cause the CW value to be maintained in situations where the interfering radar pulse was not at a high enough power level to cause ACK (CTS) reception error, but where a collision from another frame occurred. Since there was a collision from another node's frame, it is desirable to double the CW to reduce the chance of a collision occurring again. But the CW Preservation scheme prevents this from happening because the radar interference was above the power threshold. In contrast, if the radar power threshold is set too high, the CW retention mechanism may not trigger in circumstances where it would be beneficial.

A Wi-Fi node using the CW Preservation scheme would still be able to communicate successfully with legacy nodes, as no essential operations are altered.

5.2 Concurrent Frame-Radar Avoidance

If a Wi-Fi node can predict when interfering radar pulses will arrive, it can avoid transmitting frames during times of radar interference. This is main idea behind the *Concurrent Frame-Radar Avoidance* mitigation scheme that the author has devised. This scheme as-

sumes that Wi-Fi nodes can accurately estimate radar pulse duration, PRF, scan rate, and signal power. If a node has radar detection capability such that it can measure pulse duration and power with sufficient precision, it could do some fairly straightforward processing to calculate PRF and scan rate estimates. For example, after taking a few measurements of the radar pulse duration and time length between pulse arrivals, the node could estimate the PRI and thus the PRF.

Algorithm 2 Concurrent Frame-Radar Avoidance

```
1: if ready to send data frame or RTS frame then
2:    if frame duration predicted to overlap with radar pulse above power threshold then
3:       enter random backoff
4:    else
5:       transmit frame
6:    end if
7: end if
```

The Concurrent Frame-Radar Avoidance scheme applies when a Wi-Fi node is ready to send a data frame or RTS frame. Assuming the node knows the time duration of the frame it wants to transmit, it checks to see if the transmission period for the frame will overlap in time with an interfering radar pulse based on its PRF and scan rate predictions. If it is determined that overlap will occur and the predicted radar pulse power is higher than the power threshold, the node runs the random backoff procedure instead of transmitting. By doing so, the node avoids the impending radar interference at the receiver of the destination node[1], which could cause the frame to be received unsuccessfully. Instead of potentially having to wait for the ACK or CTS to time out and doubling the CW before going into random backoff, the node just enters random backoff right away (keeping the same CW

[1]Since a single Wi-Fi network typically does not cover a very wide area, it can be reasonably assumed that radar pulses will interfere with all the nodes in a network at close to the same time and power level. Therefore, it is assumed that if a radar pulse interferes with a sender node during transmission, it also interferes with the destination node during reception at a similar interference power level.

value). This can increase throughput by reducing the wasted time caused by a missed ACK or CTS.

The radar power threshold level for the avoidance mechanism to trigger has similar trade-offs to the one for the Contention Window Preservation scheme (discussed in Section 5.1). If the threshold is set too low, the avoidance mechanism may be triggered too frequently, such that transmission is deferred even when the imminent radar interference power is so low that it is extremely unlikely to cause frame receive error. On the other hand, an exceedingly high power threshold can cause the avoidance mechanism not to trigger in situations where it would be advantageous.

The proposed algorithm for Concurrent Frame-Radar Avoidance is only applied to data and RTS frames because deferring ACK and CTS frames can easily lead the sender node to reach its ACK or CTS timeout. Using the avoidance mechanism for ACK and CTS frames could cause timeouts to happen frequently, so in the author's opinion it would often be best to just transmit despite impending radar interference, since there is still a chance that the ACK or CTS frame could be received successfully by the sender node.

Like the CW Preservation scheme, nodes employing the Concurrent Frame-Radar Avoidance scheme will still be compatible with others that are not.

5.3 Time Interleaving

Whereas the two previously discussed mitigation schemes involve MAC layer adjustments, the next two schemes to be proposed concern the PHY layer. Time interleaving can be employed in a Wi-Fi system to reduce the bit error rate. Wi-Fi generally only supports frequency interleaving (described in Subsection 3.3.1), but time interleaving would be useful because a radar pulse (assuming it is significantly shorter in time duration than a Wi-Fi frame) is more likely to interfere with one OFDM symbol than with many. So if data and FEC bits for a codeword are interleaved across OFDM symbols, it is more likely that a radar pulse does not interfere with all the bits pertaining to the codeword, increasing the likelihood of successful reception.

Time interleaving in Wi-Fi would come with some downsides however, including increased latency from having to wait for more than one OFDM symbol to be received before making data bit sequence decisions, as well as potential conflicts with MAC timing periods such as SIFS. Another downside is that nodes employing time interleaving would not be compatible with legacy nodes that do not support it.

5.4 Selective Erasure of Outliers

In [16], a mitigation scheme called *Selective Erasure of Outliers* is proposed for radar interference on LTE systems using soft decision decoding at the receiver. Typical soft decision decoders use noise variance estimates from known symbols (often pilots) to calculate the log likelihood ratio (LLR) for each received bit, which at a high magnitude value indicates low uncertainty in bit identity (0 or 1) and at a low magnitude value indicates high uncertainty. The LLRs are employed in the soft decoder's decision process.

Radar interference upsets LLR calculations, tending to make the LLR much less indicative of the actual identity of the bits. Thus, it makes the soft decoder's decision less likely to be correct. In addition, radar interference frequently causes very large LLR values. This can be exploited to identify outlier LLR values that were impacted by radar interference. The Selective Erasure of Outliers scheme erases the LLR values (sets it to 0) of identified outliers. Since LLR values of 0 signify no certainty in bit identity, this scheme prevents detected outliers (from radar interference) from contributing to the soft decision process. In doing so, bit error and hence frame error probability is reduced. Figure 5.1 displays the effects and performance gain of the scheme for simulated down-link LTE with two different modulation schemes. In the cases shown, Selective Erasure of Outliers yields up to an effective 20 dB reduction in radar INR—a huge improvement. This short description of the mitigation scheme is summarized from [16], which should be consulted for more detail.

As LTE and Wi-Fi are both OFDM systems, the Selective Erasure of Outliers scheme should also be effective in improving Wi-Fi (with soft decoding) communication performance amidst radar interference. To achieve high performance gains from this mitigation scheme,

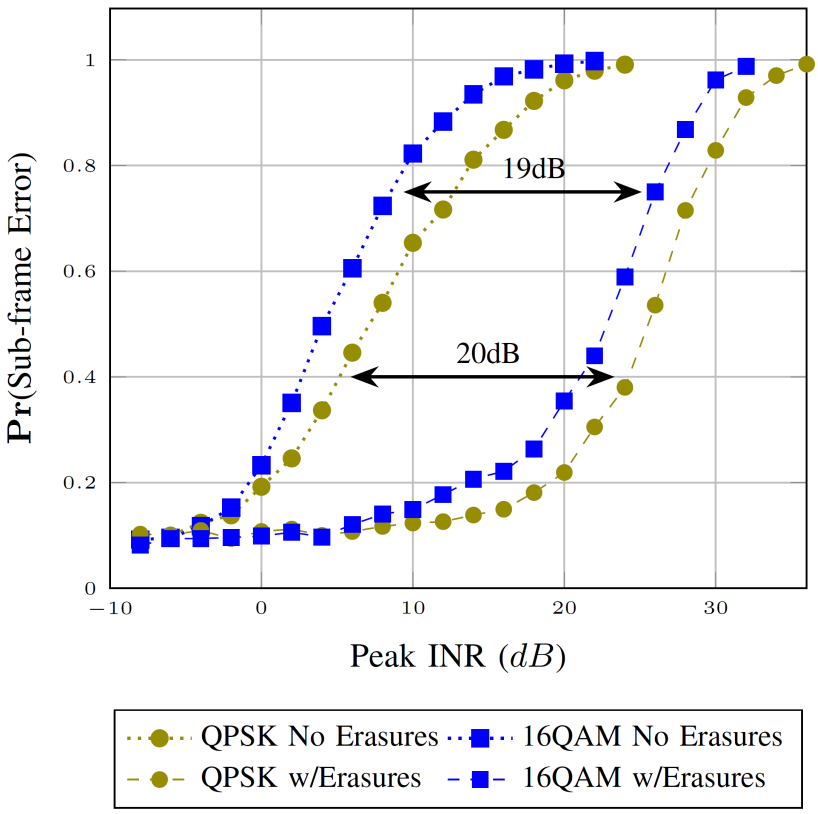

Figure 5.1: Effect of Selective Erasure of Outliers on simulated down-link LTE. Retrieved from [16]

time interleaving can be applied to the Wi-Fi PHY system, as discussed in Section 5.3, but that would prevent backward compatibility with older nodes.

Chapter 6
CONCLUSION

The link and system simulators provide a highly configurable software framework to analyze effects of radar interference on Wi-Fi networks. The link simulator is a detailed signal level model of the full Wi-Fi OFDM transmit and receive chains. It is easy to dive into any of the several blocks and observe how radar interference affects its operation or outcome. The ns-3 simulator provides a deep full-stack Wi-Fi model, which proves invaluable for simulating lifelike scenarios. The features added to ns-3 for radar interference allow users to leverage the wealth of features in ns-3 and apply them to Wi-Fi-radar coexistence studies. And the mapping of the link simulator results to ns-3 effectively enables the system simulations to have more control over transceiver level parameters and how they affect the network performance. The source code for both simulators will be hosted on the UW Fundamentals of Networking Laboratory website at: **http://depts.washington.edu/funlab/**.

The current results garnered by both simulators grant valuable insight into the nature of the Wi-Fi-radar spectrum sharing problem. The link simulation results indicate that the impact of radar interference on the Wi-Fi frame payload can be worse from an error rate perspective than preamble interference, which is concerning since radar pulses will interfere with the payload the vast majority of the time. They also showed that aggressive channel estimators may be able to recover well from radar interference. The system simulation results demonstrate that radar with long rest periods may not be very detrimental to Wi-Fi, even if it is at a relatively close range. An interesting observation was made that radar scan rate did not affect average throughput very much, but slower scan rates yield longer throughput time dips, which can be damaging to time sensitive applications. The radar PRF was found far and away to be the radar parameter that affects Wi-Fi performance the most.

From the analyzed results, the performance impact in some situations was so bad that it prompted the author to look into mitigation schemes for improving coexistence between radar and Wi-Fi. Two MAC layer and two PHY layer schemes were proposed, which all address different problems that tend to occur with strong radar interference.

6.1 Future Work

There is still much to be done in the area of Wi-Fi-radar coexistence, even from a simulation perspective. Future potential work on this project is highlighted in this section.

The proposed mitigation schemes in Chapter 5 have not been simulated, which should be done to prove their validity and discover how helpful they would be in improving coexistence performance.

More accurate propagation loss models could be implemented for both Wi-Fi and radar at the system level to make simulations more accurate, such as accounting for walls, obstructions, and terrain (especially for radar). To increase realism, channel models can be investigated and implemented for radar instead of using flat fading in the link sim.

Link sim additions, of which there are many options, would prove useful in obtaining more varied and accurate results. More precisely modeling the Wi-Fi receiver AGC and simulating results based on radar interfering with the AGC portion of the frame would be interesting, since it would likely cause the AGC to set the gain too low, reducing the resolution of the rest of the frame. The L-SIG portion of the frame could also be modeled for greater precision in results, as it carries necessary information that a frame could not recover from if corrupted.

Appendix A

ALL SYSTEM SIMULATION RESULTS

This appendix contains all system simulation results corresponding to settings in Tables 4.5, 4.6, and 4.7, scenarios in Tables 4.8 and 4.9, and variables in Table 4.10. In the plots of total throughput versus time, the interval between data points is 100 milliseconds, with each throughput measurement averaged over the previous 100 milliseconds. Average throughput for each scenario without radar interference can be found in the following table.

Average Wi-Fi throughput over 15 seconds with no radar interference

	Scenario 1		Scenario 2	
	MCS 0	**MCS 7**	**MCS 0**	**MCS 7**
Upstream	0.642 Mbps	6.36 Mbps	0.646 Mbps	6.55 Mbps
Downstream	1.93 Mbps	18.1 Mbps	1.94 Mbps	19.4 Mbps
Total	2.57 Mbps	24.5 Mbps	2.59 Mbps	26.0 Mbps

A.1 Scenario 1

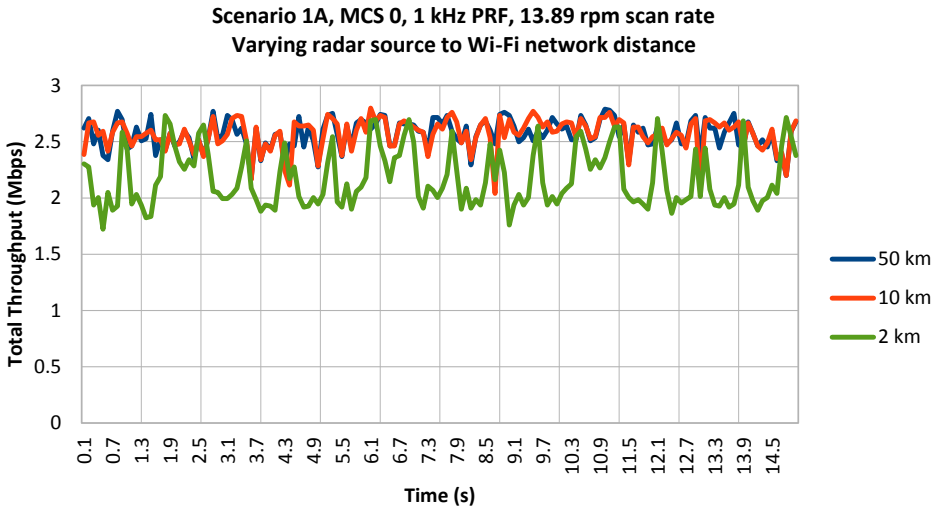

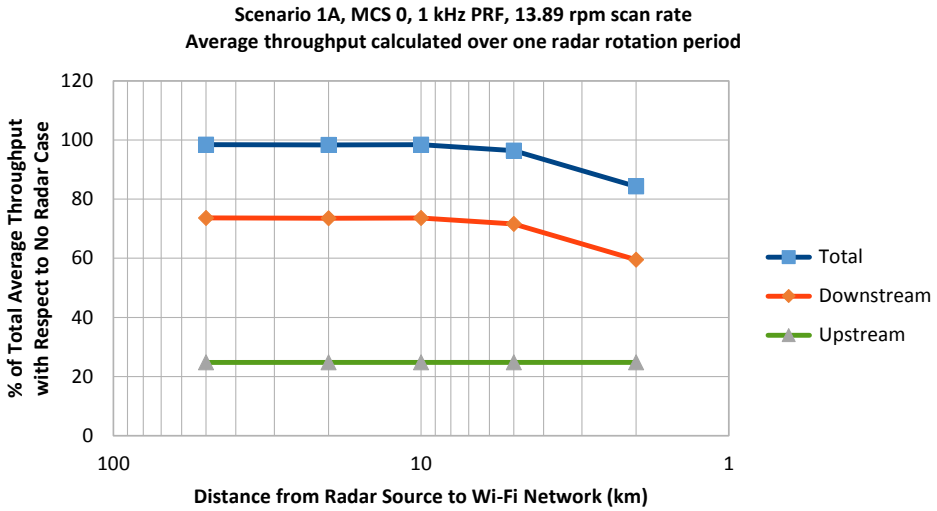

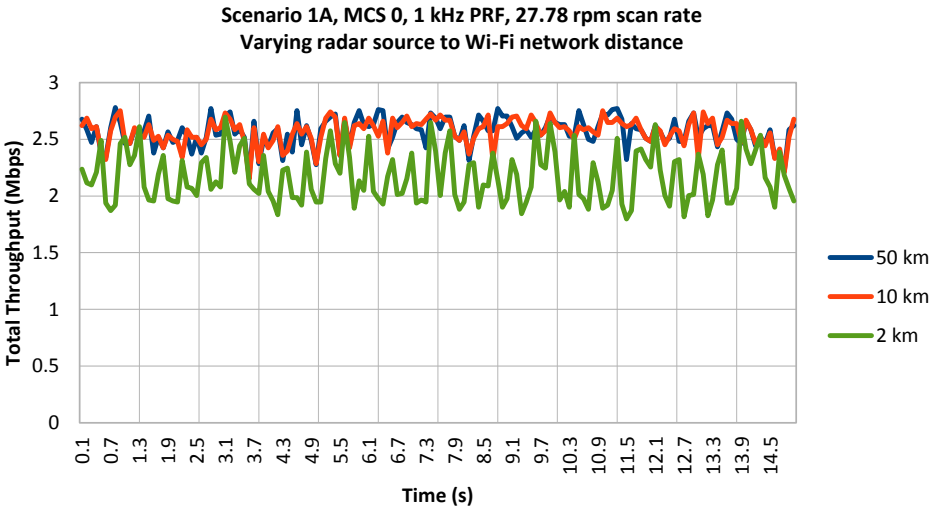
Scenario 1A, MCS 0, 1 kHz PRF, 27.78 rpm scan rate
Varying radar source to Wi-Fi network distance
Total Throughput (Mbps)
3
2.5
2
1.5
1
0.5
0
0.1
0.7
1.3
1.9
2.5
3.1
3.7
4.3
4.9
5.5
6.1
6.7
7.3
7.9
8.5
9.1
9.7
10.3
10.9
11.5
12.1
12.7
13.3
13.9
14.5
Time (s)
50 km
10 km
2 km

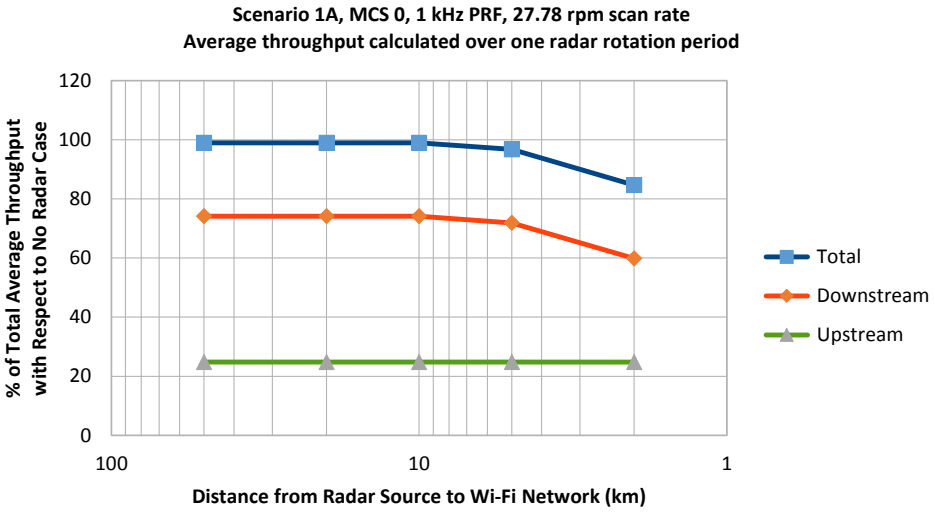
Scenario 1A, MCS 0, 1 kHz PRF, 27.78 rpm scan rate
Average throughput calculated over one radar rotation period
% of Total Average Throughput with Respect to No Radar Case
120
100
80
60
40
20
0
100
10
1
Distance from Radar Source to Wi-Fi Network (km)
Total
Downstream
Upstream

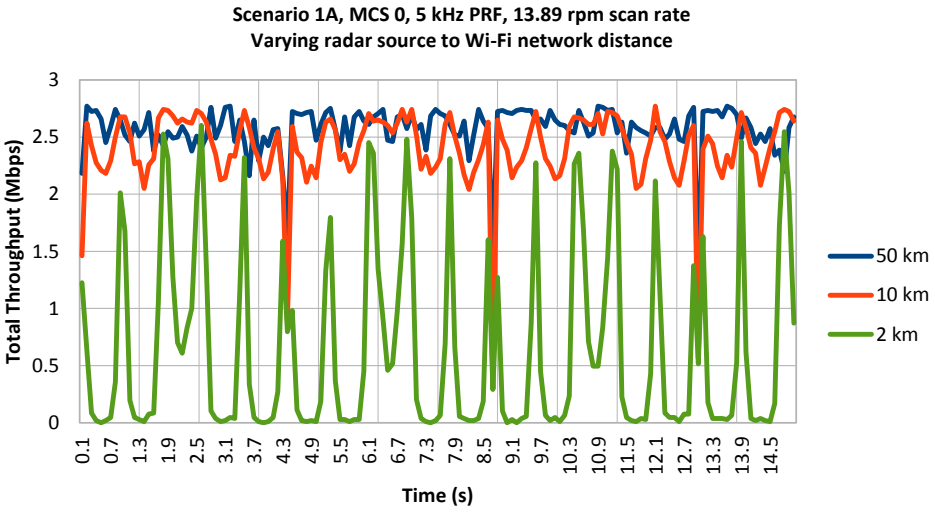
Scenario 1A, MCS 0, 5 kHz PRF, 13.89 rpm scan rate
Varying radar source to Wi-Fi network distance
Total Throughput (Mbps)
3
2.5
2
1.5
1
0.5
0
0.1
0.7
1.3
1.9
2.5
3.1
3.7
4.3
4.9
5.5
6.1
6.7
7.3
7.9
8.5
9.1
9.7
10.3
10.9
11.5
12.1
12.7
13.3
13.9
14.5
Time (s)
50 km
10 km
2 km

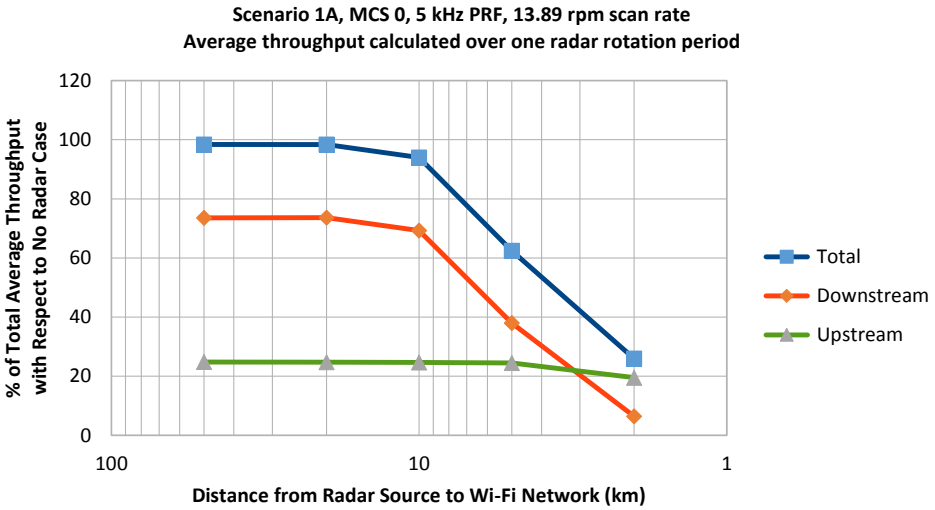
Scenario 1A, MCS 0, 5 kHz PRF, 13.89 rpm scan rate
Average throughput calculated over one radar rotation period
% of Total Average Throughput with Respect to No Radar Case
120
100
80
60
40
20
0
100
10
1
Distance from Radar Source to Wi-Fi Network (km)
Total
Downstream
Upstream

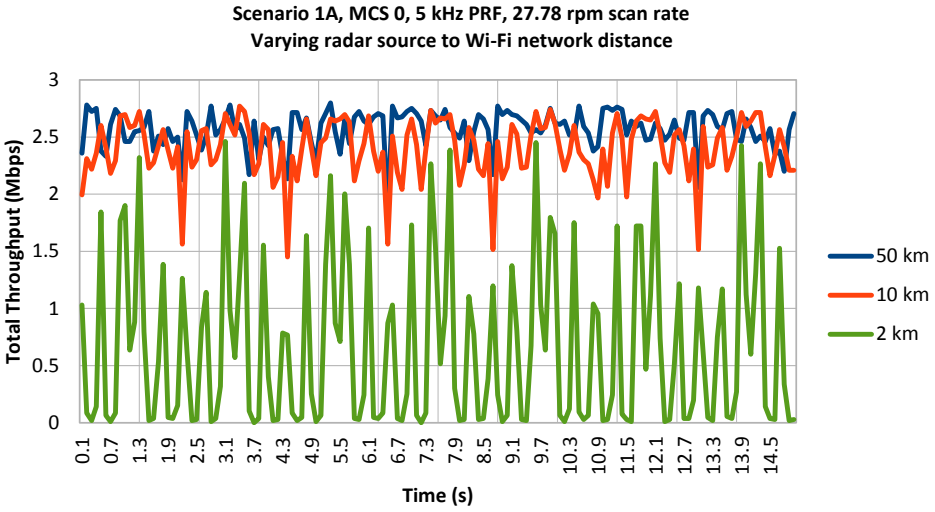
Scenario 1A, MCS 0, 5 kHz PRF, 27.78 rpm scan rate
Varying radar source to Wi-Fi network distance
Total Throughput (Mbps)
3
2.5
2
1.5
1
0.5
0
0.1
0.7
1.3
1.9
2.5
3.1
3.7
4.3
4.9
5.5
6.1
6.7
7.3
7.9
8.5
9.1
9.7
10.3
10.9
11.5
12.1
12.7
13.3
13.9
14.5
Time (s)
50 km
10 km
2 km

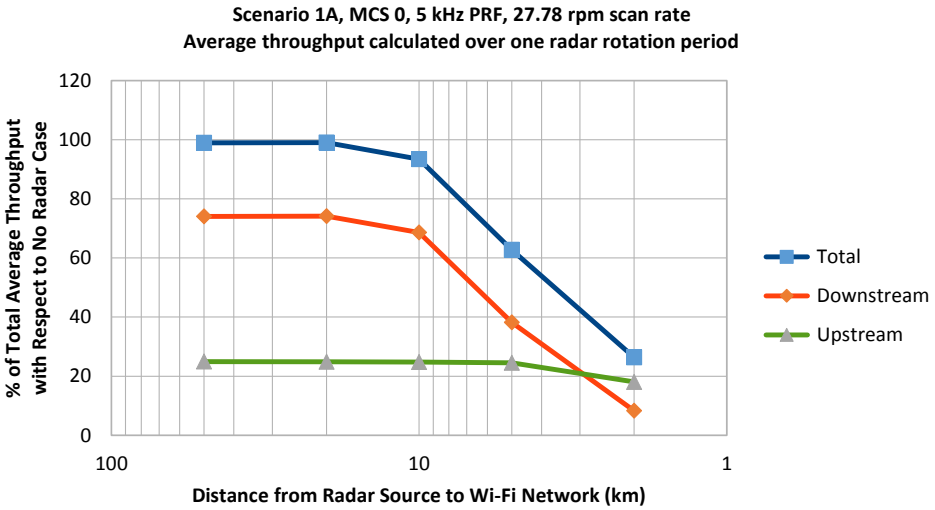
Scenario 1A, MCS 0, 5 kHz PRF, 27.78 rpm scan rate
Average throughput calculated over one radar rotation period
% of Total Average Throughput with Respect to No Radar Case
120
100
80
60
40
20
0
100
10
1
Distance from Radar Source to Wi-Fi Network (km)
Total
Downstream
Upstream

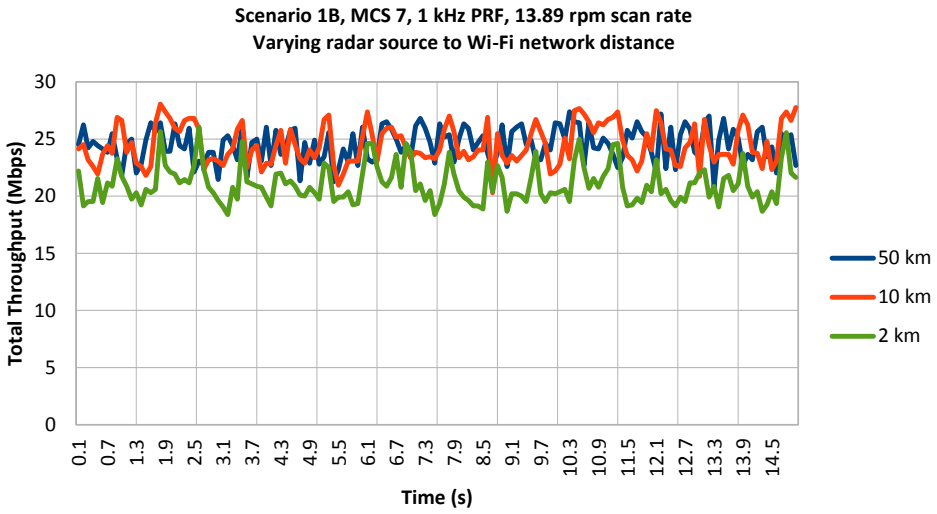
Scenario 1B, MCS 7, 1 kHz PRF, 13.89 rpm scan rate
Varying radar source to Wi-Fi network distance
Total Throughput (Mbps)
30
25
20
15
10
5
0
0.1
0.7
1.3
1.9
2.5
3.1
3.7
4.3
4.9
5.5
6.1
6.7
7.3
7.9
8.5
9.1
9.7
10.3
10.9
11.5
12.1
12.7
13.3
13.9
14.5
Time (s)
50 km
10 km
2 km

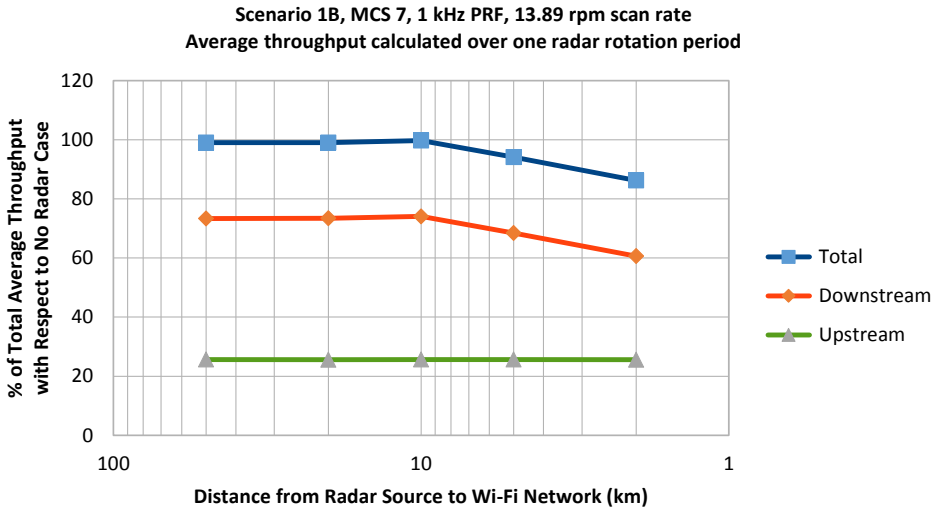
Scenario 1B, MCS 7, 1 kHz PRF, 13.89 rpm scan rate
Average throughput calculated over one radar rotation period
% of Total Average Throughput with Respect to No Radar Case
120
100
80
60
40
20
0
100
10
1
Distance from Radar Source to Wi-Fi Network (km)
Total
Downstream
Upstream

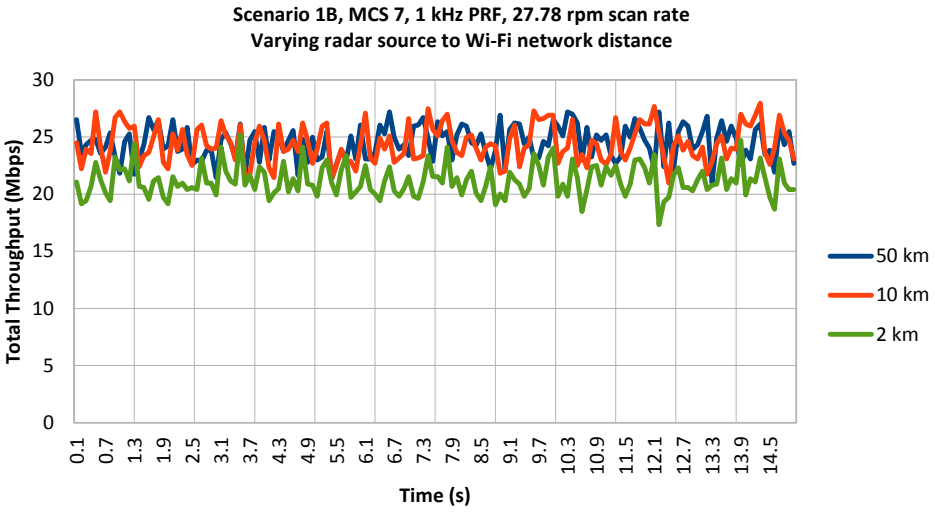
Scenario 1B, MCS 7, 1 kHz PRF, 27.78 rpm scan rate
Varying radar source to Wi-Fi network distance
Total Throughput (Mbps)
30
25
20
15
10
5
0
0.1
0.7
1.3
1.9
2.5
3.1
3.7
4.3
4.9
5.5
6.1
6.7
7.3
7.9
8.5
9.1
9.7
10.3
10.9
11.5
12.1
12.7
13.3
13.9
14.5
Time (s)
50 km
10 km
2 km

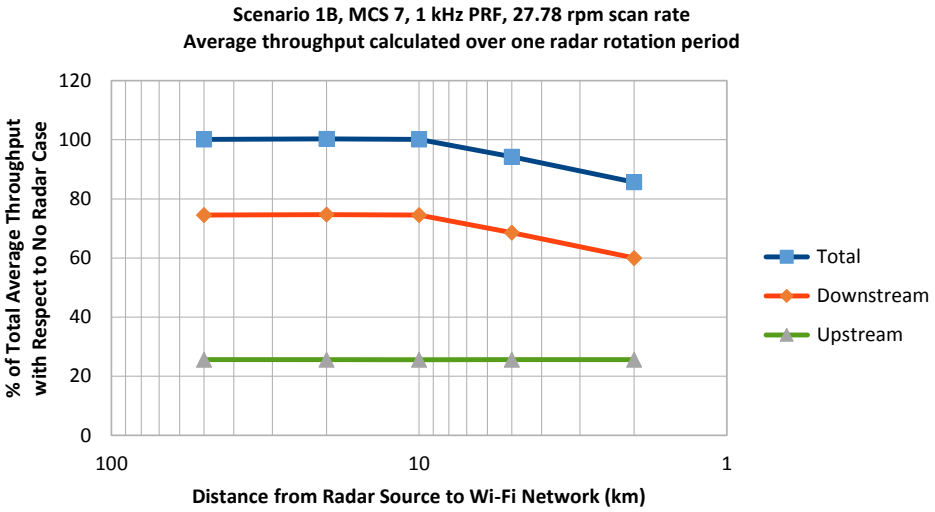
Scenario 1B, MCS 7, 1 kHz PRF, 27.78 rpm scan rate
Average throughput calculated over one radar rotation period
% of Total Average Throughput with Respect to No Radar Case
120
100
80
60
40
20
0
100
10
1
Distance from Radar Source to Wi-Fi Network (km)
Total
Downstream
Upstream

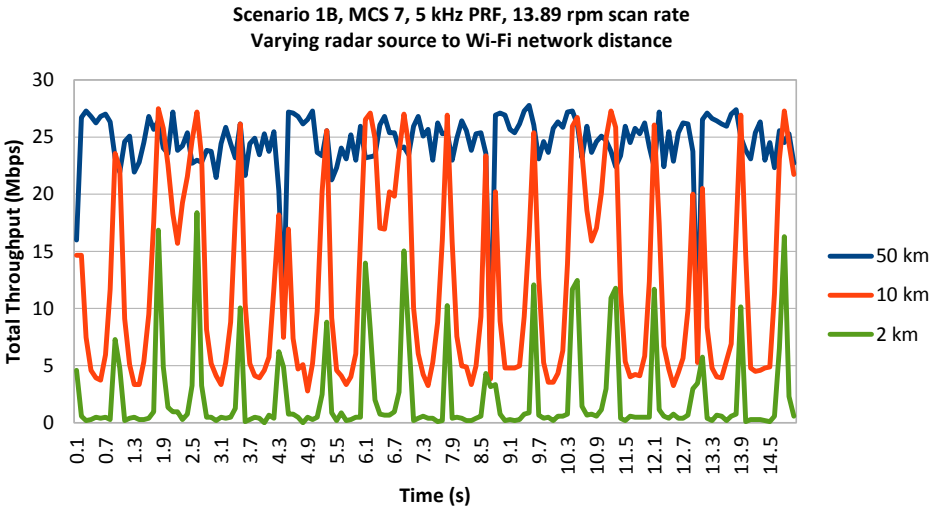
Scenario 1B, MCS 7, 5 kHz PRF, 13.89 rpm scan rate
Varying radar source to Wi-Fi network distance
Total Throughput (Mbps)
30
25
20
15
10
5
0
0.1
0.7
1.3
1.9
2.5
3.1
3.7
4.3
4.9
5.5
6.1
6.7
7.3
7.9
8.5
9.1
9.7
10.3
10.9
11.5
12.1
12.7
13.3
13.9
14.5
Time (s)
50 km
10 km
2 km

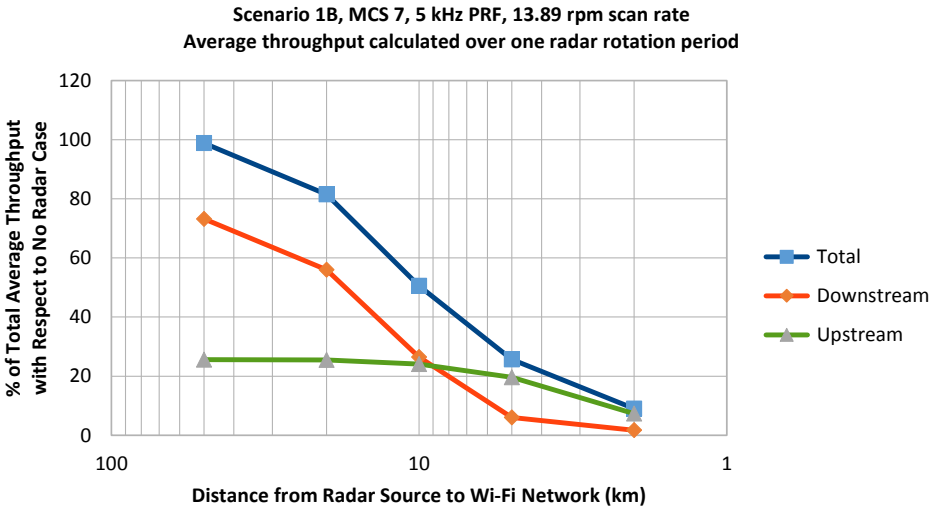
Scenario 1B, MCS 7, 5 kHz PRF, 13.89 rpm scan rate
Average throughput calculated over one radar rotation period
% of Total Average Throughput with Respect to No Radar Case
120
100
80
60
40
20
0
100
10
1
Distance from Radar Source to Wi-Fi Network (km)
Total
Downstream
Upstream

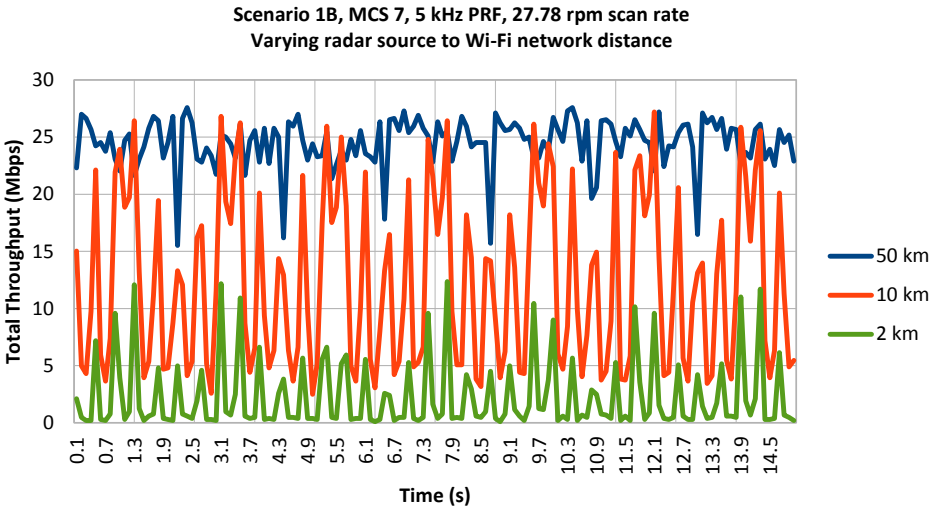
Scenario 1B, MCS 7, 5 kHz PRF, 27.78 rpm scan rate
Varying radar source to Wi-Fi network distance
Total Throughput (Mbps)
30
25
20
15
10
5
0
0.1
0.7
1.3
1.9
2.5
3.1
3.7
4.3
4.9
5.5
6.1
6.7
7.3
7.9
8.5
9.1
9.7
10.3
10.9
11.5
12.1
12.7
13.3
13.9
14.5
Time (s)
50 km
10 km
2 km

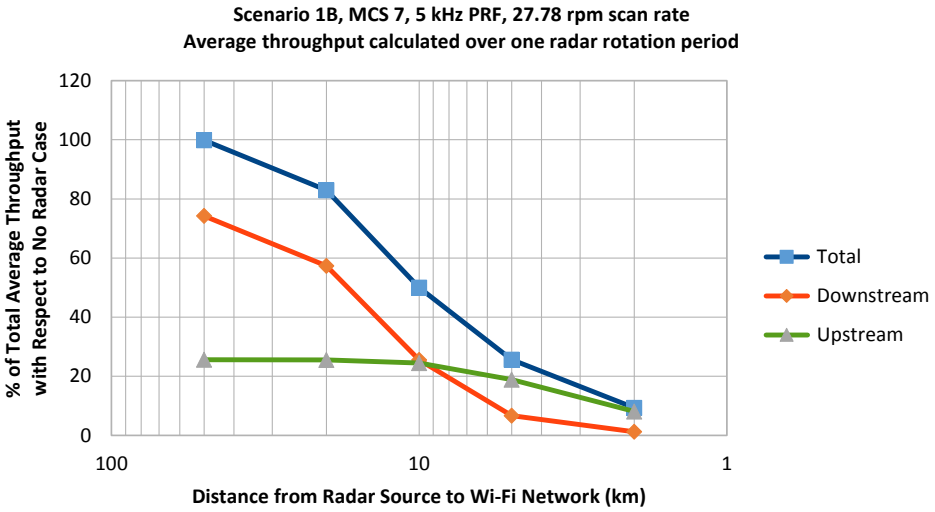
Scenario 1B, MCS 7, 5 kHz PRF, 27.78 rpm scan rate
Average throughput calculated over one radar rotation period
% of Total Average Throughput with Respect to No Radar Case
120
100
80
60
40
20
0
100
10
1
Distance from Radar Source to Wi-Fi Network (km)
Total
Downstream
Upstream

Single radar rotation period average throughput (Mbps) vs. radar distance for Scenario 1, 1 kHz PRF, 13.89 rpm scan rate

	MCS 0			MCS 7		
Distance	Upstream	Downstream	Total	Upstream	Downstream	Total
50 km	0.635917	1.89107	2.52698	6.2776	17.9616	24.2392
20 km	0.636133	1.88803	2.52417	6.26873	17.986	24.2548
10 km	0.636133	1.88977	2.5259	6.2776	18.139	24.4166
5 km	0.636133	1.83928	2.47542	6.2776	16.7624	23.04
2 km	0.636133	1.5275	2.16363	6.26873	14.8428	21.1115

Single radar rotation period average throughput (Mbps) vs. radar distance for Scenario 1, 1 kHz PRF, 27.78 rpm scan rate

	MCS 0			MCS 7		
Distance	Upstream	Downstream	Total	Upstream	Downstream	Total
50 km	0.637	1.9032	2.5402	6.26873	18.2387	24.5075
20 km	0.637	1.90277	2.53977	6.27317	18.2831	24.5562
10 km	0.637	1.9032	2.5402	6.2643	18.2432	24.5075
5 km	0.637	1.84557	2.48257	6.26873	16.7979	23.0666
2 km	0.636567	1.53617	2.17273	6.27317	14.6965	20.9697

Single radar rotation period average throughput (Mbps) vs. radar distance for Scenario 1, 5 kHz PRF, 13.89 rpm scan rate

	MCS 0			MCS 7		
Distance	Upstream	Downstream	Total	Upstream	Downstream	Total
50 km	0.636567	1.88912	2.52568	6.26873	17.9262	24.1949
20 km	0.634617	1.89042	2.52503	6.251	13.7101	19.9611
10 km	0.633317	1.77862	2.41193	5.90298	6.46823	12.3712
5 km	0.6279	0.975217	1.60312	4.81017	1.48073	6.2909
2 km	0.501367	0.1651	0.666467	1.79328	0.41895	2.21223

Single radar rotation period average throughput (Mbps) vs. radar distance for Scenario 1, 5 kHz PRF, 27.78 rpm scan rate

	MCS 0			MCS 7		
Distance	Upstream	Downstream	Total	Upstream	Downstream	Total
50 km	0.640033	1.90017	2.5402	6.26873	18.1767	24.4454
20 km	0.638733	1.9032	2.54193	6.25543	14.0492	20.3047
10 km	0.637	1.76193	2.39893	5.9983	6.251	12.2493
5 km	0.630067	0.981067	1.61113	4.6284	1.64033	6.26873
2 km	0.4641	0.2132	0.6773	1.99057	0.301467	2.29203

A.2 Scenario 2

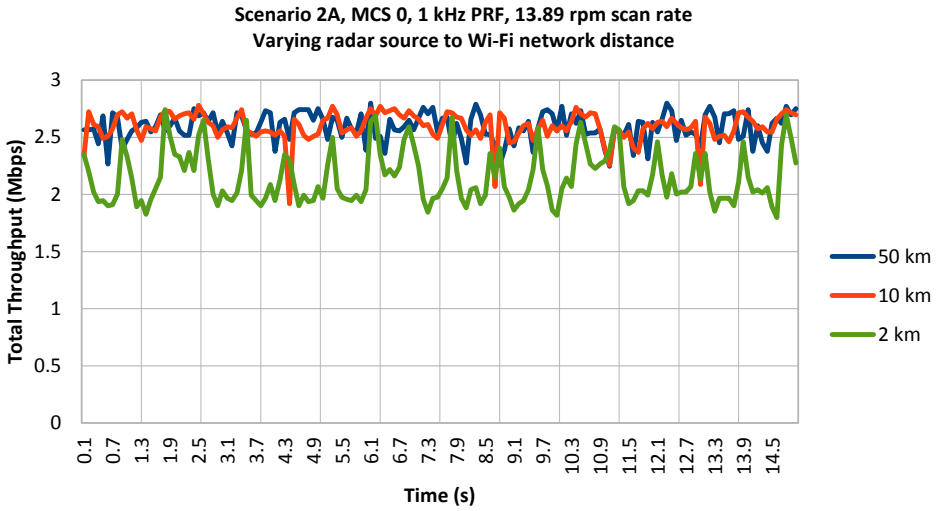

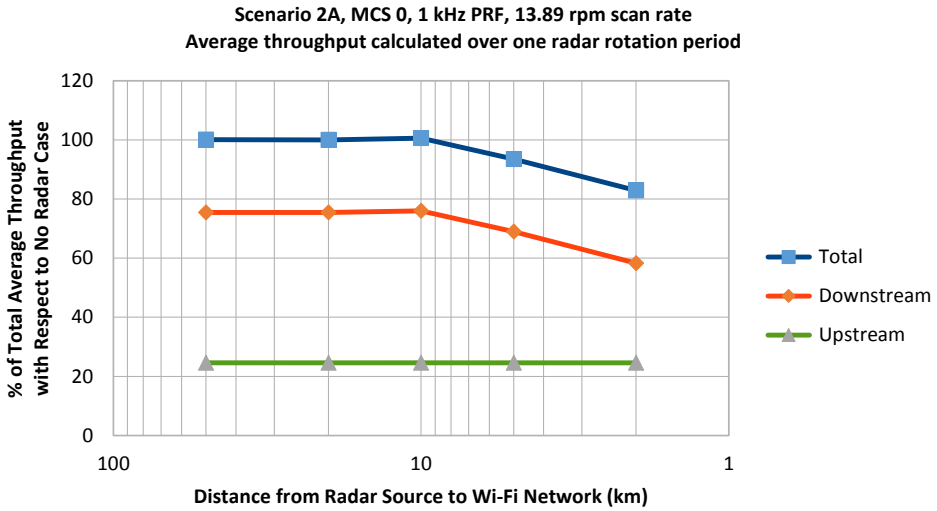

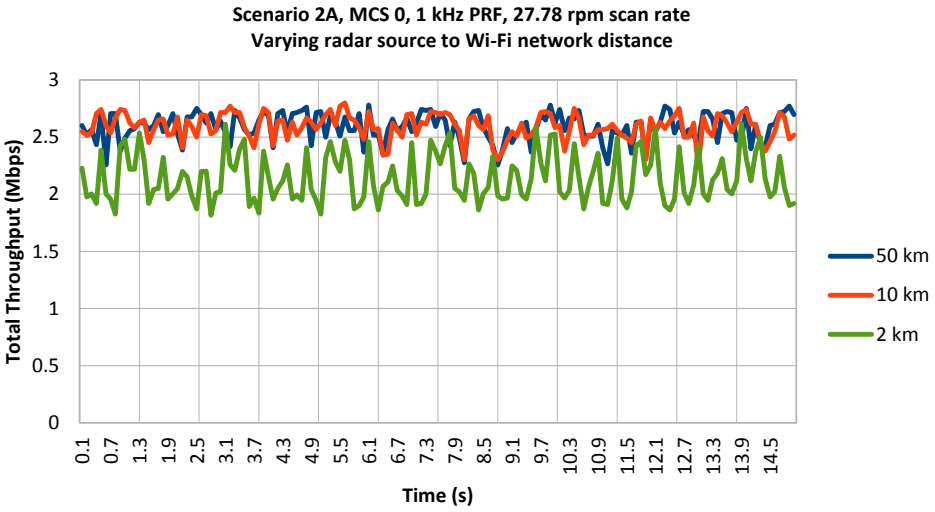
Scenario 2A, MCS 0, 1 kHz PRF, 27.78 rpm scan rate
Varying radar source to Wi-Fi network distance
Total Throughput (Mbps)
3
2.5
2
1.5
1
0.5
0
0.1
0.7
1.3
1.9
2.5
3.1
3.7
4.3
4.9
5.5
6.1
6.7
7.3
7.9
8.5
9.1
9.7
10.3
10.9
11.5
12.1
12.7
13.3
13.9
14.5
Time (s)
50 km
10 km
2 km

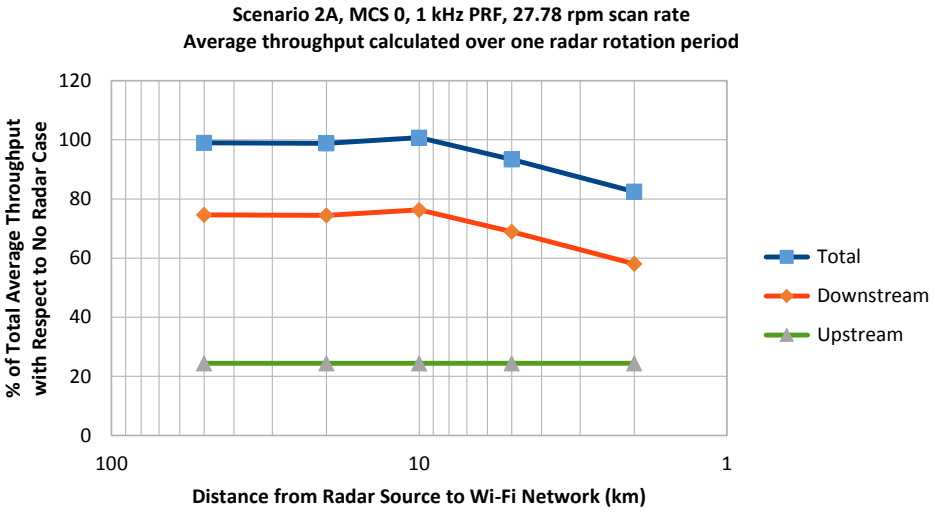
Scenario 2A, MCS 0, 1 kHz PRF, 27.78 rpm scan rate
Average throughput calculated over one radar rotation period
% of Total Average Throughput with Respect to No Radar Case
120
100
80
60
40
20
0
100
10
1
Distance from Radar Source to Wi-Fi Network (km)
Total
Downstream
Upstream

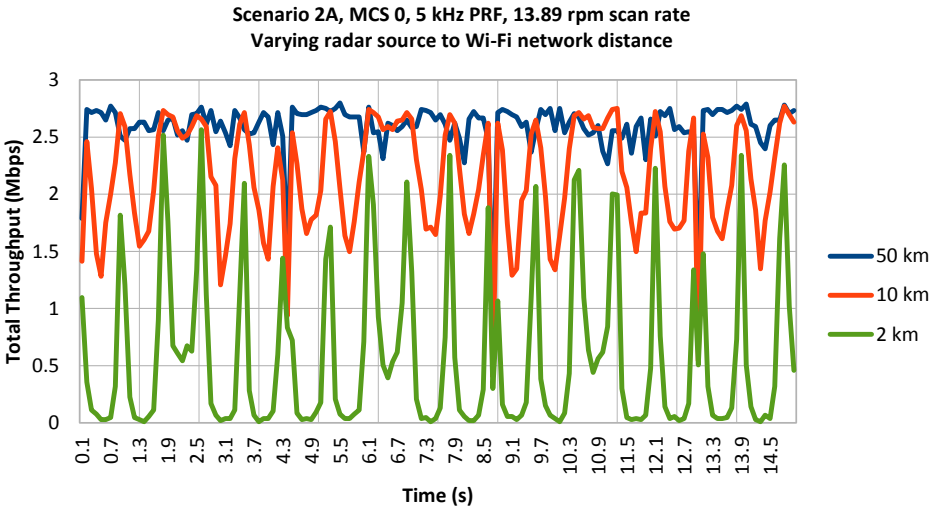
Scenario 2A, MCS 0, 5 kHz PRF, 13.89 rpm scan rate
Varying radar source to Wi-Fi network distance
Total Throughput (Mbps)
3
2.5
2
1.5
1
0.5
0
0.1
0.7
1.3
1.9
2.5
3.1
3.7
4.3
4.9
5.5
6.1
6.7
7.3
7.9
8.5
9.1
9.7
10.3
10.9
11.5
12.1
12.7
13.3
13.9
14.5
Time (s)
50 km
10 km
2 km

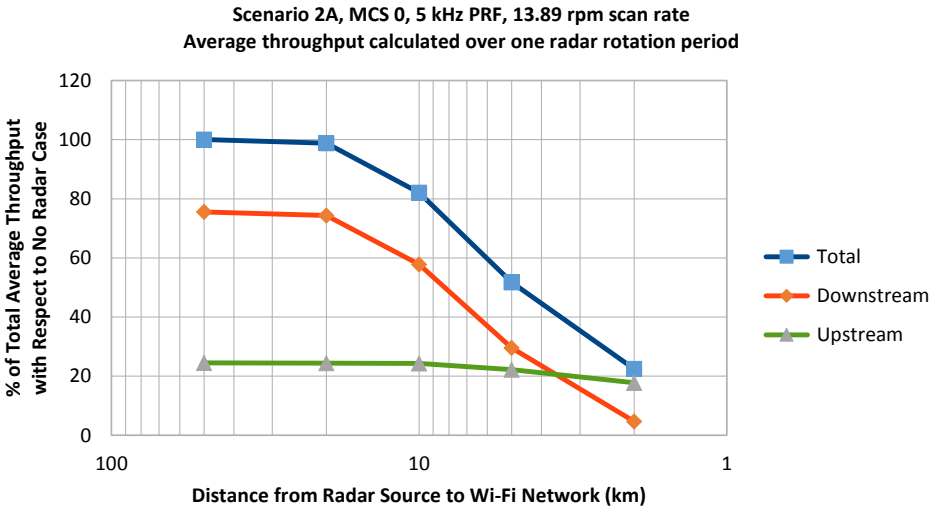
Scenario 2A, MCS 0, 5 kHz PRF, 13.89 rpm scan rate
Average throughput calculated over one radar rotation period
% of Total Average Throughput with Respect to No Radar Case
120
100
80
60
40
20
0
100
10
1
Distance from Radar Source to Wi-Fi Network (km)
Total
Downstream
Upstream

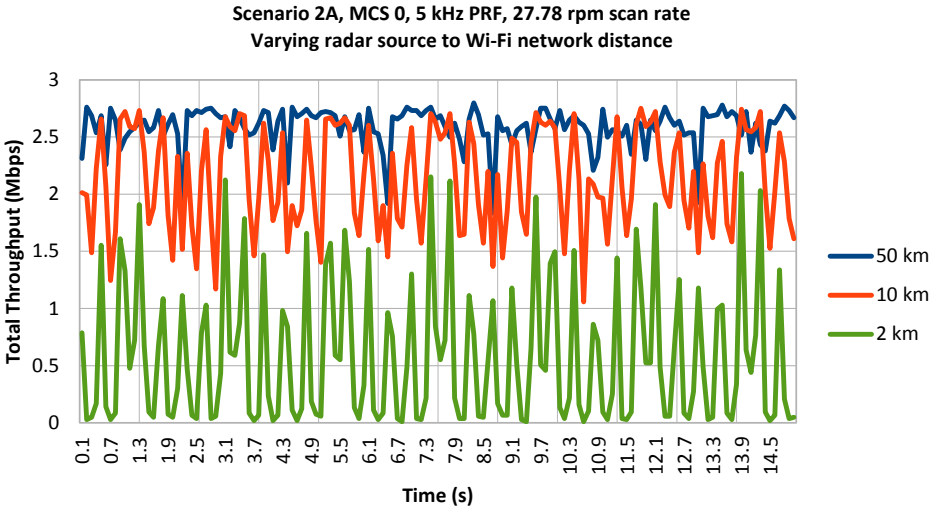
Scenario 2A, MCS 0, 5 kHz PRF, 27.78 rpm scan rate
Varying radar source to Wi-Fi network distance
Total Throughput (Mbps)
3
2.5
2
1.5
1
0.5
0
0.1
0.7
1.3
1.9
2.5
3.1
3.7
4.3
4.9
5.5
6.1
6.7
7.3
7.9
8.5
9.1
9.7
10.3
10.9
11.5
12.1
12.7
13.3
13.9
14.5
Time (s)
50 km
10 km
2 km

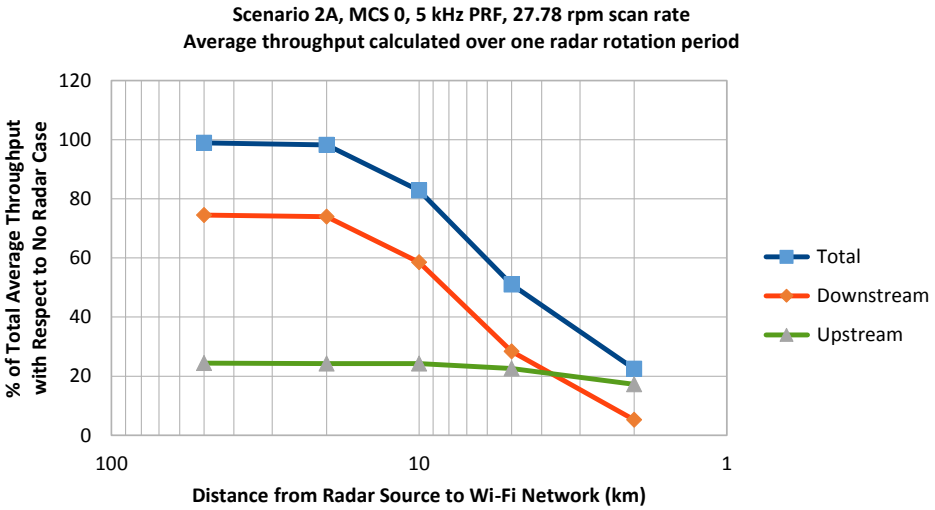
Scenario 2A, MCS 0, 5 kHz PRF, 27.78 rpm scan rate
Average throughput calculated over one radar rotation period
% of Total Average Throughput with Respect to No Radar Case
120
100
80
60
40
20
0
100
10
1
Distance from Radar Source to Wi-Fi Network (km)
Total
Downstream
Upstream

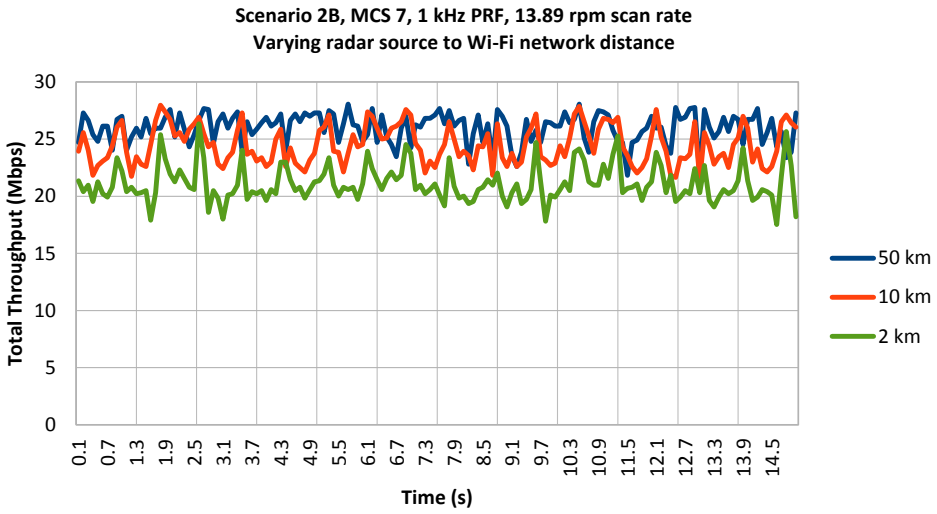
Scenario 2B, MCS 7, 1 kHz PRF, 13.89 rpm scan rate
Varying radar source to Wi-Fi network distance
Total Throughput (Mbps)
30
25
20
15
10
5
0
0.1
0.7
1.3
1.9
2.5
3.1
3.7
4.3
4.9
5.5
6.1
6.7
7.3
7.9
8.5
9.1
9.7
10.3
10.9
11.5
12.1
12.7
13.3
13.9
14.5
Time (s)
50 km
10 km
2 km

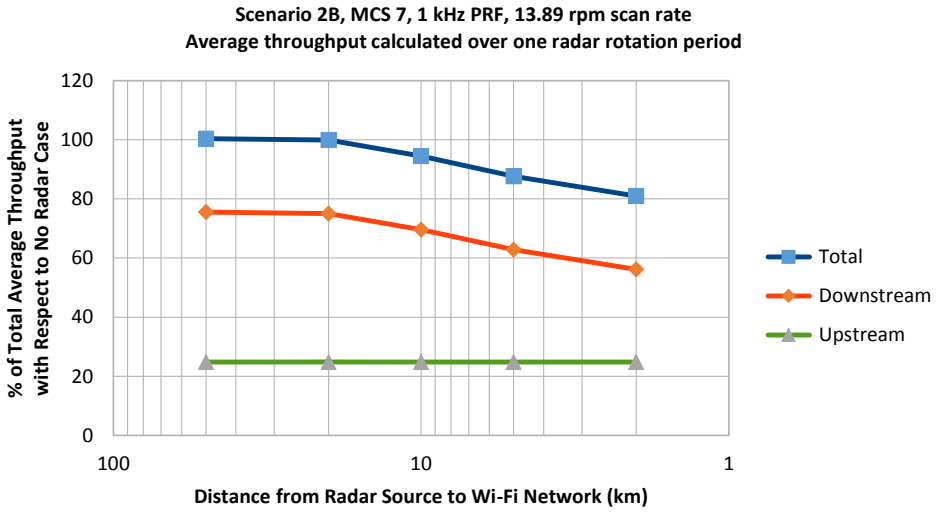
Scenario 2B, MCS 7, 1 kHz PRF, 13.89 rpm scan rate
Average throughput calculated over one radar rotation period
% of Total Average Throughput with Respect to No Radar Case
120
100
80
60
40
20
0
100
10
1
Distance from Radar Source to Wi-Fi Network (km)
Total
Downstream
Upstream

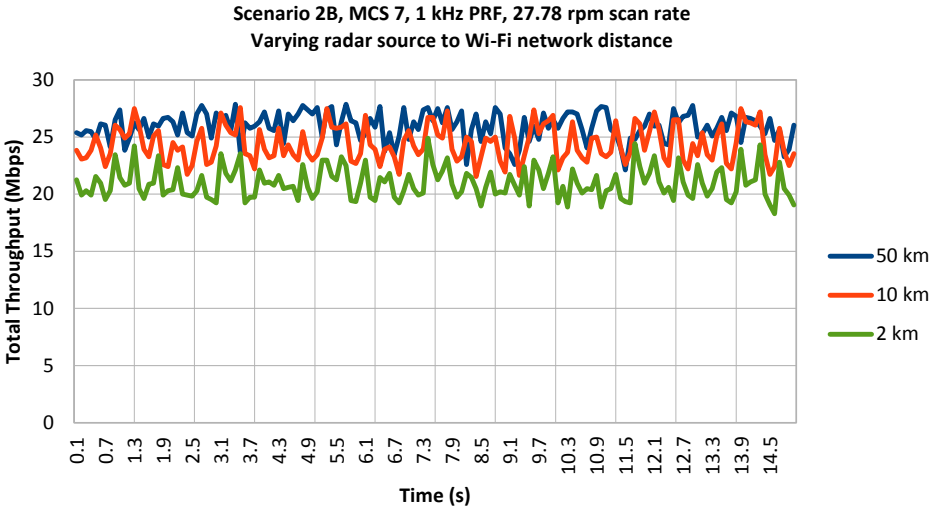
Scenario 2B, MCS 7, 1 kHz PRF, 27.78 rpm scan rate
Varying radar source to Wi-Fi network distance
Total Throughput (Mbps)
30
25
20
15
10
5
0
0.1
0.7
1.3
1.9
2.5
3.1
3.7
4.3
4.9
5.5
6.1
6.7
7.3
7.9
8.5
9.1
9.7
10.3
10.9
11.5
12.1
12.7
13.3
13.9
14.5
Time (s)
50 km
10 km
2 km

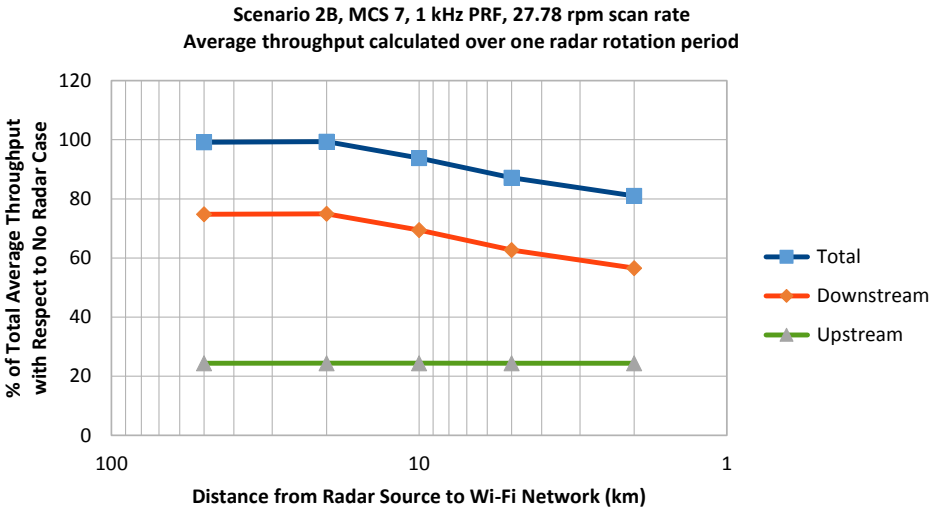
Scenario 2B, MCS 7, 1 kHz PRF, 27.78 rpm scan rate
Average throughput calculated over one radar rotation period
% of Total Average Throughput with Respect to No Radar Case
120
100
80
60
40
20
0
100
10
1
Distance from Radar Source to Wi-Fi Network (km)
Total
Downstream
Upstream

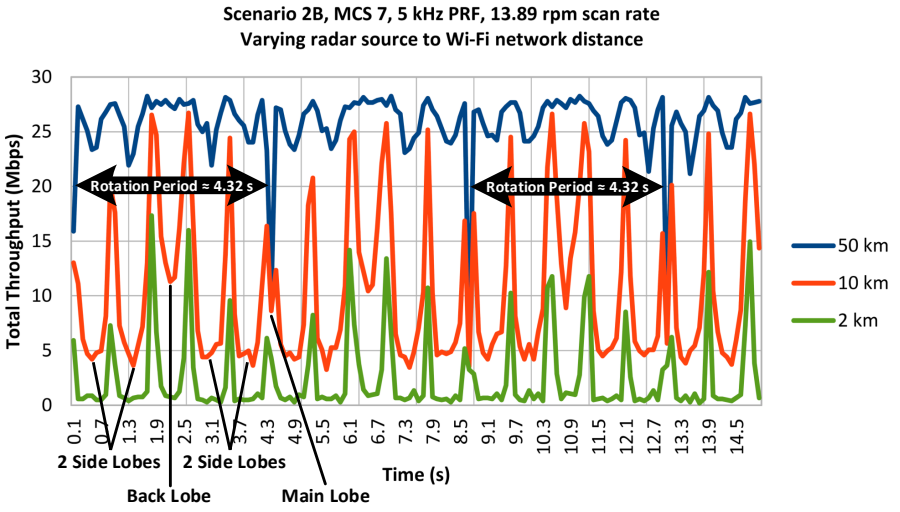
Scenario 2B, MCS 7, 5 kHz PRF, 13.89 rpm scan rate
Varying radar source to Wi-Fi network distance
Total Throughput (Mbps)
30
25
20
15
10
5
0
Rotation Period ≈ 4.32 s
Rotation Period ≈ 4.32 s
0.1
0.7
1.3
1.9
2.5
3.1
3.7
4.3
4.9
5.5
6.1
6.7
7.3
7.9
8.5
9.1
9.7
10.3
10.9
11.5
12.1
12.7
13.3
13.9
14.5
Time (s)
2 Side Lobes
2 Side Lobes
Back Lobe
Main Lobe
50 km
10 km
2 km

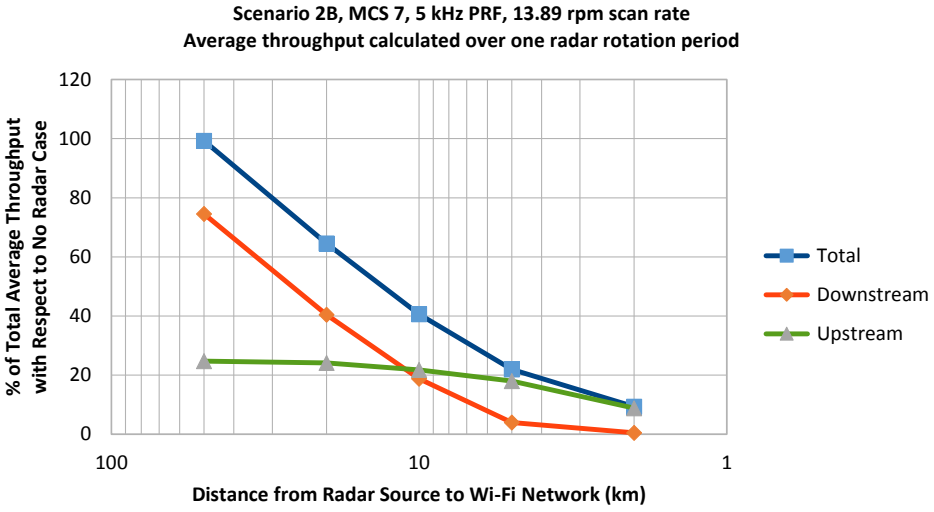
Scenario 2B, MCS 7, 5 kHz PRF, 13.89 rpm scan rate
Average throughput calculated over one radar rotation period
% of Total Average Throughput with Respect to No Radar Case
120
100
80
60
40
20
0
100
10
1
Distance from Radar Source to Wi-Fi Network (km)
Total
Downstream
Upstream

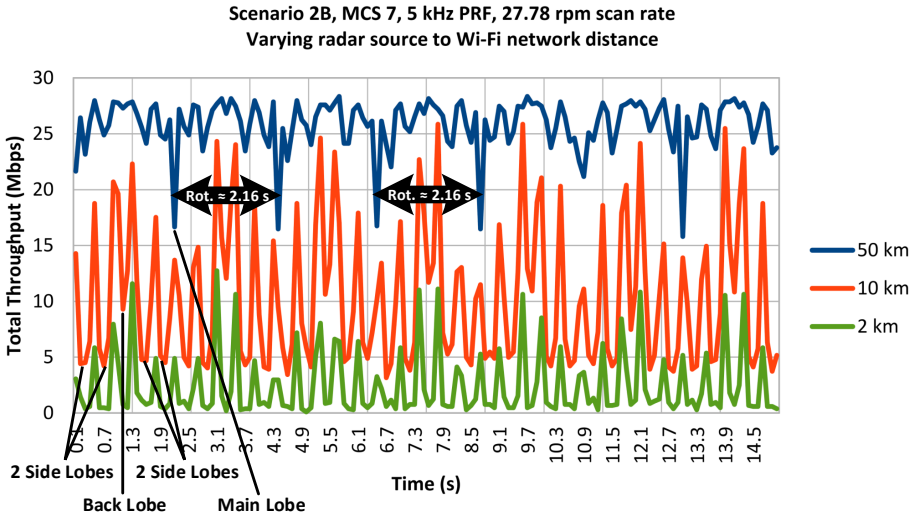
Scenario 2B, MCS 7, 5 kHz PRF, 27.78 rpm scan rate
Varying radar source to Wi-Fi network distance
Total Throughput (Mbps)
30
25
20
15
10
5
0
Rot. ≈ 2.16 s
Rot. ≈ 2.16 s
0.1
0.7
1.3
1.9
2.5
3.1
3.7
4.3
4.9
5.5
6.1
6.7
7.3
7.9
8.5
9.1
9.7
10.3
10.9
11.5
12.1
12.7
13.3
13.9
14.5
Time (s)
2 Side Lobes
2 Side Lobes
Back Lobe
Main Lobe
50 km
10 km
2 km

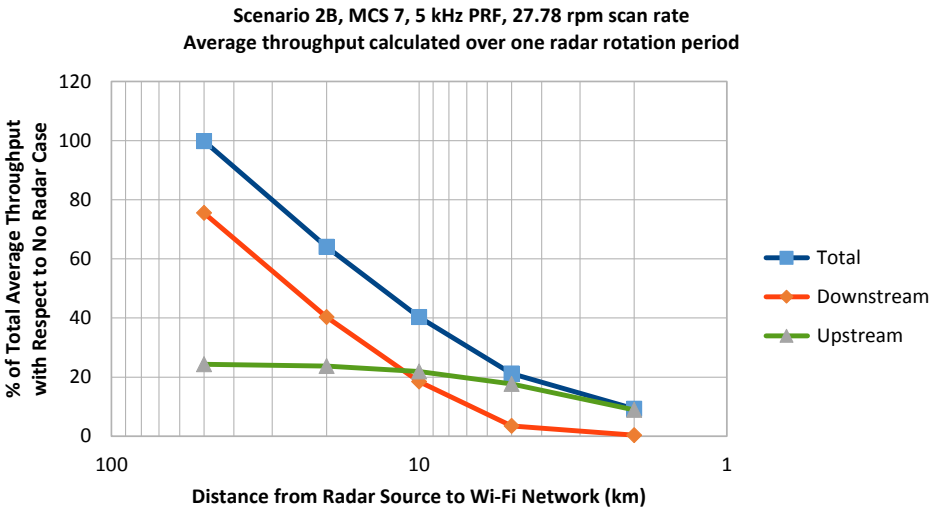
Scenario 2B, MCS 7, 5 kHz PRF, 27.78 rpm scan rate
Average throughput calculated over one radar rotation period
% of Total Average Throughput with Respect to No Radar Case
120
100
80
60
40
20
0
100
10
1
Distance from Radar Source to Wi-Fi Network (km)
Total
Downstream
Upstream

Single radar rotation period average throughput (Mbps) vs. radar distance for Scenario 2, 1 kHz PRF, 13.89 rpm scan rate

	MCS 0			MCS 7		
Distance	Upstream	Downstream	Total	Upstream	Downstream	Total
50 km	0.636783	1.9539	2.59068	6.44385	19.602	26.0458
20 km	0.636567	1.95325	2.58982	6.44828	19.4756	25.9239
10 km	0.636783	1.96733	2.60412	6.44163	18.0658	24.5075
5 km	0.636567	1.7849	2.42147	6.44163	16.3124	22.7541
2 km	0.636567	1.50908	2.14565	6.44163	14.5701	21.0118

Single radar rotation period average throughput (Mbps) vs. radar distance for Scenario 2, 1 kHz PRF, 27.78 rpm scan rate

	MCS 0			MCS 7		
Distance	Upstream	Downstream	Total	Upstream	Downstream	Total
50 km	0.632233	1.93093	2.56317	6.3308	19.4091	25.7399
20 km	0.632233	1.9266	2.55883	6.33967	19.4446	25.7843
10 km	0.632233	1.976	2.60823	6.33967	18.0126	24.3523
5 km	0.632233	1.78533	2.41757	6.33523	16.2792	22.6144
2 km	0.632233	1.5028	2.13503	6.33523	14.6832	21.0184

Single radar rotation period average throughput (Mbps) vs. radar distance for Scenario 2, 5 kHz PRF, 13.89 rpm scan rate

	MCS 0			MCS 7		
Distance	Upstream	Downstream	Total	Upstream	Downstream	Total
50 km	0.6344	1.95585	2.59025	6.41947	19.3382	25.7577
20 km	0.632233	1.92595	2.55818	6.26873	10.4826	16.7514
10 km	0.630067	1.49522	2.12528	5.65693	4.87667	10.5336
5 km	0.575683	0.76505	1.34073	4.66165	1.03962	5.70127
2 km	0.4602	0.121333	0.581533	2.29425	0.128567	2.42282

Single radar rotation period average throughput (Mbps) vs. radar distance for Scenario 2, 5 kHz PRF, 27.78 rpm scan rate

	MCS 0			MCS 7		
Distance	Upstream	Downstream	Total	Upstream	Downstream	Total
50 km	0.632233	1.92877	2.561	6.3175	19.5909	25.9084
20 km	0.628767	1.9149	2.54367	6.1712	10.4494	16.6206
10 km	0.628333	1.51537	2.1437	5.67467	4.79243	10.4671
5 km	0.585867	0.7345	1.32037	4.59293	0.908833	5.50177
2 km	0.446767	0.135633	0.5824	2.3142	0.0975333	2.41173

BIBLIOGRAPHY

[1] "Cisco visual networking index: Global mobile data traffic forecast update, 2014-2019," White Paper, Cisco Systems, 2015. [Online]. Available: http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white_paper_c11-520862.pdf

[2] U. S. O. of Spectrum Management, "United states frequency allocations : the radio spectrum," U.S. Dept. of Commerce, National Telecommunications and Information Administration, Office of Spectrum Management, Aug. 2011. [Online]. Available: http://www.ntia.doc.gov/files/ntia/publications/spectrum_wall_chart_aug2011.pdf

[3] Q. Zhuang. (2009, May) Crystal-oscillator psr improves 802.11n evm performance. Mobile Dev & Design. [Online]. Available: http://mobiledevdesign.com/learning-resources/crystal-oscillator-psr-improves-80211n-evm-performance

[4] E. Perahia and R. Stacey, *Next Generation Wireless LANs: 802.11n and 802.11ac*, 2nd ed. Cambridge University Press, 2013.

[5] R. van Nee and R. Prasad, *OFDM for Wireless Multimedia Communications*, ser. Artech House Universal Personal Communications. Artech House, 2000.

[6] U. Madhow, *Fundamentals of Digital Communication*. Cambridge University Press, 2008.

[7] Y. S. Cho, J. Kim, W. Y. Yang, and C. G. Kang, *MIMO-OFDM Wireless Communications with MATLAB*. IEEE Press : John Wiley & Sons (Asia), 2010.

[8] 802.11n ofdm signal structure. Keysight Technologies. [Online]. Available: http://rfmw.em.keysight.com/wireless/helpfiles/n7617a/mimo_ofdm_signal_structure.htm

[9] F. Hessar, "Spectrum sharing in white spaces," Ph.D. dissertation, University of Washington, 2015.

[10] (1996) 4.2 real aperture radar. Japan Association of Remote Sensing. [Online]. Available: http://wtlab.iis.u-tokyo.ac.jp/~wataru/lecture/rsgis/rsnote/cp4/cp4-2.htm

[11] (2015, Mar.) Interleaving explained. Kitz. [Online]. Available: http://www.kitz.co.uk/adsl/interleaving.htm

[12] (2015, Apr.) Wifi – model library. [Online]. Available: https://www.nsnam.org/docs/models/html/wifi.html

[13] F. Hessar, "Signal simulator for radars," Tech. Rep., Oct. 2010.

[14] S. Rickman. (2014, Jan.) dirplot. [Online]. Available: http://www.mathworks.com/matlabcentral/fileexchange/1251-dirplot

[15] (2015, Mar.) Pyviz - nsnam. [Online]. Available: https://www.nsnam.org/wiki/PyViz

[16] H.-A. Safavi-Naeini, C. Ghosh, E. Visotsky, R. Ratasuk, and S. Roy, "Impact and mitigation of narrow-band radar interference in down-link lte," in *Proc. IEEE ICC'15*, 2015.

[17] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 5: Enhancements for Higher Throughput*, IEEE Std. 802.11n, 2009.

[18] V. Erceg, L. Schumacher, P. Kyritsi *et al.*, "Tgn channel models," IEEE, Tech. Rep. 802.11-03/940r4, May 2004.

[19] T. S. Rappaport, *Wireless Communications: Principles and Practice*, 2nd ed. Prentice Hall, 2002.

[20] Shared spectrum access for radar and communications (ssparc). United States Defense Advanced Research Projects Agency. [Online]. Available: http://www.darpa.mil/Our_Work/STO/Programs/Shared_Spectrum_Access_for_Radar_and_Communications_%28SSPARC%29.aspx

[21] (2015) ns-3. [Online]. Available: http://www.nsnam.org/

[22] T. K. Sarkar, "Spectrum management for simultaneous transmission of wifi and radar signals," in *Proc. ICEAA'14*, 2014.

[23] Z. Horváth and D. Varga, "Channel allocation technique for eliminating interference caused by rlans on meteorological radars in 5 ghz band," *Infocommunications Journal*, vol. LXIV, 2009.

[24] C. Pisa, A. Garcia-Saavedra, and D. J. Leith, "Investigating bit error patterns for radar pulse detection in ieee 802.11," in *Proc. ICWMC'14*, 2014.

[25] "In the matter of: Amendment of the commission's rules with regard to commercial operations in the 3550-3650 mhz band; reply comments of google inc. on the further notice of proposed rulemaking," GN Docket No. 12-354, Federal Communications Commission, Aug. 2014. [Online]. Available: http://apps.fcc.gov/ecfs/document/view;jsessionid=22SnJGFQqg9Mwq6dQpnFhyNM9Jp3hxvkQLq2Vdvhv1Z12ZsJ1Bfv!44062279!-58662085?id=7521768445

[26] F. H. Sanders, J. E. Carroll, G. A. Sanders, and R. L. Sole, "Effects of radar interference on lte base station receiver performance," U.S. Dept. of Commerce, National Telecommunications and Information Administration, Tech. Rep. 14-499, Dec. 2013.

[27] R. Saruthirathanaworakun, "Opportunistic sharing between rotating radar and cellular," *IEEE Journal on Selected Areas in Communications*, Nov. 2012.

[28] E. Krouk and S. Semenov, Eds., *Modulation and Coding Techniques in Wireless Communications*. John Wiley & Sons, 2011.

[29] B. H. Walke, S. Mangold, and L. Berlemann, *IEEE 802 Wireless Systems: Protocols, Multi-hop Mesh/Relaying, Performance and Spectrum Coexistence*. John Wiley & Sons, 2006.

[30] M. Gast, *802.11 Wireless Networks: The Definitive Guide*, 2nd ed. O'Reilly Media, 2005.

[31] N. Levanon, *Radar Principles*. John Wiley & Sons, 1988.

[32] (2004) Doppler radar. University of Illinois at Urbana-Champaign Department of Atmospheric Sciences. [Online]. Available: http://www.atmos.illinois.edu/courses/atmos410-fa04/Doppler%20radar.ppt

[33] M. A. Richards, J. A. Scheer, and W. A. Holm, *Principles of Modern Radar, Volume I - Basic Principles*. SciTech Publishing, 2010.

[34] (1998) Radar systems. Federation of American Scientists. [Online]. Available: http://fas.org/man/dod-101/navy/docs/es310/radarsys/radarsys.htm

[35] J. L. Eaves and E. K. Reedy, Eds., *Principles of Modern Radar*. Springer US, 1987.

[36] K. Wang, M. Faulkner, J. Singh, and I. Tolochko, "Timing synchronization for 802.11a wlans under multipath channels," Centre for Telecommunications and Microelectronics, Victoria University. [Online]. Available: http://cegt201.bradley.edu/projects/proj2006/sradio/archives/Timing_sync_802_wlan_multipath.pdf

[37] E. W. Weisstein. Circular segment. MathWorld–A Wolfram Web Resource. [Online]. Available: http://mathworld.wolfram.com/CircularSegment.html

[38] N. Friedman, *The Naval Institute Guide to World Naval Weapon Systems*, 5th ed. Naval Institute Press, 2006.

[39] H.-A. Safavi-Naeini and S. Roy, "Radar detection in wi-fi networks," 2015, unpublished.