Antifragile Communications

Marc L. Lichtman

Dissertation submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Electrical Engineering

Jeffrey H. Reed, Chair
T. Charles Clancy
Jaime De La Ree
Michael J. Roan
Dennis G. Sweeney

July 14, 2016
Blacksburg, Virginia

Keywords: Jamming, Electronic Warfare, Jammer Exploitation, Machine Learning,
Cognitive Radio, Wireless System Security

Antifragile Communications

Marc L. Lichtman

(ABSTRACT)

Jamming is an ongoing threat that plagues wireless communications in contested areas. Unfortunately, jamming complexity and sophistication will continue to increase over time. The traditional approach to addressing the jamming threat is to harden radios, such that they sacrifice communications performance for more advanced jamming protection. To provide an escape from this trend, we investigate the previously unexplored area of *jammer exploitation*.

This dissertation develops the concept of antifragile communications, defined as the capability for a communications system to improve in performance due to a system stressor or harsh condition. Antifragility refers to systems that increase in capability, resilience, or robustness as a result of disorder (e.g., chaos, uncertainty, stress). An antifragile system is fundamentally different from one that is resilient (i.e., able to recover from failure) and robust (i.e., able to resist failure). We apply the concept of antifragility to wireless communications through several novel strategies that all involve exploiting a communications jammer. These strategies can provide an increase in throughput, efficiency, connectivity, or covertness, as a result of the jamming attack itself. Through analysis and simulation, we show that an antifragile gain is possible under a wide array of electronic warfare scenarios. Throughout this dissertation we provide guidelines for realizing these antifragile waveforms. Other major contributions of this dissertation include the development of a communications jamming taxonomy, feasibility study of reactive jamming in a SATCOM-type scenario, and a reinforcement learning-based reactive jamming mitigation strategy, for times when an antifragile approach is not practical.

Most of the jammer exploitation strategies described in this dissertation fall under the category of jammer piggybacking, meaning the communications system turns the jammer into an unwitting relay. We study this jammer piggybacking approach under a variety of reactive jamming behaviors, with emphasis on the sense-and-transmit type. One piggybacking approach involves transmitting using a specialized frequency-shift keying (FSK) waveform, tailored to exploit a jammer that channelizes a block of spectrum and selectively jams active subchannels. To aid in analysis, we introduce a generalized model for reactive jamming, applicable to both repeater-based and sensing-based jamming behaviors. Despite being limited to electronic warfare scenarios, we hope that this work can pave the way for further research into antifragile communications.

(GRANT INFORMATION)

This material is based on research sponsored by the Air Force Research Laboratory (AFRL) under grant number FA9453-13-1-0237 and FA9453-14-1-0222. The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the AFRL or the U.S. Government.

Acknowledgments

I would first like to thank Dr. Jeffrey Reed, for his guidance and support through my time in grad school. I am fortunate for the opportunity to work with someone that has immense experience in the area of wireless communications. He has encouraged me to have faith in my own research ability, and has been a source of motivation and inspiration. I would also like to thank Dr. T. Charles Clancy, for helping me make sure my research was always pointed in the right direction, and planting the seed that led me to my dissertation topic. I am thankful for the pleasure of having my first ECE professor and favorite professor during undergrad, Dr. Jaime De La Ree, as one of my committee members. Additional thanks to my other committee members, Dr. Michael Roan and Dr. Dennis Sweeney, for taking the time and effort needed to serve on my committee. While not part of my committee nor a student, I would like to thank Dr. Vuk Marojevic, for helping out with papers, coordinating proposal writing, and organizing presentations.

Next, I have had the privilege of working with several great students here at Virginia Tech. I would like to thank Dr. SaiDhiraj Amuru for all the help with brainstorming and reviewing papers. Thanks to Raghunandan Rao and Mina Labib for working on LTE related publications with me. Thanks to Matt La Pan and Chowdhury Shahriar for including me in a journal paper when I was early in grad school, and all of the other collaboration. Tad Czauski, I appreciate all the valuable feedback during group meetings, sharing desk space, and for riding electric skateboards with me when we needed a break from work. Thanks to Paul David for always being available to chat about machine learning.

I was lucky to have a research assistantship that was flexible enough to cover the research I wanted to do, and that would not have been possible without Dr. Steven Lane and the Air Force Research Laboratory (AFRL). Thank you Steven for supporting my graduate research through a stable source of funding, and giving me the opportunity to work in Albuquerque for a summer. A special thanks to Paul Tilghman, for allowing me to attend an amazing technology conference hosted by DARPA, and giving me the opportunity to discuss my Ph.D. research with the Secretary of Defense (I'm still confused as to how that happened).

I would also like to thank the Wireless@VT staff, Nancy Goad and Hilda Reynolds, for all the help over the years.

Lastly, I would like to thank my wife, Lindsey, and the rest of my family, who have been a source of love, support, and encouragement during this whole process.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Jamming is an ongoing threat that plagues wireless communications in contested areas. Unfortunately, jamming complexity and sophistication will continue to increase over time. This is in part due to the availability of powerful and low-cost software-defined radios. The current approach to countering such a threat revolves around *jammer detection* and *jammer mitigation*. As such, an increase in jammer complexity requires an increase in countermeasure complexity. This leads to extremely hardened radios that sacrifice communications performance for more advanced jamming protection. To provide an escape from this trend, we investigate the previously unexplored area of *jammer exploitation*. Unlike mitigation (i.e., anti-jamming), the more complex an enemy jammer, the more potential there is for exploitation. It is for this reason that the antifragile paradigm should be applied to wireless communications. A strategy that exploits a jamming attack to provide a communications gain, such as reducing the bit error rate or increasing the rate of communication, can be labeled as "antifragile communications". This is because a gain is achieved by harnessing the presence of a stressor, which in this context is the jammer. This dissertation introduces and investigates the concept of antifragile communications.

1.1 Thesis

In this work, we show that a wireless communications system has the ability to improve its performance by exploiting a hostile jammer, as opposed to simply mitigating the impact of the jammer. The ability to improve system performance in the presence of a system stressor, such as a hostile jammer, is considered an antifragile property. An antifragile system, in general, is fundamentally different from one that is resilient (able to recover from failure) or robust (able to resist failure). An antifragile system is not synonymous with one that is adaptive (able to change to fit some purpose or situation) or cognitive (incorporating artificial computational processes that act like a person), but may incorporate some of these principles in its operation. An antifragile communication system is something has not yet been realized. This dissertation revolves around quantifying potential antifragile gains in the face of an enemy jammer, in order to motivate its implementation at some future date.

The following is my one-line thesis, in the spirit of *lolmythesis.com*:

On occasion, when the enemy jams our radios, we can communicate better than normal.

1.2 Contributions

In this dissertation, I provide the following contributions:

1. Development of a novel *antifragile waveform*, used for manipulating a reactive jammer into relaying information, whereby a gain (relative to a non-jammed case) in throughput, efficiency, connectivity, or covertness can be achieved.
2. Introduction of a generalized model for reactive jamming, applicable to both repeater-based and sensing-based jamming behaviors.
3. Design of a frequency-shift keying (FSK)-based antifragile waveform, used to exploit a sense-and-transmit type reactive jammer that selectively jams active subchannels.
4. Design of a novel approach for reactive jammer mitigation, based on reinforcement learning and Markov decision processes.
5. Formulation of a communications jamming taxonomy, which involves four key capabilities that can define the fundamental behavior of a jammer, as well as a second tier of jamming parameters.
6. Vulnerability analysis of LTE and LTE-A to jamming and RF spoofing.
7. Development of a sense-and-transmit type reactive jammer implementation that selectively jams active subchannels, based on OFDM concepts.
8. Analysis of the feasibility of reactive jamming in a satellite communications scenario, for both uplink and downlink jamming.
9. Investigation into how practical capacity-approaching codes can be abstracted for the sake of simulation using a large table of SNR thresholds.

1.3 Publications

Below is a list of my peer-reviewed publications, with the ones directly relevant to the work presented in this dissertation highlighted in blue.

— M. Lichtman, J. H. Reed, "Reactive Jammer Piggybacking: Achieving Antifragile Electronic Warfare," *IEEE Military Communications Conference (MILCOM)*, Nov. 2016. [Accepted]

[1] M. Lichtman, M. T. Vondal, T. C. Clancy, J. H. Reed, "Antifragile Communications," *IEEE Systems Journal*, 2016.

[2] M. Lichtman, R. P. Jover, M. Labib, R. Rao, V. Marojevic, J. H. Reed, "LTE/LTE-A Jamming, Spoofing, and Sniffing: Threat Assessment and Mitigation," *IEEE Communications Magazine*, April 2016.

[3] M. Lichtman, J. Poston, S. Amuru, C. Shahriar, T. C. Clancy, R. M. Buehrer, J. H. Reed, "A Communications Jamming Taxonomy," *IEEE Security & Privacy*, Feb. 2016.

[4] M. Lichtman, J. H. Reed, "Anomaly-Based Intrusion Detection of Protocol-Aware Jamming," *IEEE Military Communications Conference (MILCOM)*, Oct. 2015.

[5] M. Lichtman, J. H. Reed, "Analysis of Reactive Jamming against Satellite Communications," *International Journal of Satellite Communications and Networking*, May 2015.

[6] M. Lichtman, T. Czauski, S. Ha, P. David, J. H. Reed, "Detection and Mitigation of Uplink Control Channel Jamming in LTE," *IEEE Military Communications Conference (MILCOM)*, Oct. 2014.

[7] J. Kakar, K. McDermott, V. Garg, M. Lichtman, V. Marojevic, J. H. Reed, "Analysis and Mitigation of Interference to the LTE Physical Control Format Indicator Channel," *IEEE Military Communications Conference (MILCOM)*, Oct. 2014.

[8] M. Lichtman, J. H. Reed, "Reinforcement Learning for Reactive Jamming Mitigation," *Journal of Cyber Security and Mobility*, July 2014.

[9] C. Shahriar, M. La Pan, M. Lichtman, T. Clancy, R. McGwier, R. Tandon, S. Sodagari, J. Reed, "PHY-Layer Resiliency in OFDM Communications: A Tutorial," *IEEE Communications Surveys & Tutorials*, Aug. 2014.

[10] S. Dudley, W. C. Headley, M. Lichtman, E. Imana, X. Ma, M. Abdelbar, A. Padaki, A. Ullah, M. Sohul,T. Yang, J. H. Reed, "Practical Issues for Spectrum Management with Cognitive Radios," *Proceedings of the IEEE*, March 2014.

[11] M. Lichtman, J. H. Reed, T. C. Clancy, M. Norton, "Vulnerability of LTE to Hostile Interference," *IEEE Global Conference on Signal and Information Processing*, Dec. 2013.

[12] M. Lichtman, W. C. Headley, J. H. Reed, "Automatic Modulation Classification under IQ Imbalance using Supervised Learning," *IEEE Military Communications Conference (MILCOM)*, Nov. 2013.

[13] T. C. Clancy, M. Norton, M. Lichtman, "Security Challenges with LTE-Advanced Systems and Military Spectrum," *IEEE Military Communications Conference (MILCOM)*, Nov. 2013.

[14] M. La Pan, M. Lichtman, T. C. Clancy, R. W. McGwier, "Protecting Physical Layer Synchronization: Mitigating Attacks against OFDM Acquisition," *Global Wireless Summit 2013*, June 2013.

1.4 Contents

This dissertation begins with a brief background of electronic warfare, and the term antifragile. We then provide a comprehensive communications jamming taxonomy, in Chapter 3, meant to provide additional background into communications jamming.

In Chapter 4, we begin applying the concept of antifragility to wireless communication systems. Specifically, we propose a waveform meant to exploit the presence of a jammer to achieve a communications gain relative to a jammer-free case. This should not be confused with jamming mitigation. Antifragile electronic warfare takes mitigation one level further, by providing a boost during the jamming attack. This chapter is presented using a systems-level point of view, taking into account many different reactive jamming behaviors. In Chapter 5, we focus solely on the sense-and-transmit type jammer model, which monitors a large number of subchannels, and reactively jams each one independently. In other words, the jammer only jams a specific subchannel when it senses energy or a signal on that subchannel. By doing so, the jammer greatly reduces its power consumption, and remains difficult to detect, because it is always hidden under or adjacent to another signal. We introduce an FSK-based waveform design to exploit this type of jammer. Chapter 6 wraps up the antifragile electronic warfare material by discussing antifragile strategies that take place on the network layer (NET). This includes one strategy that can exploit a non-reactive jammer to produce a network-wide coarse timing synchronization signal.

In Chapter 7 we analyze the feasibility of performing and mitigating reactive jamming in a satellite communication (SATCOM) type scenario. The analysis in this chapter can be used to determine in what applications an antifragile waveform is even worth considering. Chapter 8 describes a novel approach for reactive jammer mitigation, based on reinforcement learning and a Markov decision process (MDP), for instances when an antifragile gain is not achievable. Chapter 9 discusses a specific reactive jammer design, based on OFDM concepts. This type of reactive jammer is especially suited for launching a sense-and-transmit type attack against a large number of adjacent subchannels, across a wide bandwidth. By investigating a specific reactive jammer design, we gain additional insights into the possible secondary behaviors the jammer may present. Chapter 10 concludes this dissertation.

Lastly, we note that you can mouse-over most acronyms used throughout this dissertation to view the unabbreviated version (although not all PDF readers support this).

Chapter 2

Background

2.1 Electronic Warfare

Electronic warfare (EW) is any action intending to preserve the use of the electromagnetic spectrum for yourself and allies or to deny its use to the enemy [15]. EW also encompasses directed energy weaponry, which are weapons that emit a high amount of focused energy toward a target, in order to damage it. Even though directed energy falls under EW (because it involves using electromagnetic radiation), we will not be discussing it in this section, and instead focusing on the control of spectrum. EW is most often applied to radar and wireless communications, and its applications span land, air, sea, and space. Traditionally, EW was split into the following three categories [16]:

1. Electronic countermeasures (ECM)- primarily jamming or causing deception
2. Electronic counter-countermeasures (ECCM)- primarily anti-jamming
3. Electronic support measures (ESM)- primarily the search, interception, identification, and/or localization of sources of electromagnetic energy

However, the current accepted terminology, at least within NATO, is Electronic Attack (EA), Electronic Protection (EP), and Electronic Support (ES). These three subdivisions correspond to ECM, ECCM, and ESM respectively, and their definitions have no significant difference in meaning.

At the heart of EW is jamming, which can be performed against radar and communications. Jamming is defined as deliberately blocking or interfering with signals involved in a radar (including heat-seeking weapons) or communications system. Jamming can be used to remotely disable a radar, or simply fool the radar into thinking it sees something that is not there. In the former case, this is achieved by saturating the radar's receiver with noise. Within wireless communications, jamming is used to simply deny or degrade the transfer of information, usually by injecting noise into the destination node's receiver. In Chapter 3 we dive into the communications side of EA, through the formulation of a communications jamming taxonomy, as well as an investigation into jamming attacks against LTE.

EP is any action taken to counter an electronic attack such as jamming. This protects personnel, facilities, and equipment from enemy EA. EP capabilities are included in the design and implementation of "protected" radars or radios. This kind of protection can take many different forms, and is usually tailored to a certain form of EA. An example of EP is using a phased array receive antenna with adaptive nulling, which can spatially null out interference, as long as it does not originate from the same location as the communications signal of interest. A novel EP technique used to protect against reactive jamming is described in Chapter 8.

ES, on the other hand, exists in order to perform threat recognition, targeting, and planning of operations. If a system's ES fails in the face of a sophisticated EA, then one solution is to kinetically destroy the source of EA, which means identifying the signal and locating the

source. No additional discussion on ES is included in this dissertation, as literature related to signal detection, classification, and localization is significantly more abundant than literature on EA and EP combined, because it is used in more applications than EW.

Antifragile Electronic Warfare (AEW), the subliminal title of this dissertation, is a concept introduced in [1], which is Chapter 4 of this dissertation. AEW is a step beyond standard EP, occurring when a communications link being jammed actually increases in capability as a result of a jamming attack. We show that this is only possible under certain circumstances such as reactive forms of jamming. Chapters 4, 5, and 6 propose and analyze several antifragile EW strategies.

2.2 The Meaning of Antifragility

Until recently, there has been no word in the English language that expressed the opposite of "fragile". To many, the opposite of fragile is simply robust or resilient. To illustrate why this may not be the case, consider a box with a delicate glass object in it, and on the side is written "please handle with care". If dropped, the glass object has a high likelihood of breaking. A box with a robust object in it may be perfectly fine after being dropped several times, experiencing no change whatsoever. However, in terms of logic, the *exact opposite* of a fragile package would be one that benefits from being mishandled. This package's worst-case scenario would be one in which it is entirely unharmed. The term antifragile expresses this concept of reverse fragility.

Antifragility is a concept popularized by Nassim Nicholas Taleb and is a term he coined in his 2012 book *Antifragile* [17]. Antifragility refers to systems that benefit from some form (and correct dose) of disorder. In his book, Taleb explains how antifragility is fundamentally different from the concepts of resiliency (the ability to recover from failure) and robustness (ability to resist failure). "The resilient resists shocks and stays the same; the antifragile gets better" [17]. This fragile-robust-antifragile triad is heavily used by Taleb in his book, and is emphasized through Figure 2.1.

The concept itself is nothing new. Consider the ancient Greek myth of Hydra, the multi-headed serpent. For every head chopped off, Hydra would regrow two more heads, causing an even stronger Hydra and a hopeless struggle for anyone other than the hero (usually Hercules). This myth that dates back to 700 B.C.E. embodies the concept of antifragility. A common example of antifragility is the human muscle, which when confronted with fatigue builds strength, exemplifying the saying "what doesn't kill you makes you stronger". An example of antifragility within finance is being able to exploit a volatile stock market for your gain. For more information about the origins of the term, or more examples of the concept, we refer the reader to Taleb's book [17].

In his book, Taleb defines what he calls the "Extended Disorder Family or Cluster": (i) uncertainty, (ii) variability, (iii) imperfect, incomplete knowledge, (iv) chance, (v) chaos, (vi) volatility, (vii) disorder, (viii) entropy, (ix) time, (x) the unknown, (xi) randomness, (xii) turmoil, (xiii) stressor, (xiv) error, (xv) dispersion of outcomes, and (xvi) unknowledge [17]. An antifragile system benefits in some way as a result of disorder, and the purpose of this disorder cluster is to provide various ways in which disorder can manifest itself. Because of the large number of elements in this cluster, we have organized them into categories based on what we believe are similar or synonymous terms, as shown in Figure 2.2.

With the popularization of the antifragility concept has come related literature in various fields. While this includes fields such as finance, business, health, and policy, for this work we are most interested in literature within science and engineering. Table 2.1 organizes peer-reviewed literature within the science and engineering domains that includes the word antifragile or antifragility (as of March 2015). The term must appear in the title, abstract, or

appear heavily in the body of the text. The purpose of this literature review is to investigate how existing literature uses the term "antifragile", and how well its use aligns with Taleb's definitions and intent of the word. We did not include literature that appeared only in prepublication databases such as arXiv, without mention of being accepted in a peer-reviewed conference or journal. The search engines used included Google Scholar, IEEE Xplore, and the ACM Digital Library. Also worth noting is that there exists a workshop within the computer science domain called "Antifragile", which began in 2014, and is part of the International Conference on Ambient Systems, Networks and Technologies (all papers published within this workshop are included in Table 2.1). As mentioned before, the concept itself has existed for thousands of years, and there are likely papers that incorporate the concept of antifragility written before the word was popularized, but unfortunately there are no direct synonyms that can be used for search.

After performing this literature survey, it is clear that the term antifragile has taken on a wide array of meanings, some of which more closely align with Taleb than others. In this section, we are most interested in two things:

1. How existing literature within science and engineering uses the term "antifragile"
2. The purpose of the term antifragile within wireless research, that is, to provide a label to put on a concept that previously had no label

The second item is especially important, as synonyms add very little to technical vocabulary, and in some cases can do harm (e.g., when performing a literature search). To illustrate the wide array of usage of the term antifragile, we provide two examples of the term used within wireless literature; one that could easily be replaced with preexisting vocabulary, and one that uses the term antifragile under Taleb's definition (in our opinion).

For our first example we investigate [43], in which the authors propose an "antifragility metric" for complex networks made up of a large number of nodes. This metric is used to quantify the impact of each node on failure in network connectivity. They define this metric as one that "provide[s] information to the network to be able to predict connectivity failures and adapt itself to avoid them" and "promotes self-learning and self-adaptation of the network" [43]. "The connectivity antifragility metric is designed to identify the most vulnerable connections and quantify the robustness in the network. [43]" Wireless systems that are adaptive and/or self-learning have been discussed in literature for years, and when used in the context of the lower layers it is typically referred to as cognitive radio [47]. In addition, the authors even

Figure 2.1: The fragile-robust-antifragile triad is heavily used by Taleb.

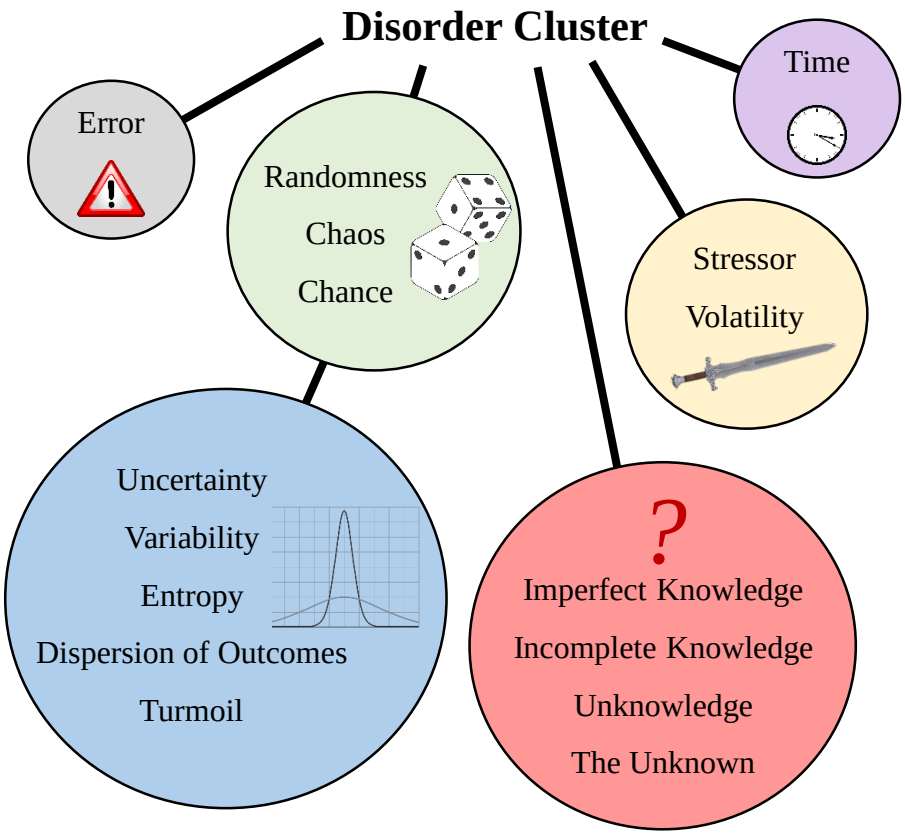

Figure 2.2: Our version of Taleb's Extended Disorder Cluster.

Table 2.1: List of peer-reviewed literature within science and engineering that includes the word “antifragile” or “antifragility”.

	Paper Title	**Type of Paper**
[17]	Antifragile: Things that gain from disorder	[Book]
[18]	Antifragile Systems [Reflections]	[IEEE Magazine]
[19]	Antifragility = Elasticity + Resilience + Machine Learning Models ... Open System Fidelity	[Antifragile 2014]
[20]	Antifragility and tinkering in biology- flexibility provides an epigenetic way to manage risk	[Journal]
[21]	Applying systems and safety engineering principles for antifragility	[Antifragile 2014]
[22]	Engineering Antifragile Systems: A Change In Design Philosophy	[Antifragile 2014]
[23]	The antifragile organization	[ACM Magazine]
[24]	Toward antifragile cloud computing infrastructures	[Antifragile 2014]
[25]	Promoting the confluence of tropical cyclone research	[Journal]
[26]	Toward Anti-fragility: A Malware-Halting Technique	[IEEE Magazine]
[27]	An initial approach towards the implementation of human error ... antifragile systems	[IEEE Conference]
[28]	Flexible and efficient aggregation framework for antifragile wireless mesh networks	[Springer Journal]
[29]	The chloroplast genome hidden in plain sight, ... anti-fragile distributed data sources	[Journal]
[30]	Anti-Fragile Information Systems	[Conference]
[31]	Stress to impress	[Wiley Magazine]
[32]	Detecting antifragile decisions and models ... of aging vehicles	[IEEE Conference]
[33]	Value and context in data use: Domain analysis revisited	[Wiley Conference]
[34]	Model engineering for cyber complex adaptive systems	[Conference]
[35]	Positioning antifragility for clouds on public infrastructures	[Antifragile 2014]
[36]	On environments as systemic exoskeletons: crosscutting optimizers and antifragility enablers	[Springer Journal]
[1]	Antifragile Communications	[IEEE Journal]
[37]	Designing for an Innovative Learning Organization	[IEEE Conference]
[38]	Fostering Progress in Performance Evaluation ... Robotic and Automation Systems	[IEEE Magazine]
[39]	Identifying signs of systems fragility: A crowdsourcing requirements case study	[IEEE Conference]
[40]	Software theory change for resilient near-complete specifications	[Antifragile 2015]
[41]	A framework for trustworthiness assessment based on fidelity in cyber and physical domains	[Antifragile 2015]
[42]	Automatic resource allocation for high availability cloud services	[Antifragile 2015]
[43]	Resilience and Knowledge in a Metric for Heterogeneous Wireless Connectivity	[Technical Report]
[44]	On resilient behaviors in computational systems and environments	[Springer Journal]
[45]	A Proposal for an Antifragile Software Manifesto	[Antifragile 2016]
[46]	Taxonomy and issues for antifragile-based multimedia cloud computing	[Springer Journal]

admit that their antifragile metric is an indication of robustness in the network, which goes against the fragile-robust-antifragile triad. The authors justify using the term antifragile because the network is able to learn from failures and adapt itself, but overall the authors use the term antifragile synonymously with resilient and robust.

Our second example involves an opinion article [18], in which the author reflects on electronic systems we have today that could be considered antifragile, including power supplies that harvest energy from random vibrations. These power supplies benefit from chaos, which seems to fit Taleb's definition of the term antifragile. The author also points out how for the first half of the last century, multipath phenomena was thought of as harmful. When signals traveling along different paths meet at the receiver, they often interfere destructively, cause intersymbol interference, and can cause link degradation. However, multiple-input multiple-output (MIMO) techniques, which uses multipath to enhance system performance, have been created that exploit multipath fading and randomness in the wireless channel. Under this logic, MIMO can be considered an example of antifragile communications. We believe that [18] uses a definition of antifragility that closely aligns with Taleb's.

Another great example of antifragility is in [24], where the authors discuss how the antifragile concept inspired new techniques in cloud computing, which typically involves a large number of virtual machines. After considering how biologic systems require perturbations and disorders to evolve and gain maturity, the authors experimented with the idea of randomly killing off virtual machine (VM) instances within a compute cloud. When a VM has a failure or is killed, a new and healthy VM is created to take its place. It would be very taxing on the cloud system to be constantly restarting VMs, but a healthy dose of randomized VM killing turned out to be beneficial.

There are several papers that claim their system is antifragile because it can benefit from failures, but only in terms of how it is able to deal with or prevent future failures. For now we will view this as "pseudo-antifragile". If there were never failures, the network would perform at its best, and no number of failures will cause this system to perform better than this baseline. This is in contrast to the techniques proposed in this dissertation, which cause a wireless link to increase above the baseline during a jamming attack. In the case of these pseudo-antifragile systems, it is difficult to argue that they really benefit from failure. One possible exception is illustrated by an Antifragile 2014 paper [19], which claims their system is antifragile because the more their system is subjected to threats and challenging conditions, the more insight it will acquired on how to respond to *new and more threatening* situations (either for itself or at the genotypic level). It could be argued that if the system never experienced minor threats (which may be quick to overcome), it would not do as well at combating severe threats that were not anticipated at the time of design.

One reason for the large number of papers misusing the term may stem from the fact that there was a workshop soliciting papers related to antifragility, and it is not uncommon to take existing work and attempt to mold it into a certain theme, even if it does not suite that theme. In addition, the workshop's call for papers may have been counterproductive in promoting the

term antifragile under Taleb's definition. Several topics listed include the phrase "...resilient and antifragile...", as if resilient and antifragile are synonymous. Some papers submitted under this workshop unarguably use the term antifragile to mean adaptive. For example, [22] describes several aerospace related components that were once static, but now are being designed to be adaptive. Other papers, such as [42], do not claim to be antifragile, but still appear in the Antifragile workshop for some reason. Regardless, a major purpose of the workshop is to determine whether it makes sense to apply antifragility to computing systems in the first place, and how to evolve modern engineering practices in order to develop systems that are antifragile by construction. There is no doubt that this is an open question, and the establishment of an antifragile workshop will lead to more researchers tackling the problem.

We have seen that the term antifragile is often used synonymously with self-learning. While this may make sense in certain areas of research, the area of wireless communications already has a widely used term for this concept: cognitive radio (for more information see Chapter 4). Introducing another term into wireless communications that is synonymous with cognitive radio adds nothing. However, the term antifragile can be applied to wireless research in a much more Taleb-centric way, as seen in the remainder of this work (in the form of stressor exploitation), as well as some of the aforementioned references.

In this section we have discussed the ways Taleb defines antifragile, and the manner it has been used in literature. We conclude this section by stating our own definition of the term antifragile, which is more tailored to use within the context of science and engineering, and addresses some common misuses of the word. We define an antifragile system as one that has the potential to exhibit an increase in performance, capability, resiliency, or robustness as a result of stress, randomness, harsh conditions, volatility, or uncertainty. Notice that we have left out "time", as many engineering systems have mechanisms in place to allow an increase in performance over time. We have left out "error" because, by definition, and error is something wrong, and if it made an improvement then it was not an error. Lastly, we note that the material in this dissertation is certainly not a perfect embodiment of antifragility. Antifragility within engineering systems is a great challenge, as engineering systems by definition will not benefit from most of the items on Taleb's disorder cluster. For this research, we have simply used the concept of antifragility as inspiration for the techniques developed during the course of this research, and as a way to view the problem of anti-jamming from a different perspective. We hope that this work can pave the way for further research into antifragile communications.

2.3 Signal Power, Energy per Bit, and Noise Power

Throughout this dissertation one can find many references to SNR, energy per bit (E_b), and noise spectral density (N_0). In this section we briefly review the meanings of these metrics/parameters that often appear in link budgets and information theory. Researchers within wireless communications may prefer to skip this short section.

Let us begin with E_b, which is the energy per bit, in units of joules (or J). A joule can be defined as one watt of power for one second, or one watt-second (W·s). Often times E_b means the energy per *user* bit or *information* bit, so that overhead such as forward error correction can be taken into account. For an example of the usefulness of E_b, one can consider a BER vs E_b/N_0 curve, which shows BER irrespective of the data rate and bandwidth used. Note that E_b almost always refers to *received* energy per bit, not transmitted.

E_b is most often confused with either energy per symbol, E_s, or received signal/carrier power, often denoted as S, C, or P_R. Energy per symbol is simply the energy per bit multiplied by the bits per symbol, or $\log_2 M$ where M is the number of different modulation symbols, e.g., four for quadrature phase-shift keying (QPSK). If E_b is the number of information bits per symbol, then $E_s = E_b r \log_2 M$ where r is the code rate. Signal power, on the other hand, is measured in watts and is found by multiplying the energy per symbol with the number of symbols per second (or symbol rate), usually denoted as R_s or R_s. When dealing with a single modulated carrier, which is often the case in wireless communications, signal power and carrier power are synonymous.

Often confusingly referred to as noise power, N_0 is the noise spectral density, which is the noise power over 1 Hz of bandwidth, measured in watts per hertz (W/Hz) or joules. If it was simply noise power then it would be measured in watts. Because E_b and N_0 have the same units, the ratio E_b/N_0 is dimensionless.

N_0 should not be confused with noise power. In link budgets, noise power, often denoted as N, P_n, or P_{noise}, represents the noise floor and is compared to the signal power. P_n is often found using the "kTB method", where Boltzmann's constant k (J/K) is multiplied by noise temperature T in kelvin (K), and then by channel bandwidth B (Hz). The result is in units J·Hz, J/s, or simply, watts. N_0 is simply equal to kT (as long as the noise is one-sided white noise), and it can be converted to P_n by multiplying by bandwidth. Two sided power spectral density is typically expressed as $N_0/2$. Equations for received signal typically contain a noise term n, and when it represents AWGN, n has zero mean and variance $\sigma^2 = N_0/2$.

While E_b/N_0 may be useful from an information theoretic perspective, link budgets typically are concerned with SNR, which is the ratio of two powers (in watts), and is thus dimensionless. Without knowing anything about the specific signal being received, one can estimate SNR, as long as observations of noise power N are found before the signal appears.

Single carrier digital signals often use pulse-shaping to limit the signal's bandwidth. A common pulse-shaping filter is the raised-cosine (RC) filter. The main parameter of a RC

filter is the roll-off factor, denoted as α or sometimes β, and is between zero and one, and often around 0.25. The roll-off factor measures the excess bandwidth of the filter, relative to the Nyquist bandwidth, for a given symbol rate R_s. The Nyquist rate is often thought of as providing the minimum sampling rate required to sample a band-limited signal (for a signal of bandwidth B, you must sample at $2B$ or higher). However, the same theorem provides an upper bound for the symbol rate across a bandwidth-limited channel. In other words, for a baseband signal, the bandwidth must be at least $B > R_s/2$. To achieve $B = R_s/2$ the pulse-shaping filter must be of infinite length, which is not practical, so we sacrifice extra bandwidth to produce a practical filter. The excess bandwidth parameter α is now included, and the bandwidth equation becomes $B = (R_s + \alpha R_s)/2$ or $B = R_s(\alpha + 1)/2$. When the signal is moved to RF, we must take into account the region $[-B, B]$ instead of just $[0, B]$, so the equation becomes $B = R_s(\alpha + 1)$. This expression is useful because it allows us to approximate the noise power, which requires knowing the channel bandwidth B, for a given data rate (recall that the signal power itself does not require knowing the bandwidth). Note that this assumes the receiver filters out all frequencies other than the signal of interest which occupies bandwidth B. If the receiver has poorly designed filters, then the noise power will be excessively large.

The following are the most important equations that appear in this section, for reference:

$$E_s = E_b \log_2 M \qquad \text{or} \qquad E_s = E_b r \log_2 M \tag{2.1}$$

$$S = C = P_R = E_s R_s \tag{2.2}$$

$$N_0 = kT \tag{2.3}$$

$$N = P_n = P_{noise} = kTB = N_0 B \tag{2.4}$$

$$B = R_s(\alpha + 1) \tag{2.5}$$

2.4 Modeling Practical Capacity-Approaching Codes

In this section we propose a method for modeling practical capacity-approaching codes, such as turbo and LDPC codes, so that channel coding can be taken into account during simulation in a computationally efficient manner. This method can be used as part of a system-level communications simulation framework in order to abstract a large portion of the physical layer (PHY).

Link budgets are very common in analysis and simulation of wireless communications systems; they are most often used to determine SNR given a set of system parameters. However, once a communications system is created, the metric of most interest is often throughput, usually in bits per second. That raises the question of how to link a SNR with throughput for a given system. When we assume noise is Gaussian, the Shannon limit tells us how many bits per second can be reliably communicated given a certain bandwidth and SNR. Despite the fact that many modern channel coding techniques (a.k.a., forward error correction) are labeled as "capacity approaching", using the Shannon limit during numerical simulation is not always sufficient. In addition, a real system is not able to achieve capacity approaching performance at any given SNR, because it uses a limited amount of modulation and coding schemes. For example, most modern communications protocols use a small set of modulation schemes, such as QPSK, 16-quadrature amplitude modulation (QAM), and 64-QAM in LTE. Similarly, a given protocol will only support a certain set of code rates. Because the decoder takes such a large portion of silicon resources (e.g., space on an FPGA or ASIC), puncturing and repetition techniques can be used to achieve many different code rates for a single decoding implementation in silicon. However, using puncturing and repetition can reduce performance compared to implementing a separate decoder for each code rate.

When paired together, each modulation and coding scheme achieves a certain data rate. As an example, we provide LTE's standard modulation and coding schemes (MCSs) for the downlink shared channel, taken from Release 8, shown in Table 2.2. The downlink shared channel uses turbo codes of various rates. By multiplying the modulation order and code rate for each MCS, we find the number of bits carrying information for each research element. We see that the designers of LTE have used fairly uniform steps between each MCS. Channel quality indicator (CQI) index is proportional to SNR, while MCS Index is simply a way for 3GPP to have an index for each combination of modulation scheme, coding scheme, and number of spatial streams (only rows corresponding to one spatial stream are shown in the table).

While a set of MCS tells us the achievable discrete levels of throughput, it does not indicate at what SNR each MCS requires for reliable communications. This information is typically acquired through simulation or real-world testing, by finding the bit error rate curve for each MCS. As an example we have provided a copy of a block error rate curve (BLER) produced by NIST as part of an investigation into the performance of the LTE physical layer [48], shown in Figure 2.3. These are simulation results that take into account the ETU70 channel

Table 2.2: LTE's Modulation and Coding Schemes (MCSs) for the Downlink Shared Channel.

CQI Index	MCS Index	Modulation Order	Modulation Rate	Code Rate	Info. Bits per Resource Element
1	-	QPSK	2	0.076	0.15
2	0	QPSK	2	0.12	0.24
3	2	QPSK	2	0.19	0.38
4	4	QPSK	2	0.3	0.6
5	6	QPSK	2	0.44	0.88
6	8	QPSK	2	0.59	1.2
7	11	16-QAM	4	0.37	1.5
8	13	16-QAM	4	0.48	1.9
9	15	16-QAM	4	0.6	2.4
10	18	64-QAM	6	0.46	2.8
11	20	64-QAM	6	0.55	3.3
12	22	64-QAM	6	0.65	3.9
13	24	64-QAM	6	0.75	4.5
14	26	64-QAM	6	0.85	5.1
15	28	64-QAM	6	0.93	5.6

model, which is an urban multipath fading channel model.

BLER is effectively the BER when taking into account channel coding, which is performed in blocks. If a block is received in error, then the entire block is discarded and a retransmission may be performed if the information is not real-time. Thus, a BLER curve can help provide accurate throughput for a given MCS and SNR.

LTE performs adaptive modulation and coding to attempt to cause the BLER to be less than 10%, while maximizing throughput. Using this logic, one could say that the minimum SNR needed for each MCS is simply where each curve intersects the 10% mark. However, looking at Figure 2.3 (or any BLER curve of a modern coding scheme), we see that the curve is very steep within the usable region (e.g., less than 10% BLER). What this means is that at a certain SNR, the link goes from being almost completely error-free to an error rate high enough to cause a broken link, within roughly 3 dB for most of the curves shown in Figure 2.3. It is this feature that we will take advantage of in our approach to modeling channel coding.

The question we seek to answer is how to take into account these performance curves in system-level analysis and simulation of communications systems. Simulating the channel coding and decoding process is extremely computationally intensive. Another option is to use previously found BLER curves directly, as a sort of lookup table. However, we believe that this is not necessary considering how steep the curves are for modern coding schemes, unless extreme precision is needed.

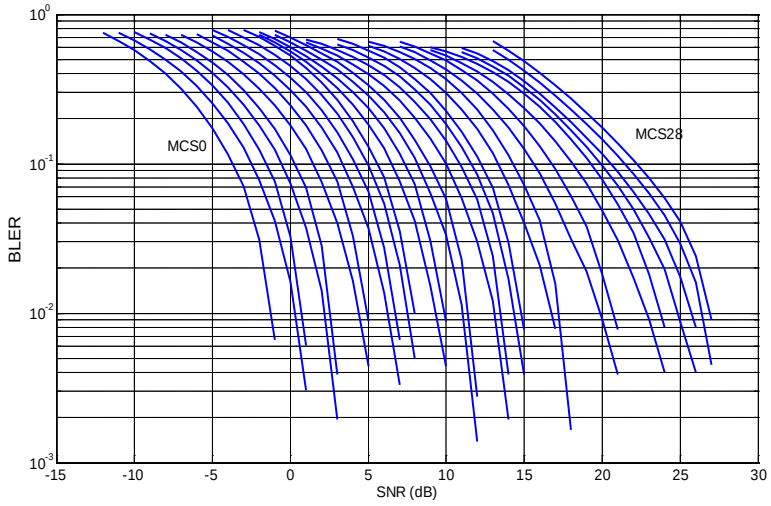

Figure 2.3: Copy of a block error rate curve (BLER) produced by NIST as part of an investigation into the performance of the LTE physical layer [48].

In order to abstract a large portion of the PHY during simulation, we propose a method of modeling capacity-approaching codes, such as turbo and LDPC codes, whereby the BLER curves are approximated as vertical lines. This translates into each curve corresponding with a SNR threshold that indicates the point at which the link goes from zero error to total error (i.e., 50% BER or 100% BLER). An example is shown in Figure 2.4, using three different MCSs whose BLER curves were placed arbitrarily, but match the typical curve shape of LDPC and turbo coding schemes. We chose to place the vertical lines toward the center of each curve, as opposed to where the curve hits 10%, so as to reduce the sum error introduced by this model. The SNR threshold for each MCS is simply where the dotted red line hits the SNR axis.

This abstraction approach can be directly used in a simulation framework, through a table of SNR thresholds for a large number of MCS schemes. Ideally, this table would have hundreds of entries, to help smooth out curves generated through simulation. The table is used to find the highest possible data rate that has an SNR threshold below the SNR being queried. In creating the table, one does not have to worry about only using the top performing coding schemes, because they simply will never be chosen. What is most important is that the BLER curves used to create the SNR thresholds and table are accurate, and observed under conditions that go along with the table. For example, there could be a table corresponding to an urban multipath fading channel. LDPC BLER curves are widely available online due to their use in protocols such as DVB-S2, and turbo codes are gaining popularity in literature due to their use in LTE.

In this dissertation we make use of the approach described in this section in the numerical results provided in Section 4.7. The throughput versus SNR (or JSR) curves that we show for the primary results are characterized by their staircase-like shape, showing sharp transitions

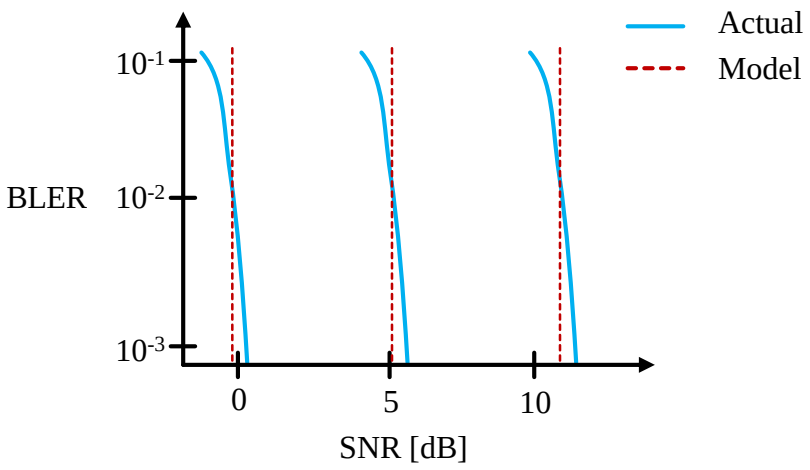

Figure 2.4: Depiction of proposed approach to modeling channel coding using a SNR threshold.

every time the MCS ratchets up or down. A larger table would help smooth out this staircase pattern. While we considered using interpolation to entirely smoothen the curve, we chose not to, because it reduces the main advantage gained by using the proposed modeling (i.e., the fact that a realizable MCS scheme exists that allows communications at the given SNR).

While the method of abstraction described in this section is fairly simple, we believe it has a lot of potential for allowing simple yet effective simulation frameworks for wireless communications systems. Treating each MCS with a minimum SNR threshold is nothing new, but in this section we have investigated what it really means to do so, and why it should be a valid method of modeling high-performance coding schemes.

Chapter 3

A Communications Jamming Taxonomy

A portion of the material in this chapter has been previously published in [3] and [2], and is being reproduced in this dissertation with the consent of all co-authors, as well as the original publisher if required.

3.1 Introduction

The inherent openness of the wireless medium makes it susceptible to adversarial attacks. The vulnerabilities of a wireless system can be largely classified based on the capability of an adversary: a) an eavesdropping attack in which the eavesdropper (passive adversary) can listen to the wireless channel and try to infer information b) a jamming attack in which the jammer (active adversary) can transmit energy to disrupt reliable data transmission and c) a higher-layer active attack that threatens *integrity* and *confidentiality* of a link. In this work, we study jamming attacks, principally those at the physical layer (PHY), intended to cause Denial of Service (DoS) to one or more users, thus compromising the *availability* of a link.

Jammers may employ a wide range of behaviors to cause DoS, and a sampling of literature related to jamming will show numerous jamming models and assumptions. These models or behaviors can span in complexity from a constant source of continuous wave interference to an intelligent jammer that has the capability of sensing and making decisions in real-time to increase effectiveness and covertness of the attack.

In this chapter, we propose a taxonomy that covers the communications side of jamming (as opposed to radar jamming or attacks against radio navigation). Research on electronic warfare (EW) and jamming dates to World War II, an era when jammers needed to be categorized by signal type, because each signal type had to be constructed from distinct radio circuitry. In the present era of software-defined radio (SDR), however, the historical approach burdens the understanding of jamming with unnecessary restrictions. Today, the important questions to answer are what information does the jammer possess and what is the jammer's capacity to act on that information.

The intent of this jamming taxonomy is to help researchers place newly discovered jamming or anti-jamming strategies within a larger context of known strategies in a way that is consistent with modern electronic warfare.

Closer to the technological theme of this taxonomy, the Common Attack Pattern Enumeration and Classification (CAPEC) [49] is a catalog and taxonomy of cyber-attack patterns, created to assist in the building of secure software. Each attack pattern provides a challenge that the attacker must overcome, common methods used to overcome that challenge, and recommended methods for mitigating the attack. The taxonomy is organized at the top-level by mechanisms of attack (e.g., abuse of functionality, exploitation of authentication, malicious code execution) and domains of attack (e.g., hardware, software, social engineering). While the jamming taxonomy proposed in this chapter is fundamentally different in structure, CAPEC represents a cyber security equivalent.

There are general similarities in the strategies of EW and cyber-attack. An early jamming technique included *barrage* jamming (explained in Section 3.4) that, qualitatively, resembles the approaches of early Internet DoS flooding attacks. More recently, however, both EW and cyber-attack often begin with a reconnaissance phase in order to better understand

the technical characteristics of the target and craft a tailored attack. The EW literature, reflecting its military heritage, referred to this preliminary stage as *signals intelligence* (SIGINT). As holds true for cyber-attacks, jamming can serve a larger purpose than just denying communications. For example, it could deny wireless users access to a network with strong authentication and privacy mechanisms, but permit association with another network having inadequate security measures, thereby setting the stage for breaches of confidentiality, integrity, and identity. The full discussion of all of these parallels between EW and cyber-attack, however, is beyond this chapter's scope of introducing a new jamming taxonomy.

As this taxonomy only covers jamming, we distinguish between a jammer and cyber-attack based on the intended mode of failure, and the type of attack vector used by the adversary. Traditionally, jamming is performed using an RF vector while a cyber-attack is launched through a network vector. The blurry area occurs when dealing with reactive jamming, described in Section 3.3.1, where the jammer both receives and transmits a signal. We will assume that a jamming signal is *not* a valid frame or packet, because such attacks are rarely classified as jamming. However, we make no limitations to the receiving capabilities of the jammer. For example, the jammer could process the received waveform at the media access control layer (MAC) and network layer (NET), in order to target a certain type of frame or packet. However, to remain within the common definition of jamming, the transmitter portion of the jammer must either inject noise into the communications link, or transmit what looks like a real PHY-layer signal (as discussed in Section 3.3.4). Otherwise the attack should be classified under the cyber security domain. This classification is not meant to limit the capability of jammers, but rather to put a label on a given attack and better define the scope of this taxonomy.

It is important to note that this chapter does not discuss malicious node detection, anti-jamming strategies, jammer detection, or jammer localization. Likewise, it does not cover radar jamming or radio navigation (a.k.a. positioning navigation and timing) jamming or spoofing such as attacks on GPS. The goal of the work is to shed light on the broad characteristics that the jammer may possess and also provide the right references for someone interested in pursuing research related to jamming.

This chapter is organized as follows. Section 3.2 highlights related works. In Section 3.3 we identify key capabilities that distinguish major classes of jamming, which forms the core basis for the jammer taxonomy and also provides a parametric framework that covers the second tier jammer characteristics. Section 3.4 samples several jamming attacks found in open literature. In Section 3.5 we provide a thorough vulnerability analysis of LTE to jamming, showing several specific LTE-based attacks. Finally, conclusions are provided in Section 3.6.

3.2 Related Works

The comprehensive references of Adamy [15] and Poisel [50] for the most part reflect the historical tradition of distinguishing jamming by signal type (e.g., noise, tone, pulse). Poisel's work has more of a communications focus than Adamy's and includes *smart jamming* techniques that, in this taxonomy, we term *protocol-aware* jamming. In contrast to [15, 50], this taxonomy emphasizes behaviors and attributes a jammer could have and then discusses specific jamming techniques characterizing a given behavior. For example, a jammer having one tone versus multiple tones is just expressing an adjustable parameter within the same overall jamming behavior.

Another categorization of jammers is provided in [51], where the authors use the categories of: constant jammer, deceptive jammer, random jammer, and reactive jammer. The authors describe a constant jammer as one that sends out random bits without following any MAC-layer protocols. Their deceptive jammer (termed as *spoofing* in this chapter) is one that transmits regular packets into the channel, following the PHY and MAC layer protocol used by the target. They define random jamming as the jammer turning on and off with a random or fixed period. Lastly, reactive jamming is a jammer that senses the target channel and only transmits when there is activity. While these categories are well-suited for the analysis performed by the authors, they omit distinctions for whether or not the jammer is adapting its signal based on information it has *a priori* or has acquired. A similar concern applies to jamming literature surveys such as [52].

As noted above, most of the previous work only studied specific aspects of the jamming problem and did not provide a complete overview of the potential jamming attacks that can be performed depending on the information available to the jammer. In this regard, our taxonomy is not only more comprehensive than those in the above references, but unique in the sense that it is based on the information the jammer possesses and the jammer's capacity to act on that information.

3.3 Key Jammer Capabilities

The primary delineation of the taxonomy is by jammer capabilities that define the fundamental behavior of the jammer. A secondary refinement of the taxonomy by parameters is explained in the next section. A jammer can have one or more of the following major capabilities:

1. Reactive a.k.a. Time Correlated
2. Protocol-Aware
3. Ability to Learn
4. Signal Spoofing

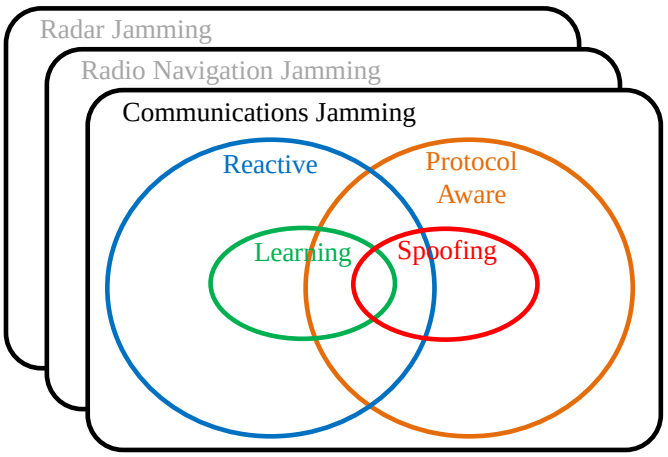

Figure 3.1: Key capabilities of a jammer and how they relate.

Figure 3.1 shows how the jammer capabilities are interrelated.

These four capabilities were chosen based on a survey of jammer models that exist in literature, with an emphasis on complex forms of jamming. For example, a learning jammer (a.k.a. cognitive jammer) may not represent the majority of what is found in current-day operations, but it is a topic of interest in recent research and will likely become more prevalent over the next decade. For reactive jamming, we are referring to correlation in the time domain specifically, because it is implicit that a jammer's signal needs to have some correlation in the frequency domain with the victim's desired signal to be successful (i.e., at least be aware of the spectrum being used by the victim and thereby perform jamming attacks over this spectrum). The remainder of this section provides more details about each capability.

3.3.1 Reactive Jamming

A time correlated jammer, known as a reactive jammer in most literature [51, 53–58], is one that transmits a jamming signal that is correlated to the target signal in time, in some fashion. A reactive jammer implies the jammer can listen to the transmitter's signal, leading to the geometrical configuration shown in Figure 3.2. The implementation of this capability may be by alternately receiving then transmitting or, for simultaneous receive and transmit operation, the jammer may cancel its own signal or use separate directional antennas.

This class of jammer could take on a wide range of specific models. For example, it may sense a block of subchannels and jam those that contain energy significantly above the noise floor, or the jammer may retransmit a manipulated replica of what it receives for its attack as in the case of a digital radio frequency memory (DRFM) or repeater jammer. While reactive jamming is a very broad category of jamming, it acts as a good characteristic to quickly identify the complexity of the jammer, as a reactive jammer must have some form of

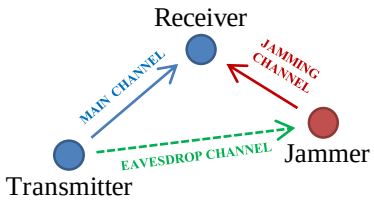

Figure 3.2: Geometrical configuration of a reactive jamming scenario, showing the three channels involved.

a receiver. Because there is significant engineering that goes along with receiving capability (e.g., a full RF chain, sampling, processing), any reactive jamming attack corresponds to a more complex attack. For the remainder of this chapter we will refer to a jammer that isn't reactive as "non-correlated".

One may ask how jamming is possible without a receiver; how does a jammer know which signals to jam? When we discuss a jamming attack, we are referring to the specific attack being launched against a signal. Before any attack is launched, the following steps must occur (see below). The reactive capability comes into play during the actual attack, not the signal awareness step. Obtaining signal awareness surely requires receiving capability, but a reactive attack consists of the jammer being tightly synchronized to the target signal.

1. **Signal Awareness:** sensing and detecting signals across the spectrum of interest
2. **Threat Assessment:** for each signal, a decision must be made as to whether or not it will be jammed
3. **Attack Selection:** for each signal to be jammed, the best attack must be selected

We continue to discuss reactive-type jamming in Section 4.4.2.

3.3.2 Protocol-Aware

The term protocol-aware simply means the jammer is aware of the protocol of the target signal. Information about the signal's protocol is obtained during the **Signal Awareness** step and used in the **Attack Selection** decision-making. For example, the jammer may identify that a particular signal is a Wi-Fi or LTE signal, which due to the open nature of specifications allows the jammer to know almost everything about the PHY and MAC layers. A jammer could use *a priori* knowledge of the protocol to exploit weaknesses in the protocol, and launch a jamming attack that is more effective and may be harder to detect than non-protocol-aware jamming. We note that a signal does not have to belong to a specific technology to be open to a protocol-aware attack. For example, the jammer may only know a signal uses orthogonal frequency-division multiplexing (OFDM) with pilots in certain

locations, which would be considered protocol-aware if it knew exactly where the pilots were placed. Although this work is focused on jamming, we note that essentially all cyber-attacks must be protocol-aware to some extent (by definition).

As discussed in literature, if a jammer knows the specific protocol being used, it can increase in effectiveness by jamming a PHY or MAC layer mechanism instead of data directly (sometimes referred to as *selective jamming*). In most wireless protocols, the data takes up the largest portion of time and frequency resources. Thus, if a jammer targets something besides the data, it will likely result in an attack that uses less power and is harder to detect (as long as the targeted mechanism is essential for communications). Possible mechanisms that could be targeted in a protocol-aware attack (taken from open literature) include:

- Control channels/subchannels
- Control frames or packets (e.g., ACKs)
- Pilots (a.k.a. reference symbols)
- Synchronization signals
- Cyclic prefix in OFDM

For a survey of protocol-aware jamming attacks against Wi-Fi and LTE we refer the reader to [59] and [11], respectively.

3.3.3 Ability to Learn

In this dissertation, we will define the term "learning" in the machine learning (ML) sense: "systems that can learn from data, rather than follow only explicitly programmed instructions" [60]. A jammer that has the ability to learn is one that may modify its behavior in real-time in response to its experiences (i.e., instances of successful or unsuccessful jamming actions/decisions) [61]. However, a learning system has capabilities beyond an adaptive system that is limited to following a pre-programmed sequence of change in response to stimulus. A simple test to determine if a given jammer has the ability to learn is to see if it evolves its behavior in response to a target's behavior and adaptation. Learning jammers go beyond simply detecting the target's waveform type and choosing from a pre-programmed set of jamming waveforms. Rather, a jammer that learns may detect that the target has initiated an anti-jam strategy, and then the jammer can explore different strategies of its own to circumvent this anti-jam defense. This category corresponds to some of the most complex jammers, for the following two reasons:

1. Learning algorithms (e.g., supervised learning algorithms such as the popular Support Vector Machine (SVM) or artificial neural networks (ANNs)) are complex, with high computational complexity during training.

2. Determining the success of a certain jamming attack may be difficult for the jammer, as it may not have access to the channel feedback information. This is an area where traffic analysis may be used.

Often the ability to learn leads to the label of "cognitive". However, a cognitive jammer that is capable of learning should not be confused with "cognitive radio jamming", i.e., a jammer designed to deny a cognitive radio network (e.g., the primary user emulation attack [62]). In some cognitive radio jamming literature, the term "cognitive jammer" is used, even though the primary user emulation attack rarely involves learning and often is not even reactive.

In some situations a learning jammer may target radios that are also capable of learning, such as cognitive radios in the Mitola sense [63] (as opposed to dynamic spectrum sharing radios). The jammer can exploit this fact using a belief manipulation attack thereby causing the targeted system's adaptation processes to seek a poor operating point [64]. If you can metaphorically convince a radio that "up is down", and "down is up", you can severely impact how it behaves and reacts to particular situations.

In terms of how presence of learning relates to the other key capabilities, a jammer capable of learning is almost surely reactive, because learning involves observing the target signal. We consider learning and protocol-aware as independent features, leading to the relationship shown in Figure 3.1.

3.3.4 Spoofing (a.k.a. Protocol Emulation)

Spoofing is broadly defined as a situation in which one entity successfully masquerades as another by falsifying data and/or signals in order to gain an illegitimate advantage. Typically spoofing targets a PHY-layer mechanism by emulating a signal. In terms of jamming, which is assumed to be a physical layer adversary, we will define spoofing as **the action of transmitting a signal that is meant to look like a legitimate signal**. To distinguish physical layer spoofing from, for example, transmitting fake frames or packets, we will confine spoofing to be on the physical layer. In other words, the spoofed signal need not have any properties that make it look like a valid frame or packet. Rather, the spoofed signal must be intended to fool the signal processing level of the target. Spoofing may or may not be reactive, although in literature it is more often *not* reactive.

Protocol-aware jamming may or may not involve spoofing, but if spoofing occurs then the jammer is almost surely protocol-aware, because it needs to know what to spoof. Determining whether a given adversary is spoofing is rather simple. One must check whether it is transmitting noise, or transmitting something that looks legitimate to the target's PHY layer. The difference between physical layer spoofing and higher layer spoofing is less clear, although if the adversary is transmitting what looks like a valid packet or frame, then the attack is definitely not confined to the PHY layer, and the attack would fall under the category of cyber-attack.

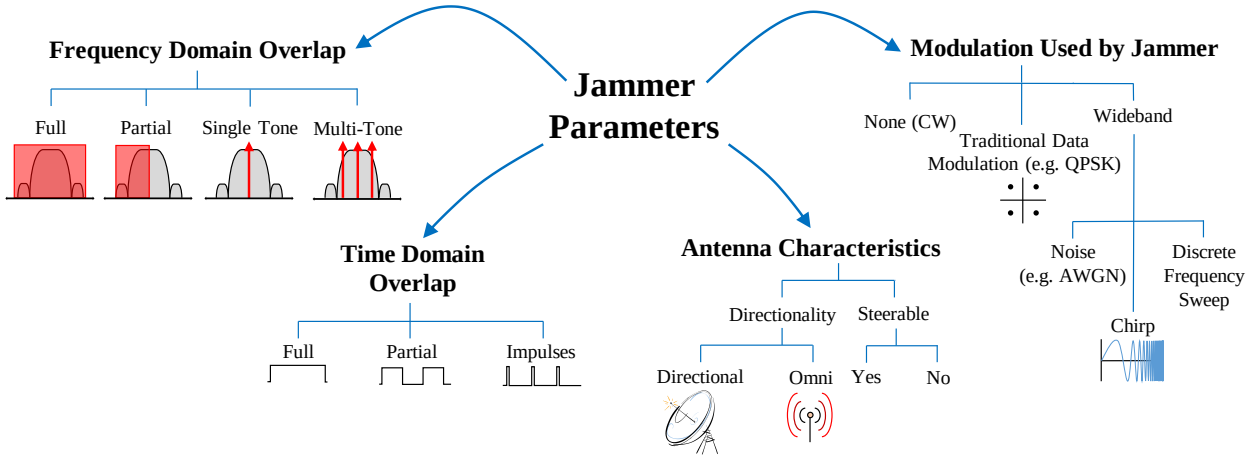

Figure 3.3: Jammer parameters organized into trees.

Cognitive radio jamming techniques, such as primary user emulation, may be considered spoofing depending on the specific waveform the jammer transmits. In forms of primary user emulation where the secondary users only utilize an energy detector, the jammer must only transmit noise at a particular frequency for the secondary users to evacuate the band and avoid using the spectrum. Other forms of primary user emulation could involve the jammer transmitting a signal meant to look like the primary user's signal (e.g., the pilots associated with a radio station), in which case it is PHY layer spoofing.

3.3.5 Jammer Parameters

Building upon the definition of major categories of jammer capabilities, a second tier refinement comes from the choice of physical parameter values. As illustrated in Fig. 3.3, example parameters include frequency, time overlap with jamming target, antenna directionality, and the jammer's waveform or modulation. In this way, jammer types that, in early literature, were treated as distinct technologies can be understood now as minor variations on a common algorithm.

3.4 A Sampling of Specific Jamming Attacks

In this section we provide several example jamming attacks that can be found in open literature, and discuss how they are classified with respect to the taxonomy we have developed in the previous two sections, as shown in Figure 3.4.

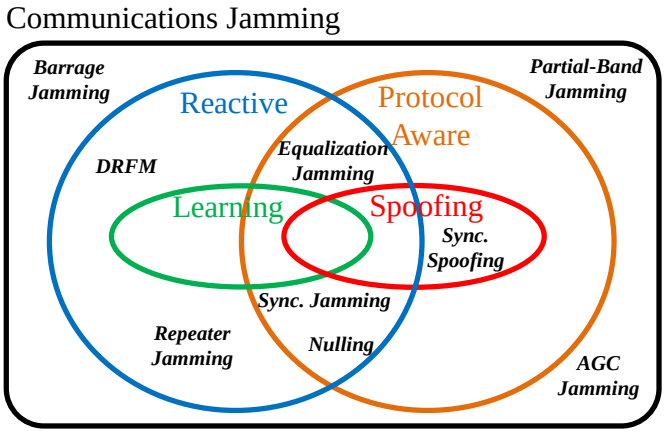

Figure 3.4: Specific jamming techniques discussed in literature, mapped according to key jammer capabilities.

3.4.1 Barrage Jamming

Barrage jamming is the simplest form of jamming and is usually defined as a jammer which transmits noise-like energy across the entire portion of spectrum occupied by the target with 100% duty cycle in time. Thus, it is non-correlated and non-protocol-aware. Barrage jamming has been shown game theoretically and information theoretically to be the best a jammer can do in the absence of any knowledge of the target signal [65].

3.4.2 Partial-band Jamming

When jamming a single-carrier signal, it has been shown that jamming gains can be achieved by not jamming the entire signal in the frequency domain, but rather jamming a fraction of the signal. This is known as partial-band jamming, and it is usually considered a non-correlated jamming attack because the jammer transmits continuously in time. Performing partial-band jamming against an OFDM waveform does not make sense because strong forward error correction could allow the data to be reconstructed from the unjammed subcarriers.

3.4.3 Automatic Gain Control Jamming

The automatic gain control (AGC) mechanism in a receiver adjusts the input gain such that the received signal comes in at a proper level to best utilize the range of the analog-to-digital converter(s). A jamming attack that targets the AGC mechanism is one that uses a very low duty cycle (e.g., 2%) but with extremely high instantaneous power. By not transmitting continuously, the jammer can save power and remain harder to detect in some situations [15]. AGC jamming is non-correlated, although the specific period and duty cycle used are

important parameters. Aside from the assumption/knowledge that the target receiver uses AGC, it is non-protocol-aware.

3.4.4 Equalization Jamming

Equalization jamming involves targeting any mechanism related to equalization. Wireless transmitters insert known data symbols (a.k.a. reference symbols) into the transmitted waveform to estimate the channel's frequency response and equalize the effect of the channel at the receiver prior to demodulation. These known symbols are called pilot symbols in multicarrier communications such as OFDM or single carrier-frequency division multiple access (SC-FDMA) and channel sounding symbols in multiple-input multiple-output (MIMO) systems [66]. For example, in OFDM, pilot tone jamming is simply the process of jamming pilot tones, which may reside on certain subcarriers (in the case of 802.11) or may be multiplexed in time and frequency with data (in the case of LTE). Pilot jamming is protocol-aware because the jammer must know where the pilots are located. If the pilots occur on a dedicated subcarrier then the attack is non-correlated, but if they are multiplexed in time then it must be correlated in order to surgically jam the pilots. It was found that pilot jamming can be energy efficient and similar degradation in target receiver's BER can be achieved using roughly one-tenth of the energy [66]. The pilot jamming process is similar in the case of SC-FDMA, which is the single-carrier variant of OFDM and used in the uplink of the LTE air-interface. In MIMO systems, known reference signals are used for channel sounding and thus can be jammed as well as long as they are known by the jammer *a priori*.

3.4.5 Synchronization Jamming

For a communications link to function, the receiver must synchronize to the incoming signal in both time and frequency. To aid in this task, a synchronization signal, or synchronization symbols, are usually designed into the PHY layer protocol. For example, in LTE there are two different synchronization signals that each appear every 5 ms. Synchronization jamming (a.k.a. synchronization signal jamming) is simply the process of surgically jamming one or more synchronization signals. This jamming technique is unique in the sense that it may only prevent radios from establishing a communications link, and thus it won't cause immediate DoS [11]. However, synchronization signals tend to be very sparse with respect to the entire signal, thus providing a significant jamming gain. Synchronization jamming must be protocol-aware, in order to know where the synchronization signal is located. It must be time-correlated, assuming the synchronization signal is multiplexed in time with data and other signaling.

3.4.6 Nulling

The nulling attack is similar to spoofing, except instead of confusing the target radio with a masquerading signal, the jammer attempts to cause destructive interference in such a way that the received energy at the target receiver is driven as close to zero as possible [66]. It requires transmitting what the target transmitter just sent (with anticipation to take into account propagation delay), and requires phase coherency, so that the jamming signal is received with a π-radian phase shift relative to the target signal. When performed successfully, this nulls out the target signal and leaves only channel noise. Nulling attacks are extremely difficult under a controlled environment, and considered infeasible in real-world scenarios because the jammer must have near-perfect knowledge of the channels involved. This is not likely possible given the varying nature of the wireless channel, and the lack of a mechanism to estimate the channels at the jammer. Even if the target signal includes pilots and synchronization symbols, that would only provide accurate knowledge of the channel between the jammer and transmitter (to perform nulling, the jammer would also need to know the jammer-receiver and transmitter-receiver channels, as shown in Figure 3.2). Nulling also requires *a priori* knowledge of the signal, which is possible in some circumstances (e.g., the value of pilot sequences in Wi-Fi and LTE is openly published).

Even though they are presently believed to be infeasible in practice, nulling attacks are included in this taxonomy due to their presence in academic literature. Pilot nulling against OFDM, which was introduced in [66], involves nulling the pilots at the target receiver. Channel sounding singularity attacks (another name for nulling sometimes used in literature) against MIMO systems has also been investigated.

3.4.7 Repeater Jamming

Repeater jamming (a.k.a. DRFM jamming or follower jamming) is the simplest form of reactive jamming when the jammer has no knowledge of the protocol. In repeater jamming, the jammer transmits when it senses energy on the channel. This may be in the form of the jammer retransmitting what it receives with noise added, or sensing a series of subchannels and transmitting noise when it senses energy on one or more subchannels. Regardless of the specific model used, repeater jamming can "follow" a signal if it hops around in frequency, negating the anti-jam gains associated with frequency-hopping spread spectrum (FHSS).

When there are large distances between the transmitter, receiver, and jammer, a repeater jamming attack may fail to achieve time-correlation with the target signal. However, a repeater jamming attack may still be used in order to decrease probability of detection, because the jamming signal will resemble the target's communications. In addition, simply replaying the target signal can aid in achieving frequency correlation, assuming the target is not hopping or changing channels too quickly. We will refer to this form of repeater jamming, which is too slow to overcome the gain associated with FHSS, as a replay attack.

3.4.8 Protocol-Aware Jamming Against Wi-Fi

There are several intelligent jamming attacks against Wi-Fi (IEEE 802.11) found in open literature [59], most likely due to the popularity of Wi-Fi and length of time Wi-Fi has been widely used. *Clear to Send (CTS) Jamming* is when a jammer waits for there to be a Request to Send (RTS) packet transmitted over the channel, and then transmits a burst of noise after waiting for the Short Interframe Space (SIFS) interval which is defined in the specifications [52]. *Acknowledgment (ACK) Jamming* works the same way, except the jammer waits for a data packet to be transmitted, then after waiting for a SIFS interval it transmits a burst of energy with the intent to jam the ACK [52].

While these previous two attacks are both protocol-aware and reactive, it is possible to have a protocol-aware and non-correlated attack in 802.11, using an attack known as *DIFS Wait Jamming* [59]. This works by transmitting periodic pulses that repeat with a frequency based on the 802.11 DCF Interframe Space (DIFS) duration. This duration determines how long a node senses the channel in order to decide whether or not the channel is idle. Thus, this attack causes a "busy channel" while saving power in a non-correlated manner. A protocol-aware jamming strategy that incorporates learning is proposed in [67].

3.5 Protocol-Aware Jamming Against LTE/LTE-A

Because the vulnerability of LTE to jamming has been a large part of my research, a large section has been added to the taxonomy dedicated to its discussion.

3.5.1 Introduction

The Long Term Evolution (LTE) has been standardized by the 3rd Generation Partnership Project (3GPP) to meet the growing demand in cellular data traffic. LTE offers better coverage, enhanced system capacity, higher spectral efficiency, lower latency, and higher data rates than its predecessors in a cost effective manner. True to its namesake, LTE has been be able to keep pace with the rapid evolution of technology by introducing LTE-Advanced (LTE-A) for even higher peak data rates and capacity, more reliable coverage, and greater spectral efficiency. LTE-A leverages techniques such as carrier aggregation to increase the peak downlink data rate to 3 Gbps. Currently, there are 422 commercially launched LTE networks in 143 countries, out of which 88 operators have commercially launched LTE-A carrier aggregation systems. LTE/LTE-A is unarguably the primary standard for 4G cellular technology and is well on its way to becoming the primary global cellular standard. In addition to mobile user communications, cellular networks are used to broadcast emergency information, announcing natural disasters and other crises. Over the next decade we will likely become further dependent on commercial cellular networks based on LTE, which is

why we must ensure it is secure and available when and where we need it. Unfortunately, like any wireless technology, disruption through deliberate radio frequency (RF) interference, or jamming, is possible.

In the U.S., LTE is being used as a framework for the nationwide public safety network known as FirstNet. The objective of FirstNet is to provide a nationwide interoperable public safety broadband network that provides reliable communications among first responders. Of greatest concern are emergencies caused by an adversary, such as a terrorist organization, whose attack may involve radio jamming against cellular networks (including FirstNet) to ensure disarray and cause further panic. As such, anti-jamming countermeasures need to be considered to ensure communication networks remain available when they are most needed.

The U.S. military has considered using ad-hoc LTE-based networks to keep soldiers on the battlefield connected, as well as for shipborne communication with naval aircraft. Unlike military standards, cellular standards are publicly available, meaning that adversaries may leverage this knowledge and target weak points in the protocol to enhance the efficacy of their attacks. Radio jamming attacks are a serious threat to any military or battlefield communications link and must be accounted for.

Attacks on LTE can be grouped into two broad categories: DoS and information extraction. Jamming attacks are typically used to cause DoS; attacks that extract information or cause DoS by targeting the higher layers fall under the category of cyber-attacks. Radio jamming is broadly defined as an attack in which a jammer transmits energy to disrupt reliable data transmission. Jamming is performed through a RF attack vector, while cyber-attacks use network attack vectors. In this article we are only concerned with jamming. An important property of jamming is that jamming attacks always target the receiver (as opposed to the transmitter), regardless of how close the jammer is to the transmitting node. Thus, jamming the LTE downlink, the signal transmitted by a base station and received by mobile devices, targets the mobile devices, whereas jamming the uplink targets the base station.

Protocol-aware jamming attacks against LTE networks are primarily enabled by the openness of the protocol. Moreover, the broadcast messages transmitted by LTE base stations do not use any means of encryption. As a result, all sorts of essential network configuration details can be easily eavesdropped with low-cost software radios, aiding attackers in optimizing and crafting attacks against LTE-based networks, such as radio jamming and rogue base stations.

The objectives of this article are to outline and motivate the need for high availability LTE networks, provide insight into physical layer vulnerabilities of LTE, and survey mitigation techniques that can harden the LTE physical layer of next generation LTE and LTE-A deployments. The remainder of this article is organized as follows. Section II provides a brief background on the physical layer of LTE. In Section III we investigate the individual channels and signals within LTE, and how vulnerable they are to jamming and spoofing. Section IV offers a comparison of attacks in terms of effectiveness and complexity, Section V surveys mitigation techniques found in literature, and Section VI concludes.

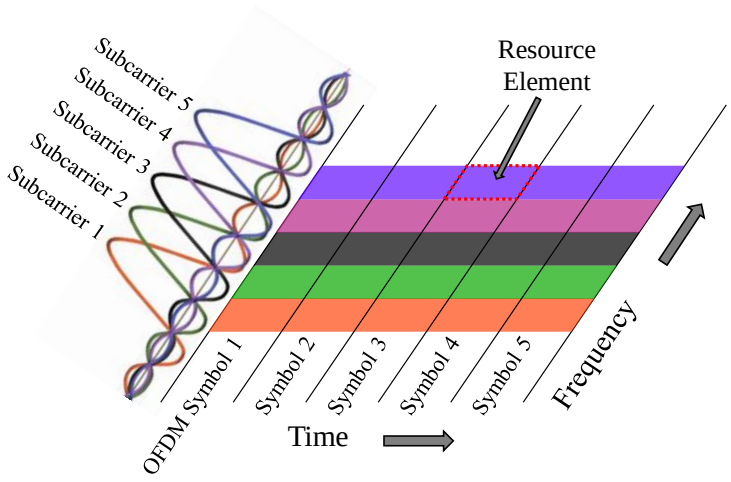

Figure 3.5: Depiction of the OFDM time-frequency lattice

3.5.2 Background of LTE

Orthogonal frequency-division multiple access (OFDMA) is the channel access scheme used in the LTE downlink. OFDMA uses OFDM as the underlying modulation scheme and transmits a large number of parallel subcarriers with different blocks designated to different users. For example, when LTE is configured for a 10 MHz bandwidth (the most common configuration in the U.S.), there are 600 subcarriers in the downlink signal. Each subcarrier carries a separate stream of information, resulting in information being mapped in both the time and frequency domains. This leads to the OFDM time-frequency lattice, which is a two-dimensional grid used to represent how information is mapped to physical resources. In LTE, one subcarrier over one OFDM symbol interval is called a resource element, as shown in Figure 3.5. The entire frame is 10 milliseconds long, and frames repeat continuously.

SC-FDMA is the multiple-access scheme used for the LTE uplink. An equivalent time-frequency grid is defined for the uplink transmission, either in a different band (frequency division duplex mode) or in the same band (time-division duplex mode). However, unlike in OFDMA, information traveling through the air is spread across several subcarriers.

LTE user devices–cellphones, tablets, and dongles, among others–are known as user equipment (UE). The UE accesses the LTE network by connecting to an LTE base station, called an evolved NodeB or eNodeB. A UE attaches to only one eNodeB at a time, but constantly monitors the surrounding cells for the purpose of assisting the network in the handover process, most often performed when a UE is moving and begins to enter a different cell. In addition, UE can sometimes roam (depending on the network's policy) in other 4G, 3G, or 2G networks when their home LTE network is unavailable, which can be considered an anti-jam strategy in itself.

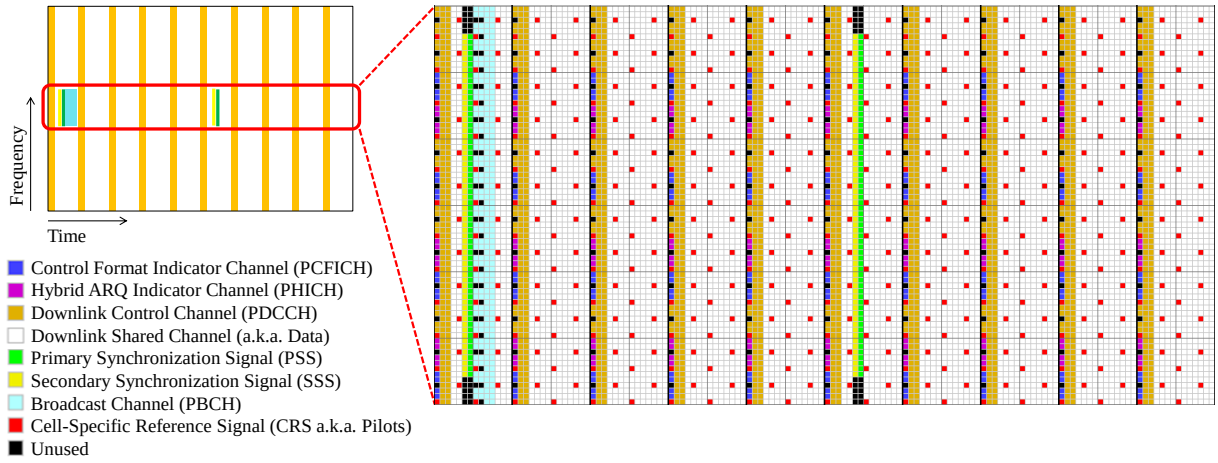

Figure 3.6: The LTE Downlink Signal, showing one full 10ms frame (left), and the central 1.4 MHz of the same frame (right)

The LTE downlink and uplink signals are made up of "physical channels" and "physical signals". These physical channels and signals are multiplexed together in time and frequency, in a non-overlapping manner and mapped onto the time-frequency frame lattice. The mapping of physical channels within the frame is defined in the broadcast messages sent by each base station. This method of information mapping allows a jammer to selectively jam information contained in specific resource elements (REs) or interfere with specific physical downlink channels or signals. Figures 3.6 and 3.7 show the mapping of downlink and uplink LTE signals respectively, when using frequency division duplex mode. Each color represents a different physical control channel or signal, whereas the white spaces represent data.

LTE-A networks are an evolution of LTE. They use the same resource structure as LTE (shown in Figures 3.6 and 3.7) and add additional signalling and resources to support carrier aggregation, coordinated multi-point (CoMP) transmission and reception, and other LTE-A features that are beyond the scope of this dissertation.

3.5.3 Vulnerability of Physical Channels and Signals

The following subsections investigate the various LTE physical channels and signals as well as discuss potential threats that could cause communications denial. All threats analyzed in the remainder of this chapter are fundamental to the protocol and thus apply to LTE and LTE-A networks. Table 3.1 highlights the parameters associated with each physical channel and signal and will be referenced throughout the remainder of this article.

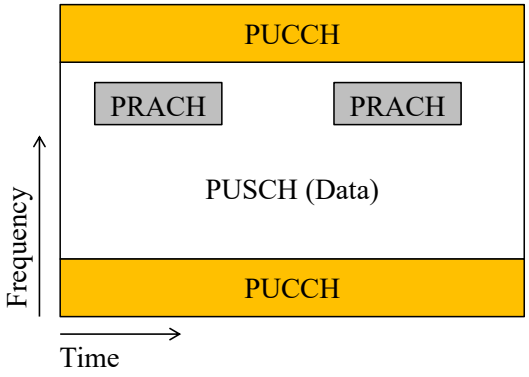

Figure 3.7: The LTE Uplink Signal

Synchronization Signals

The Primary Synchronization Signal (PSS) is a downlink synchronization signal, received by the UE in order to find and begin connecting to cells (macro-cell base stations typically have three cells, also known as sectors, each). By detecting the PSS, the UE determines the cell's physical layer identity and acquires time and frequency synchronization. The Secondary Synchronization Signal (SSS) provides the UE with the physical cell identity group. The physical cell identity group together with the physical layer identity provides the full Physical Cell Identity (PCI). Through the SSS, the UE also learns about the Cyclic Prefix (CP) type and the duplexing mode used by the cell.

Jamming the PSS or SSS requires fairly high power, because they are designed to be detectable at a low signal-to-noise ratio (SNR), so that the UE can also detect neighboring cells. It has been shown that a more effective method for attacking the PSS is to use RF spoofing, to prevent the UE from detecting the real PSS of a given cell [11]. RF spoofing refers to transmitting a fake signal meant to masquerade as an actual signal [68]. PSS spoofing essentially means that the attacker transmits a fake PSS, asynchronous to the LTE frame (i.e., not overlapping in time with the real PSS) and at higher power.

In order to understand the effect of PSS spoofing, we point out that the 3GPP LTE specification states that "the UE needs to search only for the strongest cell" at any given frequency [69]. The LTE specifications do not specify the behavior of the UE when it detects a valid PSS with no associated SSS. Hence, this will be implementation-specific. However, if the PSS and SSS are both spoofed, the 3GPP specification for the Radio Resource Control (RRC) layer [70] states that if the UE is in the idle mode and does not receive the Master Information Block (MIB) message after receiving the PSS and SSS, the UE shall treat this cell as "barred" and is allowed to select the second strongest cell within the same frequency. Since the 3GPP specifications do not allow the UE to select the second strongest cell in most cases, as mentioned before, many UE baseband chips may overlook the importance

of choosing the second strongest cell in this particular case for the sake of simplifying the interface between the PHY and the RRC layer.

Downlink Reference Signal

An OFDM receiver needs to estimate the channel and perform equalization prior to decoding information. In OFDM systems, pilots or reference symbols are therefore transmitted on specific subcarriers in parallel with the data. These reference symbols are generated at the PHY-layer and collectively called the Cell-Specific Reference Signal (CRS) in the LTE downlink (see Figure 3.6). The CRS occupies roughly 14% of the resource elements in a frame. The symbols are modulated with quadrature phase-shift keying (QPSK) and are generated from a length-31 Gold sequence, which is initialized with a value based on the cell ID. The cell ID also determines the location of the CRS in LTE resource grid.

It has been shown that jamming a subcarrier that carries pilots leads to a higher error rate than jamming one that contains only data [66, 71]. This is because the adjacent subcarriers are also affected, due to the nature of channel estimation. For a jammer to surgically transmit noise on top of LTE's CRS, it must detect the target eNodeB's PSS and SSS first to retrieve the cell ID. The jammer must also synchronize its transmissions with the target cell, using the PSS and SSS. However, it does not need to be perfectly synchronized, due to LTE's long symbol duration of 66.7 microseconds. Even if the difference in the signal path lengths to the UE were 5 miles, the propagation delay difference would only be 27 microseconds, which could easily be compensated for, if needed, by the jammer transmitting a fraction of a symbol longer each time. This applies to all synchronous jamming attacks discussed in this article.

Asynchronous multi-tone jamming of CRS is also possible for a jammer with a low-complexity transceiver, where there is no need for synchronization of the jammer with the eNodeB. This strategy involves transmitting noise on all CRS subcarriers (one third of all subcarriers) at a 100% duty-cycle. However, this would come at the cost of about seven times more power than the synchronous case and would lead to a threat that is only slightly more effective than jamming the entire downlink frame.

Downlink Broadcast Channel

After synchronizing with the cell and with the help of the CRS, the UE receives more information about the cell by decoding the Master Information Block (MIB), which is transmitted on the Physical Broadcast Channel (PBCH). The MIB contains information essential for initial access to a cell. It consists of 14 bits that contain the downlink system bandwidth, information allowing frame synchronization, and other control information [72]. It is mapped to the center 72 subcarriers, on the first 1 millisecond sub-frame of every frame. The PBCH is transmitted using QPSK, and uses a 16-bit cyclic redundancy check (CRC) as a form of error detection. Against a 10 MHz signal, PBCH jamming only requires jamming

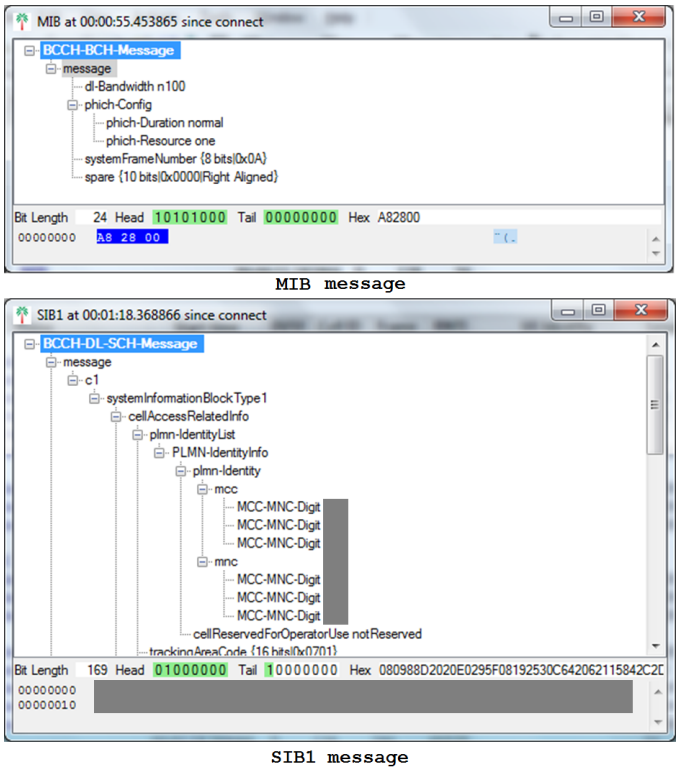

Figure 3.8: Real MIB and SIB1 Messages Captured from a Production Network

about 10% of the downlink subcarriers with a 3% duty-cycle, making it a very efficient synchronous jamming attack.

While jamming the PBCH is of concern, simply sniffing it may give the adversary information useful to more efficient attacks. Information carried over the PBCH allows the UE to determine the location of the System Information Block (SIB) messages, which are carried over the Physical Downlink Control Channel. These messages indicate the complete configuration of the cell and other critical information of the mobile network, including the eNodeB's idle timer, the configuration of the Physical Random Access Channel (PRACH), and the configuration of the Paging Channel (PCH).

As illustrated in Figure 3.8, which was obtained with the Sanjole LTE sniffing tool and processed with the WaveJudge software [73], the entirety of the information broadcasted by all eNodeBs in the MIB and SIB messages is sent in the clear. This allows an adversary to sniff this traffic and extract all details about cell and network configurations. For example, sniffing the SIB1 message allows identifying the mobile operator running the eNodeB. In the case of a public safety LTE deployment, a passive sniffer could identify the specific cells that are deployed for critical communications and distinguish them from mobile operator eNodeBs.

Having complete knowledge of the MIB and SIB messages could also be leveraged by an

attacker to determine the location of the PRACH in order to efficiently jam it, as discussed in Section 3.5.4. Other types of higher-layer network attacks [74] are enabled as well. Identifying the RRC idle timer value [70], an attacker could maximize the number of state transitions in a control plane signaling overload threat.

Downlink Control Channels

The Physical Control Format Indicator Channel (PCFICH) is used to send the UE information regarding where the Physical Downlink Control Channel (PDCCH) is located in the time-frequency lattice. Without successful decoding of this information, the UE will not be able to decode the PDCCH. The PDCCH contains information about the UE uplink and downlink resource allocation, which is vital for receiving LTE service. Although it is possible to jam the PDCCH directly, we will first discuss jamming the PCFICH.

The PCFICH appears only in the first OFDM symbol in each subframe and occupies a total of 16 REs. In other words, it is an extremely sparse channel, making it vulnerable to efficient jamming. Jamming the PCFICH consists of transmitting on top of the 16 REs that carry the PCFICH. The resource elements used for the PCFICH are shown in blue in Figure 3.6. The resource mapping is not static, but, rather, determined by the eNodeB's PCI [75], which the jammer can acquire through the PSS and SSS. This also limits a PCFICH jamming attack to a single cell, although multiple attacks could be launched by a single jammer.

Jamming the PDCCH also requires synchronization with the cell, but it is much less sparse than the PCFICH, making it a less effective jamming attack. In addition, since the PDCCH size varies between one and three OFDM symbols, the jammer needs to decode the PCFICH first in order to launch an effective attack with the least amount of power.

Hybrid-ARQ Indicator Channel

Acknowledgments (ACKs/NACKs) for uplink packets are sent on the downlink channel called the Physical Hybrid-ARQ Indicator Channel (PHICH). The PHICH uses binary phase-shift keying (BPSK) with repetition-3 coding [75]. This physical channel is fairly sparse, and thus PHICH Jamming is a threat worth considering.

3.5.4 Downlink and Uplink User Data

The Physical Downlink Shared Channel (PDSCH) and Physical Uplink Shared Channel (PUSCH) are used to transmit user data from the eNodeB to the UE and vice versa. While surgically jamming these channels is possible, the adversary might as well jam the entire signal, as these signals correspond to a majority of the resource grid. Thus, PDSCH and PUSCH jamming are two of the least important threats to consider.

However, it is possible to jam a specific user's uplink transmissions. Doing so would require extensive decoding of control information and knowledge of the user's temporary mobile identity number. This makes it an extremely complex attack that might be considered a combination of jamming and cyber-attack. We, therefore, do not include it in the vulnerability assessment.

Uplink Control Channel

The UE uses the Physical Uplink Control Channel (PUCCH) to send a variety of uplink control information (UCI) to the eNodeB, including scheduling requests, Hybrid Automatic Repeat Request (HARQ) acknowledgments, and channel quality indicators. The UCI is mapped to the resource blocks on the edges of the system bandwidth, as shown in Figure 3.7. This allows PUCCH jamming to be possible when the jammer only knows the LTE system bandwidth and center frequency. For an uplink bandwidth of 10 MHz, roughly 16 resource blocks (or 192 subcarriers) are allocated to the PUCCH [75]. Therefore, PUCCH jamming requires jamming about 25% - 30% of the uplink system bandwidth. The PUCCH is modulated with a combination of BPSK and QPSK, and uses $1/3$ rate convolutional coding. Because of its low complexity, PUCCH Jamming is an important threat to consider. Also, note that as a jamming threat against an uplink channel, its impact is on the entire cell as opposed to locally around the jammer.

Random Access Channel

After the initial cell search, the UE initiates the random access procedure with the objective to establish a RRC connection with the network. By transmitting the random access preamble on the PRACH, a UE lets the eNodeB know its presence and that it wants to connect to the cell. The specific location of the PRACH is conveyed to the UE in the SIB2 message, which is carried over the PDCCH. Therefore, to effectively jam the PRACH, the jammer must decode the SIB2 message fields. It is important to note that a successful jamming attack against the PRACH will prevent new UE from accessing a base station, but will not cause immediate DoS for active UE. However, any active UE transitioning between idle and connected RRC states will be blocked, resulting in all devices within a cell being blocked within a rather short period of time.

3.5.5 Vulnerability Assessment

We have discussed several jamming and spoofing attacks against LTE. This section compares these attacks in terms of complexity and effectiveness to quantify the vulnerability of LTE and determine its weakest links. First, we need to introduce two different ways of measuring the received jammer-to-signal ratio (J/S), that is, the ratio of the received jamming signal

Table 3.1: Physical Channel and Signal Modulation Scheme, Coding Type and Rate, Sparsity, Synchronization Requirement, and minimum J/S to cause DoS

Channel/Signal	Modulation	Coding	Coding Rate	% of REs	Sync?	J/S_{CH}	J/S_F
PDSCH	{4,16,64}-QAM	Turbo	Adaptive	$\approx$ 100%	No	0 dB	0 dB
PBCH	QPSK	Convol.	1/48	0.3%	Yes	0 dB	-25 dB
PCFICH	QPSK	Block	1/16	0.2%	Yes	0 dB	-27 dB
PDCCH	QPSK	Convol.	1/3	7%	Yes	-5 dB	-16.5 dB
PHICH	BPSK	Repetition	1/3	1.5%	Yes	3 dB	-15 dB
PUSCH	{4,16,64}-QAM	Turbo	Adaptive	$\approx$ 75%	No	0 dB	-1 dB
PUCCH	BPSK, QPSK	Convol.	1/3	$\approx$ 25%	No	-5 dB	-11 dB
PRACH	Zadoff-Chu Seq.	N/A	N/A	$\approx$ 2%	Yes	10 dB	-7 dB
PSS (Spoofing)	Zadoff-Chu Seq.	N/A	N/A	0.45%	No	3 dB	-20.5 dB
SSS	M-sequences	N/A	N/A	0.2%	Yes	15 dB	-12 dB
CRS	QPSK	N/A	N/A	5%	Yes	5 dB	-8 dB

power to the received LTE signal power. Two different J/S metrics are required because there are two different ways to observe J/S.

We will define J/S_{CH} as corresponding to a J/S that only takes into account the specific subcarriers and OFDM symbols (a.k.a. REs) of the channel or signal being jammed. For example, when jamming the broadcast channel (the light blue region in Figure 3.6), it is assumed the jammer will place its energy on top of the broadcast channel in time and frequency, and not transmit on any other REs. Thus, J/S_{CH} corresponds to the received power from the jammer divided by the received power of only the broadcast channel, not the entire downlink signal.

J/S averaged over an entire frame will be referred to as J/S_F. Using the previous example of jamming the broadcast channel, J/S_F corresponds to the received power from the jammer divided by the received power of the entire downlink signal. The J/S_F metric provides a convenient way to compare each jamming attack against the baseline attack, which is jamming the entire downlink or uplink signal.

Note that J/S alone does not given enough information to determine how large an area around the jammer is jammed (i.e., the radius of effect). Link budgets, which take into account factors like the jammer's transmit power and channel attenuation, must be created to determine such information.

The vulnerability of each channel or signal is based primarily on three factors:

1. The sparsity of the channel with respect to the entire downlink or uplink frame, i.e. the percent of REs used for the channel.
2. The jamming power needed to significantly corrupt the channel or signal, which we measure using the metric J/S_{CH}.
3. The complexity of the jammer required to perform such an attack, mostly based on whether synchronization to the cell is needed or not.

This information for each channel and signal is summarized in Table 3.1. The sparsity can

be combined with the minimum J/S_{CH} needed to cause immediate denial of the channel or signal to find an approximation for the corresponding J/S_F. This is an approximation because it assumes a uniform power spectral density across the LTE downlink or uplink signal, which is not the case in real world deployments. From the perspective of a jammer trying to minimize its power consumption and be more difficult to detect, a lower J/S_F is better.

The jamming portion of the vulnerability assessment involved a series of experiments using both simulation and tests with commercial LTE equipment. These experiments were meant to determine the approximate J/S_{CH} needed for each attack to cause DoS. First, we simulated each of the downlink jamming attacks using a system bandwidth of 10 MHz and one UE. We used the open-source, 3GPP compliant LTE emulation library known as srsLTE, a library that provides a full physical layer software radio implementation for both the LTE downlink and uplink. It allows full operation of a software-radio-based eNodeB, with ability to transmit and receive on all physical channels. We define the minimum J/S_{CH} needed for a successful attack as causing either an error rate of 10%, or a failed detection rate of 90%. At these failure rates, DoS is almost surely caused, making them fairly conservative figures. The jamming attacks not requiring synchronization to the cell were performed against a single commercial UE, connected to a commercial eNodeB running LTE Release 8. The specific eNodeB will not be disclosed due to the sensitive nature of jamming. Throughput was measured for each experiment, and the minimum J/S_{CH} was measured when throughput reached 10% relative to the baseline (no jammer) scenario. Results of these experiments are summarized in the J/S_{CH} column of Table 3.1.

To analyze the effect of RF spoofing, we built a testbed using the Rohde & Schwarz's CMW-500 as the legitimate eNodeB. To emulate PSS and SSS spoofing we used srsLTE as well, which ran on a laptop connected to a software-defined radio. A commercial LTE dongle was connected to a 2nd laptop, which monitored the UE state. For both cases of spoofing, either through PSS or through PSS and SSS, we observed that the UE was not able to camp (i.e., maintain a connection) on the legitimate eNodeB while the spoofing attack occurred, as long as the fake signal was received at a higher power level. This resulted in the UE being denied cellular service. Even though this corresponds to a J/S_{CH} of 0 dB, a 3 dB "safety-margin" from the perspective of the jammer was added, as seen in Table 3.1. Note that in the comparison we only include PSS Spoofing. Performing PSS and SSS spoofing combined requires 3 dB more power in terms of J/S_F. Likewise, PSS and SSS jamming are not included in the comparison because they require considerably more power and are not efficient attacks.

Based on the information gathered in Table 3.1, we can form an initial threat assessment of the vulnerability of LTE to jamming and RF spoofing. We compare the attacks against a baseline attack, which we define as barrage jamming over the entire LTE system bandwidth on either the downlink or uplink frame. Barrage jamming simply involves transmitting noise (typically Gaussian) over the entire LTE frame. Because there is an efficiency and complexity aspect to each attack, instead of simply ranking them, we have assembled the attacks into

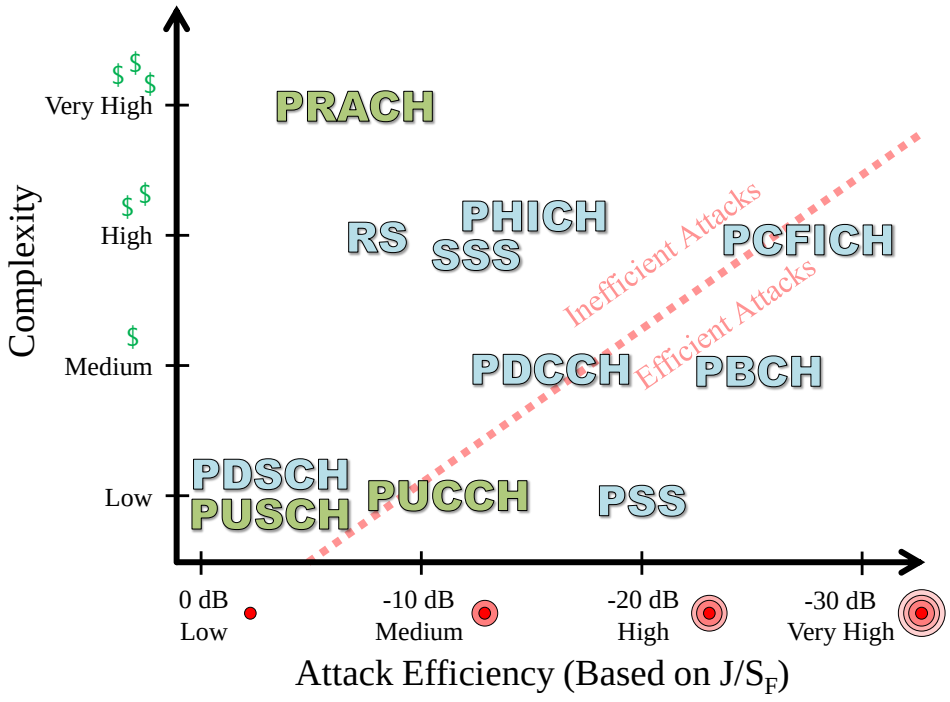

Figure 3.9: Ranking of Attacks Based on Jamming Efficiency and Complexity

a two-dimensional map, shown in Figure 3.9. From the perspective of a jammer, the most desirable attacks are towards the bottom-right. Specifically, we believe that efforts toward hardening LTE for critical communications should focus on mitigation of PSS spoofing, PUCCH jamming, PCFICH jamming, and PBCH jamming.

It is also important to note that, even the most complex attacks can be easily implemented with widely available open-source LTE libraries, low-cost software radio equipment with a budget under $1500 and basic Linux programming skills.

3.5.6 Survey of Mitigation Techniques

Before we discuss methods for mitigating jamming and RF spoofing attacks to LTE, it is important to understand the implications of the changes needed to harden LTE. The cellular technology inside of modern cellphones and other UE resides in an application-specific integrated circuit (ASIC), sometimes referred to as a system-on-chip or the chipset. On the other side of the link, the eNodeBs typically use a baseband unit that does most of the processing in software, while an RF module handles the RF chain. Thus, changes to the behavior in the eNodeB likely only require a firmware update, while changes to the UE require a new chipset to be designed and manufactured.

There is little openly available literature related to LTE jamming attacks, and even less on

mitigation of attacks. The authors of [76] propose various methods for enhancing the security of LTE networks against jamming attacks. This includes spread-spectrum modulation of the downlink broadcast channels. This strategy is meant to mitigate a jammer that targets the center 1 MHz of the downlink signal, where many important signals and channels are located. By using direct-sequence spread spectrum (DSSS), the important signals and channels can be spread across the entire available downlink bandwidth, which in most cases is 10 MHz. The authors also propose scrambling the radio resource allocation for the PUCCH with an encrypted sequence, whereby the allocation of the PUCCH is no longer on the band edges of the uplink band, but instead can appear anywhere in the uplink frame. Only legitimate users connected to the cell would know how to decrypt the scrambled sequence. Lastly, the authors of [76] propose a system in which the MIB and SIBs would be encrypted so that essential network configuration parameters are not transmitted in the clear. All three of these anti-jam strategies require changes to the UE chipset as well as the eNodeB because of the extensive modifications to the LTE protocol and signaling.

PSS spoofing can be mitigated by creating a timer for receiving the SSS [68]. If this timer expires, the UE should blacklist the PSS and choose the second strongest cell within the same frequency. PSS and SSS spoofing attacks can be mitigated by having the UE create a list of all available cells (PCIs) in the given frequency channel, along with their received power levels. The UE then could then search for the PBCH of the strongest cell, and have another timer for decoding the MIB. If this timer expires, the UE would look for the PBCH of the next strongest cell, and so forth.

A simple way of mitigating PUCCH jamming would be to provide periodic PUSCH resources to UE, even if not requested [6]. This way, the UE will send its uplink control information on the PUSCH instead of the PUCCH. Since the downlink resources are typically the bottleneck, the overhead associated with such periodic assignment of PUSCH resources might not be as critical.

The authors of [7] investigate the PCFICH jamming attack and propose a mitigation strategy called “extra-blind PDCCH decoding”. This strategy suggests that the UE decodes each PDCCH block with all three possible CFI values, instead of extracting it directly from the PCFICH. Another option is using a fixed CFI for mission-critical LTE networks or operational modes. Unfortunately both of these strategies require modifications to the UE chipset, making them unlikely to be implemented unless they are added to the 3GPP specifications.

These mitigation strategies only address a few of the attacks that we discussed in this article. Further research on mitigation techniques that require minimal changes to the UE, and the LTE standard itself, is needed.

3.6 Conclusion

As the sophistication of communications systems increases, sophisticated jamming will likely become a bigger threat in public safety, military, and other mission-critical domains. The jammer taxonomy introduced here frames the organization of jammer classes by what information they possess and their capacity to act on that information. This new view of jammers emerges naturally from the way present day wireless technology relies so extensively on software-driven behavior. In addition, understanding the key capabilities that distinguish major classes of jamming, as well as the multidimensional parameter space, can aid in the correct application of anti-jam and detection strategies.

Further research includes the design of a radar jamming taxonomy and radio navigation jamming taxonomy. It may be possible to formulate a taxonomy that applies to all forms of jamming.

Chapter 4

Throughput-Based Antifragile Gain

A portion of the material in this chapter has been previously published in [1], and is being reproduced in this dissertation with the consent of all co-authors, as well as the original publisher if required.

4.1 Introduction

Jamming is an ongoing threat that plagues wireless communications in contested areas. Unfortunately, jamming complexity and sophistication will continue to increase over time. This is in part due to the availability of powerful and low-cost software-defined radios. The current approach to countering such a threat revolves around *jammer detection* and *jammer mitigation*. As such, an increase in jammer complexity requires an increase in countermeasure complexity. This leads to extremely hardened radios that sacrifice communications performance for more advanced jamming protection. To provide an escape from this trend, we investigate the previously unexplored area of *jammer exploitation*. Unlike mitigation (i.e., anti-jamming), the more complex an enemy jammer, the more potential there is for exploitation. It is for this reason that the antifragile paradigm should be applied to wireless communications. An example of jammer exploitation includes manipulating a jammer into jamming a particular sequence of channels, where data is conveyed in the sequence of channels selected, similar to frequency shift keying. A strategy that exploits a jamming attack to provide a communications gain, such as reducing the bit error rate or increasing the rate of communication, can be labeled as "antifragile communications". This is because a gain is achieved by harnessing the presence of a stressor, which in this context is the jammer.

Antifragility is a concept popularized by Nassim Nicholas Taleb and is a term he coined in his 2012 book *Antifragile* [17]. Antifragility refers to systems that increase in capability, resilience, or robustness as a result of mistakes, faults, attacks, or failures. As Taleb explains in his book, antifragility is fundamentally different from the concepts of resiliency (the ability to recover from failure) and robustness (ability to resist failure). "The resilient resists shocks and stays the same; the antifragile gets better" [17].

As the primary focus of this dissertation, we apply the concept of antifragility to wireless communication systems. Specifically, we seek to exploit the presence of a jammer to achieve a communications gain relative to a jammer-free case. This should not be confused with anti-jamming, which seeks to mitigate jamming and perform at (or near) interference-free capability through the duration of an attack. Antifragile communications take mitigation one level further, by providing a boost during the attack. Additionally, our strategy should not be confused with self-jamming, i.e., friendly jamming, in which a secrecy channel is formed by *intentionally* transmitting noise along with the communication signal in order to prevent the eavesdropper from inferring any information.

The contributions of this chapter are summarized as follows:

1. Development of a novel *antifragile waveform*, used for manipulating a reactive jammer into relaying information, whereby a gain (relative to a non-jammed case) in throughput, connectivity, or covertness can be achieved.
2. Introduction of a generalized model for reactive jamming, applicable to both repeater-based and sensing-based jamming behaviors.

The remainder of this chapter is organized as follows. Section 4.2 discusses background information. Section 4.3 introduces the antifragile scenarios that define the scope of this chapter, while Section 4.4 defines the jammer models under consideration. Section 4.5 introduces the components of an antifragile waveform, with emphasis on the jammer piggybacking strategy. Section 4.6 develops theoretical channel capacities under each jammer model when using the antifragile waveform. Through numerical evaluation, Section 4.7 provides feasibility regions for the proposed techniques (i.e., regions in which an antifragile gain occurs for the given scenario). Section 4.8 concludes.

4.2 Background

4.2.1 Related Work

Refer to Section 2.2 for a literature review of works involving the concept of antifragility across science and engineering.

While not meant to be antifragile, research involving the establishment of a timing channel to counter a reactive jamming attack, including [77] and [78], has similarities with the approach described in this chapter. The strategy in [77] is similar to our *replace with noise* jammer exploitation approach, in terms of using the presence or absence of a signal to carry information. Such a strategy does not attempt to evade the jammer, but rather, function *in spite* of jamming, although the authors do not go as far as to *exploit* the jammer. However, the work in [78] could be considered as jammer exploitation. In an approach similar to jammer piggybacking, the authors investigate the use of a timing channel to counter reactive jamming, where information is encoded in the interval between the instant when the jammer terminates its jamming signal and the beginning of the transmission of the next packet. This strategy is depicted in Figure 4.1. The end of the jamming signal is used as a reference because the end of the previous packet is covered by the jammer's interference. While the approach in [78] involves the use of the jamming waveform to convey information, it is done using a different process and for a different reason than the approach described in this chapter. In addition, the strategy in [78] involves the source node and destination node independently estimating the precise time that the jamming signal ends, which could be very difficult in lower jammer-to-noise power level conditions, or when there is a rapidly changing channel. The precision of this timing estimation also determines how many bits per symbol can be transmitted (although the authors claim it is only based on the precision of the radio's onboard clocks). Any propagation delay also has to be taken into account, although this could be addressed through pilot symbols. Overall, this strategy lends itself to a lot of great analysis, but has many practicality issues that are not yet addressed in literature.

To the best of our knowledge, this is the first literature that specifically applies the concept of antifragility to wireless communications. In addition, this is the first literature that describes

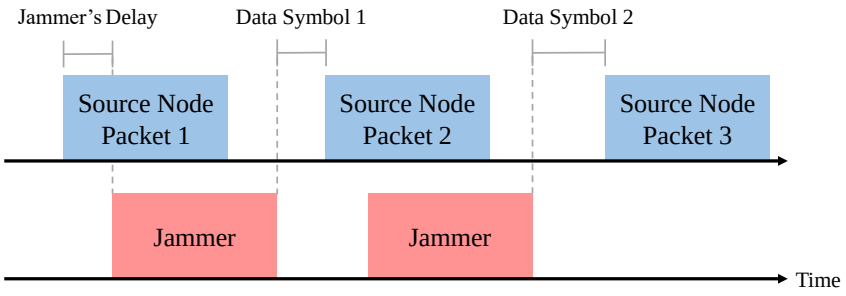

Figure 4.1: Jammer exploitation strategy described in [78], showing how information is conveyed through the time between the end of a jamming pulse and the beginning of the next packet.

a method for piggybacking off a reactive jammer, to achieve a communications gain relative to a jammer-free case. To broaden the application of the techniques introduced in this chapter, we have explored them under a wide variety of jamming behaviors. However, we make no claims as to the existence of these jammers in modern warfare.

4.2.2 Antifragility Compared to Similar Terms

To make the concept of antifragility more clear, we compare it with similar concepts.

Antifragile vs. Robust/Resilient: The concept of antifragile is most often confused with robustness or resilience. While an object or system can certainly be robust against stressors or harsh conditions, it is only antifragile if it benefits from them. In other words, compared to an established baseline performance of a given metric, a resilient system will never increase over the baseline due to harsh conditions, while an antifragile system has the ability to.

Antifragile vs. Adaptive: An adaptive system is one that changes its behavior based on information available at time of utilization (as oppose to having the behavior defined during system design). While adaptive systems allow for robustness under a variety of scenarios, they are not necessarily antifragile. In fact, it is often difficult to determine if an adaptive system is also antifragile, because baseline performance in an adaptive system is usually not well-defined, and typically based on the current conditions which may be changing.

Antifragile vs. Cognitive: We usually think of cognitive systems as incorporating artificial computational processes that act like a person. A cognitive entity is one that is capable of decision making, carrying out actions depending on its own goals and its perception of the world, and learning from experience. In the wireless domain, cognitive is most often discussed in the context of cognitive radios, which are able to perform well across a wide variety of harsh conditions. While *cognitive* characterizes how a system works, *antifragile* is more concerned with the output or performance of the system. An antifragile system can

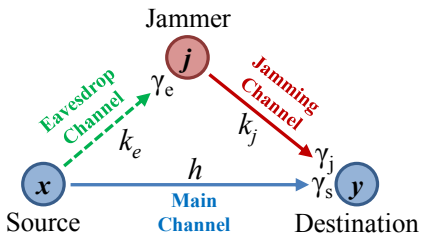

Figure 4.2: The geometrical configuration of a reactive jamming scenario illustrates the three channels involved. At each point a signal is received, a SNR, denoted with γ, is shown.

certainly contain cognitive components to it, but a cognitive system is only antifragile if it is able to satisfy the criterion of increasing in capability in some way *as a result of* a fault, attack, failure, or any negative condition. As with *adaptive*, evaluating the performance of cognitive radios is challenging, as the overall behavior (not just parameter space) of cognitive radios is not predefined and likely changes with the operating environment [79].

4.3 Antifragile Strategies

4.3.1 Motivating Scenario: Jammer Piggybacking

One approach to antifragile wireless communications is to manipulate a hostile jammer into unknowingly helping with the transmission of data. We call this strategy "jammer piggybacking". With jammer piggybacking, we seek to exploit a jammer that is correlated with the signal it is targeting, i.e., a reactive jammer. In Section 4.5 we show that information transmitted between two nodes can piggyback off the jammer, which effectively makes the jammer a relay node.

Figure 4.2 shows the geometrical configuration of a reactive jamming scenario. From an information theoretic point of view, if the transmitted signal x is jammed by signal j, an antifragile gain is only possible if the mutual information between the two is greater than zero, denoted as $I(x;j) > 0$, where $I(a;b)$ indicates the mutual information between two random variables a and b. Mutual information helps to measure how much knowing one variable reduces uncertainty about the other. In other words, if the inequality is greater than zero, then the jamming signal contains information that the source node is sending to the destination node. This information could even be manifested in the presence or absence of the jamming signal, similar to how in on-off keying (OOK) information is conveyed by the presence or absence of a carrier. To perform jammer piggybacking, the destination node must be able to separate the two signals, demodulate them, and then combine them (or simply use the one that has the most integrity).

While under some scenarios it is possible to achieve a throughput gain from jammer pig-

gybacking, with respect to a non-jammed case, an antifragile gain could also stem from other advantages. For example, if a low bit-rate signal known by the destination node is successfully relayed through the jammer, the destination node can use this signal to aid in null-steering (assuming it has that capability). By having a signal to correlate to, as opposed to just noise, the destination node can more quickly and accurately point a null towards the jammer. Another example involves solving the hidden node problem by manipulating a high-power jammer into providing a control channel (or rather a relay for a control channel) among a network that may include hidden nodes. A directional antenna should be used by the manipulating radio, so that the other radios only receive the jammer's version of the signal, and the original signal does not interfere with other transmissions. Both of these examples show how jammer piggybacking can still provide an antifragile advantage, even without achieving a higher-than-baseline bit-rate.

This concept of reactive jammer piggybacking is further developed in this chapter, as well as Chapter 5.

4.3.2 Other Example Antifragile Strategies

In addition to jammer piggybacking, we propose the following methods of antifragile communications:

Achieving Coarse Time Synchronization

A high power pulsed enemy jammer (or even enemy radar) could act as a way for multiple radios to achieve coarse time synchronization. For example, if a group of radios in the same general area use time division multiple access (TDMA), slot transitions could be governed by the pulses of the enemy transmitter. Therefore, this strategy allows a wireless network to increase in synchronization capability as a result of an attack. This strategy is further explored in Section 6.2.

Inducing Jammer Friendly Fire or Jammer Herding

Causing an enemy jammer to jam other systems friendly to it could have a clear antifragile advantage, since the total electronic attack performance increases as a result of the attack. Realizing this type of strategy would be based heavily on the specific communications protocol and jammer behaviors, and thus is beyond the scope of this work and is left for future investigation.

Hiding in an Enemy Jammer's Signal

Low-probability-of-intercept/low-probability-of-detection (LPI/LPD) is an important aspect of wireless communication waveforms meant for mission-critical use. This antifragile strategy involves increasing LPI/LPD by hiding the desired signal within a jamming signal. Clearly this is only possible if there is a way to cancel out the jamming signal at the destination node, leaving only the desired signal to be demodulated. In order to be truly hidden under a signal, the two signals must (mostly) overlap in time and frequency, and the *cover signal* must be, in general, higher power. One possible method of removing the jamming signal is to exploit cyclostationarity in the jamming signal, which could be a result of the jammer using a specific modulation scheme or using a repetitive pattern (such as a chirp signal). By hiding under the enemy jammer's signal while maintaining communications, LPI/LPD is increased as a result of an attack.

Communicating Simultaneously with Legacy Friendly Jammers

It is possible that a party "jams themselves" due to electronic warfare (EW) platforms and communications radios existing in the same spectrum. This is because legacy EW equipment may not be designed to operate alongside communications. For example, a reactive jammer may be programmed to jam a wide band, but not have a mechanism to white-list friendly signals that did not even exist at the time the jammer was designed. While muting the jammer for short periods of time may be one solution, there are scenarios where the jammer needs to be active 100% of the time for safety. An alternative is to use the jammer piggybacking techniques developed in this chapter as well as the next chapter, where the communications is piggybacked through the friendly jammer, allowing it to remain active at all times. It is assumed that the legacy jammer's behavior cannot simply be changed. However, because it is a friendly jammer, we have full knowledge of how it works and are in a great position to tailor a piggybacking strategy to it.

4.3.3 Defining Three Classes of Antifragility

Because antifragility is measured by an increase in some sort of performance, we will define three classes of antifragility, as shown in Figure 4.3. The y-axis is left unlabeled, as there are several different ways a communication system can increase in performance or capability (e.g., throughput, spectral efficiency, node connectivity). Class I antifragile systems increase in performance during an attack, but the antifragile gain does not persist once the attack ends. Most of the discussion in this chapter (and Chapter 5) involves Class I strategies. Class II systems, on the other hand, increase after the attack (usually due to information obtained during the attack). Note that in Figure 4.3, the Class II example shows what is essentially perfect resiliency during the attack, which may not always be the case (throughput may even drop to zero during the attack). Class III systems represent a combination of both Class I

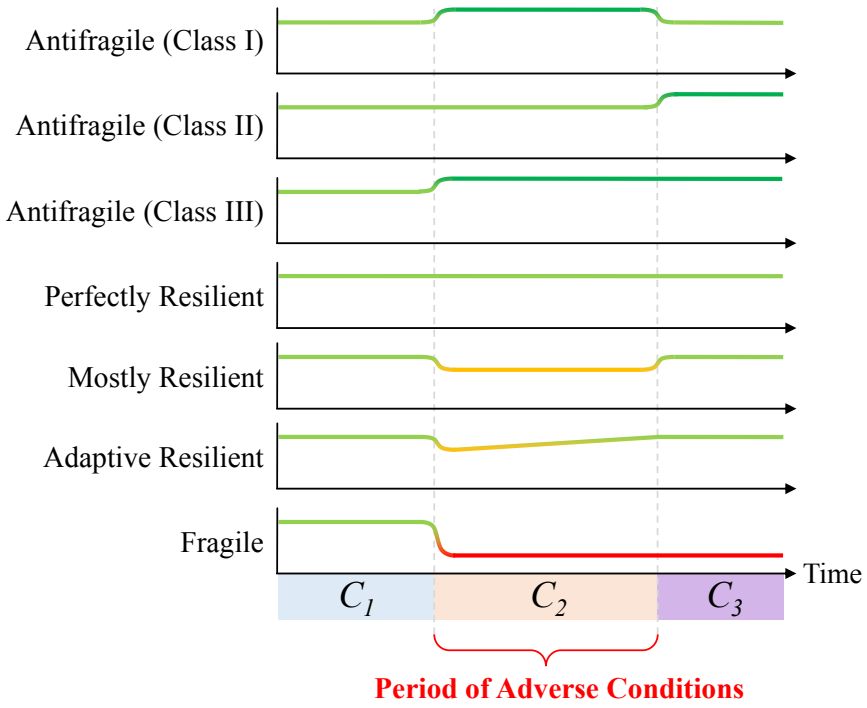

Figure 4.3: Diagram of antifragile, resilient, and fragile systems. Antifragile systems are able to maintain objective performance through a period of adverse conditions and achieve super-objective performance during or after the period of adverse conditions.

and Class II, in which the antifragile gain occurs at the beginning of the attack and persists after the attack ends.

4.4 System Model

In this section we describe the channel model and jammer models under consideration. As mentioned before, we are most interested in jamming of a reactive nature, with the goal of making the jammer act as an unwitting relay. Throughout this dissertation we refer to the transmitter as the source node, and the receiver as the destination node.

4.4.1 Channel Model

We consider a transmitted signal x that goes through a channel h with propagation delay τ_h. This signal x can be eavesdropped on by the jammer via a channel k_e with delay τ_e, and then jammed by signal j that goes through a channel k_j with delay τ_j. We assume all three channels are noisy memoryless channels, thus h, k_e, and k_j correspond to channel coefficients,

which may or may not be time-varying. In addition, the parameter τ_{jam} denotes the delay due to the jammer's RF chain and any sensing that may be performed. Noise n_d and n_j is the additive white Gaussian noise seen at the destination node's receiver and jammer's receiver respectively, with variance σ_d^2 and σ_j^2. These signals and channels are depicted in the generalized reactive jamming model, shown in Figure 4.4. Under this generalized model, $\beta(t)$ represents a time-varying transform applied to the received signal at the jammer, while $w(t)$ represents an internally generated jamming waveform.

When the jammer uses a behavior that only involves repeating the received signal with a transform $\beta(t)$ applied, $w(t) = 0$ and the received signal y is given by:

$$y(t) = \underbrace{hx(t-\tau_h)}_{\text{source signal}} + \underbrace{k_j\beta(t)\left[k_e x(t-\tau_e-\tau_{jam}-\tau_j)+n_j\right]}_{\text{jamming signal}} + \underbrace{n_d}_{\text{noise}} \tag{4.1}$$

Likewise, when the jammer uses a sensing-based behavior that involves transmitting a jamming waveform $w(t)$, $\beta(t) = 0$ and the received signal y becomes:

$$y(t) = \underbrace{hx(t-\tau_h)}_{\text{source signal}} + \underbrace{k_j w(t-\tau_e-\tau_{jam}-\tau_j)}_{\text{jamming signal}} + \underbrace{n_d}_{\text{noise}} \tag{4.2}$$

Note that while the sensing-based jammer does not retransmit x or n_j, the eavesdrop delay τ_e is still a factor because the jammer has to perform sensing.

The conditional probability of y given x and j, denoted as $p_{Y|X,J}(y|x,j)$ is assumed to be stationary and a function of the communications channels, which we will not impose a model for. The marginal distribution $p_X(x)$ is determined by the transmitter's physical layer parameters such as modulation and coding scheme.

Because the reactive jammer acts as a form of memory, we can apply Shannon's formula for capacity through a channel with memory, which is stated as [80]:

$$C = \lim_{n\to\infty} \sup_x \frac{1}{n} I(x^n; y^n) \tag{4.3}$$

Let C_1 and C_3 be the channel capacity before and after the jamming attack respectively, as labeled in Figure 4.3. In both of these cases $j(t) = 0$ and the traditional memoryless channel capacity applies. Let C_2 be the capacity during the jamming attack, when $j(t)$ takes on a certain behavior as discussed in the next subsection. Therefore, the criterion for Class I antifragility in the jammer piggybacking context is simply $C_2 > C_1$, indicating an increase in capacity as a result of jamming. Likewise, the criterion for Class II antifragility is $C_3 > C_1$. Class III antifragility requires that both inequalities be true, as depicted in Figure 4.3.

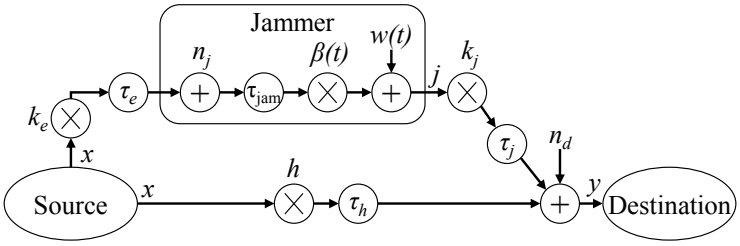

Figure 4.4: Generalized model for reactive jamming.

Table 4.1: Summary of model parameters.

Symbol	Description
h, τ_h	main channel coefficient and delay
k_e, τ_e	eavesdrop channel coefficient and delay
k_j, τ_j	jamming channel coefficient and delay
τ_{jam}	jammer's RF + sensing + intentional delay
n_j	channel noise at the jammer's receiver
n_d	channel noise at the destination node's receiver
$\beta(t)$	time-varying transform applied to signal
$w(t)$	internally generated jamming waveform
γ_s,γ_e,γ_j	SNR of main, eavesdrop, and jamming channel respectively

In Section 4.6 these capacities will be defined in terms of the SNRs of the three different channels involved in the analysis. For ease of analysis, we assume that all three SNRs are constant. The SNR of the source node signal as seen by the destination node is denoted as γ_s. The SNR of the source node signal as seen by the jammer when eavesdropping is denoted as γ_e. Lastly, the SNR of the jamming signal (not taking into account noise n_j) as seen by the destination is denoted as γ_j. These three SNRs are defined as follows:

$$\gamma_s = \frac{\mathrm{E}\left[|x|^2|h|^2\right]}{\sigma_d^2} \quad \gamma_e = \frac{\mathrm{E}\left[|x|^2|k_e|^2\right]}{\sigma_j^2} \quad \gamma_j = \frac{\mathrm{E}\left[|j|^2|k_j|^2\right]}{\sigma_d^2} \tag{4.4}$$

All parameters associated with the channel and jamming model are summarized in Table 4.1.

4.4.2 Reactive Jamming Models and Behaviors

We will model the transmitted signal $x(t)$ as a single-carrier signal with a certain carrier frequency f, phase ϕ, and amplitude A:

$$x(t) = Acos(2\pi ft + \phi) \tag{4.5}$$

Using this approach, these jammer models can apply to a communications link that uses frequency-hopping spread spectrum (FHSS), and also a traditional single-carrier based signal. We also assume the jammer does not know values of x a priori.

The basic behavior of a reactive/repeater jammer is that it receives a signal transmitted from a target transmitter and retransmits it with a possible transform applied, in a manner intended to jam the target receiver. By repeating the target signal, the jammer can follow the two radios as they hop around in frequency, countering the protection associated with FHSS (which could be over 20 dB). If the jammer has no knowledge of the frequency hopping pattern, then the best it can do is transmit across the entire hopping band. This leads to the FHSS link gaining an anti-jam advantage equal to the FHSS processing gain. However, a jammer can overcome this anti-jam advantage by following the signal in frequency, which requires receiving the signal and retransmitting it as fast as possible. In order for the jammer to increase its effectiveness, it tries to transmit the jamming signal with no significant frequency offset, in order to better align with the target signal in frequency. Outside of FHSS, reactive jamming can be used to surgically jam frames in a contention-based MAC protocol, and avoid jamming an empty channel. This lowers the power consumption of the jammer, as well as makes the jammer harder to detect, because it is effectively hidden under an actual communications signal. Literature based around this sense-and-transmit reactive jamming strategy includes [51, 53, 54].

Our first jammer model under consideration, the digital radio frequency memory (DRFM), is one that simply retransmits the target signal on a sample-by-sample basis [81, 82]:

DRFM: The jammer retransmits the received signal with a constant amplification gain β_A, such that $\beta(t) = \beta_A$:

$$j(t) = \beta_A \left[k_e x(t - \tau_e - \tau_{jam}) + n_j\right] \tag{4.6}$$

This jamming method does not require any form of frequency detector, but the jammer's bandwidth must be sufficiently high to cover the hopping range of the target link (if FHSS is in use). In the case that the noise power is not significantly higher than the repeated signal power, this jamming model can be thought of as an amplify-and-forward relay.

In some cases the jammer may want to transmit internally generated noise in place of the target signal. We generalize this jamming behavior with the following model:

Replace with noise: The jammer transmits internally generated random noise on the frequency that x was received on. When targeting a FHSS signal, this type of jammer is referred to as a follower jammer [82]. In terms of the generalized reactive jamming model, $w(t) = n_{jam}(t)$ when signal x is detected, and the jamming signal is given by:

$$j(t) = \begin{cases} n_{jam}(t - \tau_e - \tau_{jam}) & : \quad \text{x is detected} \\ 0 & : \quad \text{otherwise} \end{cases}$$

where $n_{jam}(t)$ can be modeled as a zero-mean band-limited Gaussian random process with variance σ^2_{jam}, spanning the bandwidth of signal x. This method requires some form of an energy detector to detect which frequency the target is transmitting on and if the signal is present. The delay associated with this detection process is included in τ_{jam}. Literature using this jammer model in the wireless communications domain includes [51, 53, 54, 82, 83].

While the *replace with noise* model seems like an effective method for jamming, in some scenarios the jammer may not be able to detect the frequency quick enough to successfully jam the enemy. Therefore, we will investigate some alternative behaviors that the jammer could use to prevent being an amplify-and-forward relay, without having to perform signal detection.

Phase flipping: The jammer randomly flips the phase of the received symbols:

$$j(t) = U(t) \left[k_e x(t - \tau_e - \tau_{jam}) + n_j \right] \tag{4.7}$$

where $U(t)$ is a pseudorandom sequence drawn from the set $U(t) \in \{1, -1\}$ with probability mass function given by $f_U(u) = 0.5 : U = 1, -1$. In terms of the generalized reactive jamming model, $\beta(t) = U(t)$. This could be realized by multiplying the received signal by a square wave that alternates between +1 and -1 randomly. This model is especially useful because practical modulation schemes, such as phase-shift keying (PSK), quadrature amplitude modulation (QAM), and amplitude-shift keying (ASK), can be corrupted through this approach.

Realistically, a jammer would have no reason to distinguish between symbols (symbol-level timing synchronization is not trivial), when it can simply change the value of $U(t)$ often enough to corrupt data carried in the phase. However, we use this model as a way to generalize the jamming tactic. The drawback to this jamming technique is that it spreads the transmitted signal in the frequency domain, which means there will be energy that does not overlap with the target signal. To reduce spreading, the jammer must increase the time between flipping the phase. A favorable switching period for the jammer would be one that is more frequent than the presence of equalization pilots, so that the phase shifts do not get equalized out at the destination node, but not so quick that the signal gets overly spread in frequency.

An alternative to the phase flipping approach would be to modulate the amplitude.

Modulate amplitude: Identical to the previous tactic, except the jammer randomly modulates the amplitude, and $\beta(t) = V(t)$:

$$j(t) = V(t) \left[k_e x(t - \tau_e - \tau_{jam}) + n_j \right] \tag{4.8}$$

Table 4.2: Jammer models in the context of the generalized reactive jamming model.

Jammer Model	$\beta(t)$	$w(t)$
DRFM	β_A	0
Replace with Noise	0	$\mathcal{N}(0, \sigma^2_{jam})$
Phase Flipping	$U(t) : u \in \{1, -1\}$	0
Modulate Amplitude	$V(t) : v \geq 0$	0
Replace with CW	0	$cos(2\pi ft)$

where $V(t)$ is a pseudorandom sequence of positive numbers. An example probability density function for V, in which the signal's average power is conserved, is $f_V(v) = 0.5$ where $0 \leq v \leq 2$.

If a jammer modulates the amplitude between a negative and positive value, we will consider it as also modulating the phase. In this *phase and amplitude modulating* case, the jamming signal likely resembles noise, and the *replace with noise* model is most appropriate.

Lastly, we will consider a model similar to *replace with noise*, except the jammer replaces the signal with a continuous wave (CW).

Replace with CW: The jammer transmits a CW on the frequency that the target signal was received on (i.e. $\beta(t) = 0$ and $w(t)$ is a sinusoid). Like the *replace with noise* model, this model requires some form of an energy detector. The phase and amplitude of the transmitted CW is independent of the target signal, leaving ϕ and A out of the equation for $j(t)$:

$$j(t) = \begin{cases} cos(2\pi f(t - \tau_e - \tau_{jam})) & : \quad \text{x is detected} \\ 0 & : \quad \text{otherwise} \end{cases}$$

Table 4.2 summarizes the discussed jamming models in the context of the generalized reactive jammer model shown in Figure 4.4.

Nulling-type attacks are not considered in this dissertation, because the accuracy of channel state information required by the jammer makes them infeasible [3]. Additionally, a time-delay type transformation is not included as a jammer model because it would be equivalent to the DRFM model with τ_{jam} intentionally increased. While the presented jammer models are not all-exhaustive, we feel they represent a significant portion of possible reactive jamming strategies.

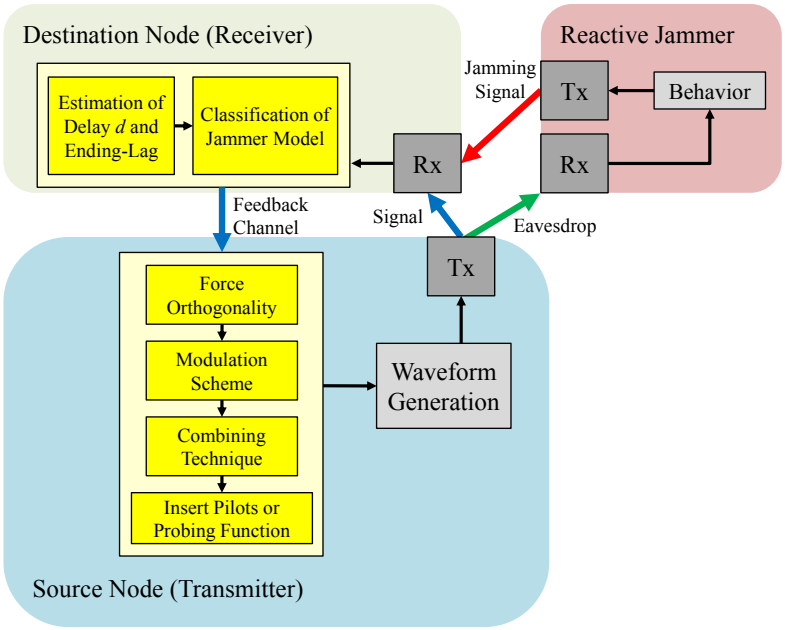

Figure 4.5: The components of the proposed antifragile system (highlighted in yellow) and how they fit into a communications system.

4.5 Components of an Antifragile Waveform

In this section we discuss components of an antifragile waveform specific to the reactive jammer piggybacking strategy, although many of the components discussed also apply to the other antifragile strategies in Section 4.3.2. The intent of an antifragile waveform is not to provide a radio with an antifragile gain at all times. Rather, an antifragile waveform grants the ability to achieve an antifragile gain when the situation allows for one. A functioning antifragile strategy is based on the scenario at hand, such as the type of communications link, the reactive jammer's characteristics, and the delays involved. An effective strategy involves implementation of a series of different schemes that exploit the jammer, and a classifier that can identify the scenario at hand and estimate parameters associated with it. Lastly, it must incorporate an engine that can assign the most effective antifragile (or mitigation) scheme to the given scenario. Figure 4.5 illustrates the components of the proposed antifragile system (highlighted in yellow) and how they fit into a communications system; each subsection in this section corresponds to one of the yellow boxes.

4.5.1 Delay Estimator

The most important question to answer is, "what is the delay associated with the reactive jammer?" In other words, the delay between when the actual signal is received and jamming

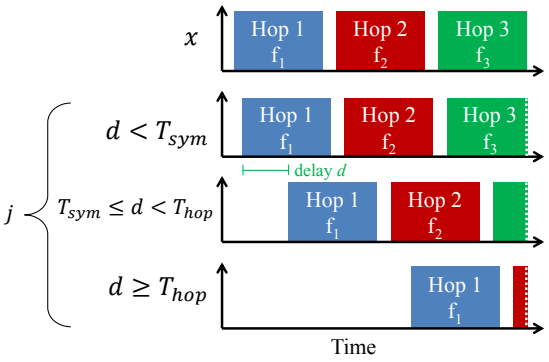

Figure 4.6: Examples of how the source and jamming signal are received at the destination node, for each of the three time delay cases.

signal is received at the destination node. We will denote this delay as d, and it is given by $d = \tau_e + \tau_{jam} + \tau_j - \tau_h$. We will split the possible delay scenarios into three cases, where T_{sym} is the period of one symbol, and T_{hop} is the period of one hop (or one packet/frame). These three cases are listed below and depicted in Figure 4.6.

1. $d < T_{sym}$
2. $T_{sym} \leq d < T_{hop}$
3. $d \geq T_{hop}$

The first case in which $d < T_{sym}$ represents a jammer that is close enough and has a low enough delay to overlap on a symbol-by-symbol basis. The third case in which $d \geq T_{hop}$ represents a slow repeater jammer that does not cause any link degradation, but can still be exploited.

To initially estimate delay d (before the jamming behavior has been classified), the destination node can simply observe the time lag between the hop/frame preamble, and the next burst of energy received. While it may take a few hops to get an accurate initial estimate, this process can be performed simultaneously with any other jammer detection on the platform. We will denote the current estimate of d as $\hat{d}$. Once the jammer is actively relaying information, updates to $\hat{d}$ can be made at the destination by observing the delay between two copies of any given preamble.

While the delay is an important measurement, it does not capture the time it takes the jammer to stop transmitting a jamming signal after the received signal at the jammer ends. We will denote this value as the *ending-lag*, and it is depicted in Figure 4.7. In cases such as an analog repeater jammer, this value of *ending-lag* would be near-zero because the jammer is simply retransmitting what it receives. However, a digital repeater jammer that uses sense-transmit cycles could have a short delay before it stops transmitting.

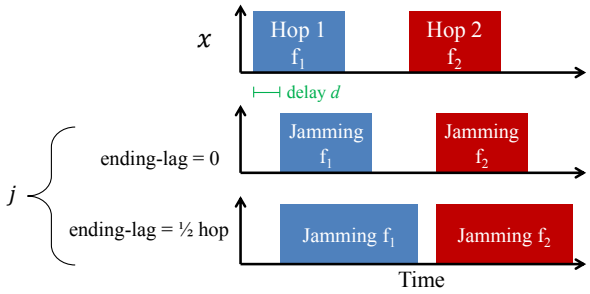

Figure 4.7: Depiction of *ending-lag*, which we define as the delay the jammer takes to stop transmitting after it stops receiving a signal.

4.5.2 Jammer Classification

Next, the radios must classify the reactive jammer's behavior. We assume that there has been time-frequency orthogonality established between the desired signal and jamming signal, a process we discuss in the next subsection. Under this assumption, the destination node can separate the two signals, extract features from the jamming signal, and perform jammer classification. The classification results must be relayed back to the source node. Therefore, it is assumed that there is, at a minimum, a low data rate channel that the destination node can use to share the classification results with the source node (typically referred to as the feedback or return channel). The number of bits associated with classification results can be reduced using a predefined lookup table. A probing function is used to send known symbols (a.k.a. pilots) that are designed to help distinguish among the various jamming models discussed in Section 4.4. We propose using machine learning classification, where the features used for classification consist of:

Demodulation of Pilot Symbols: a straightforward method of detecting if the strategy *phase flipping* or *modulate amplitude* is in use. Using QAM modulation on a selection of the pilot symbols, such as 16-QAM, would allow the destination node to determine if the jammer is modulating the phase (the amplitude portion of the constellation would remain intact) or modulating the amplitude (the phase information would remain intact).

Crest Factor (CF): a feature of the jamming signal in the frequency domain, which can be used to differentiate between *replace with CW* and the other models. CF indicates the ratio of peak value to the quadratic mean: $CF = |x|_{peak}/x_{rms}$. Note that peak-to-average power ratio (PAPR), an important parameter in wireless communications, is the square of the CF.

Cross-correlation: simply the correlation of y with y after rough timing realignment, in which a strong peak would indicate the *DRFM* model is in use. A cross-correlation (i.e., sliding dot product) is used instead of just one dot product to take into account the fact that $\hat{d}$ is likely only accurate enough for rough timing realignment, not down to the symbol or sample level. We formulate this feature as follows:

$$CC = \max_{n \in \{-\epsilon,\ldots,-1,0,1,\ldots,\epsilon\}} \sum_{m=0}^{m_{len}} y^*[m]y[m - \hat{d} + n] \tag{4.9}$$

where ϵ corresponds to this maximum anticipated error of $\hat{d}$ in either direction of time, and m_{len} is the hop/frame length.

While none of the above features explicitly detect the *replace with noise* jamming model, we can treat this model as the *default* jamming behavior that is assumed if none of the other jamming models seem to be appropriate for the scenario. This makes sense because the *replace with noise* model makes the fewest assumptions about the jammer's behavior, thus following Occam's razor.

Based on extracting these three features, the classifier can determine the specific model that most closely matches the jammer's behavior. Multiple hops worth of observations should be used for accurate classification, and features must be extracted from each hop individually. The transitory phase due to delay estimation and jammer classification may span several hops, but it is likely insignificant compared to the period of time the reactive jammer is active. The specific classifier we suggest using is the Support Vector Machine (SVM), due to its superior performance across a large range of applications [84] (for a comparison between classification methods, we refer the reader to [85]). Training is performed offline, using several instances of each jammer model (either simulated or implemented in hardware). Retraining could also be performed in the field if a new threat is discovered, or to take into account minor variations to the existing threats.

4.5.3 Forcing Orthogonality

To avoid the transmitted signal and jamming signal being received co-channel (overlapping in frequency and time), we must force orthogonality using observations of the jammer's delay and *ending-lag*. The process of causing orthogonality between a desired signal and a jamming signal is at the heart of anti-jamming, and as such, we make use of common anti-jamming strategies as part of the antifragile waveform. We assume that the communications link under study is able to null certain symbols (e.g., by assigning the value $0 + 0j$ to them). The proposed solution to forcing orthogonality is based on which of the three delay cases occur.

When $d < T_{sym}$, true time-frequency orthogonality is simply not possible because symbols are overlapping, causing co-channel signals. However, if the jammer-to-signal ratio (JSR) is significantly high, then orthogonality is not needed to achieve an antifragile gain, as combining is not necessary and the destination node can use the jamming signal and treat the actual signal as noise. Likewise, if a phase array antenna can isolate the two signals, then orthogonality is achieved *spatially*.

When $T_{sym} \leq d < T_{hop}$ and the radio is hopping as fast as it can, it must null its e last

symbols of each hop or frame, where $e = \frac{T_{hop}-d}{T_{sym}}$. In other words, the source node refrains from transmitting on symbols that would overlap with the jamming signal. If the fraction of overlapping signals is large, methods from the previous delay case can be used.

When $d \geq T_{hop}$, orthogonality can be created by hopping in such a way that there are no *collisions* in time and frequency on a hop-basis. With this approach, the actual signal and jamming signal are received orthogonally, and the signals can be separated and demodulated independently.

4.5.4 Modulation Scheme for the Antifragile Waveform

The fundamental thrust of any jammer-based antifragile communications method is to exploit the signalling dimensions enabled by the jamming system's transmissions, which may manifest in time, frequency, or space. Confining ourselves in this discussion to scenarios where the transmitted signal and jamming signal are received orthogonally, the optimal modulation scheme is intimately tied to which dimensions the jammer transform function either preserves from the transmitted signal or introduces to potential signalling by its operation. Consequently, an accurate jammer model (see Section 4.4) is critical to realizing a useful communications link, and effective classification methods are necessary to exploit encountered systems.

Several common jammer types do not preserve incoming signal amplitude and phase information, instead transmitting a newly generated signal such as a noise-like waveform or CW tone in the channel. If the jamming signal is sufficiently narrowband and its duration consistently deterministic, the source transmitter can effect a noncoherent pseudo-frequency-shift keying (FSK) relay link by modulating the frequency of its own transmissions. Similarly, a frequency-preserving jammer transform may support chirp modulation types. Alternately, a consistently deterministic jammer duration alone enables noncoherent modulation schemes such as straightforward OOK or the family of Pulse Position Modulation (PPM) types, which may accommodate differential coding, overlapping, or other features. An example of PPM transmitted *through* the jammer is shown in Figure 4.8, with the *ending-lag* highlighted in red.

Depending on the jammer architecture, the system may channelize its wideband received input in order to monitor multiple subchannels at the same time. Therefore, if the antifragile radio can determine the specific subchannel configuration used by the jammer, it has the potential to achieve simultaneous modulation on multiple subchannels, resulting in an overall transmit scheme that resembles Frequency Division Multiplexing (FDM). The radio must simply provide a guard band between exploited subchannels.

The phase flipping jammer model discussed earlier may disrupt conventional phase signalling schemes. One solution is to avoid phase information entirely by using what we call "Positive-ASK", which is ASK with constellation points only in the positive half of the real axis.

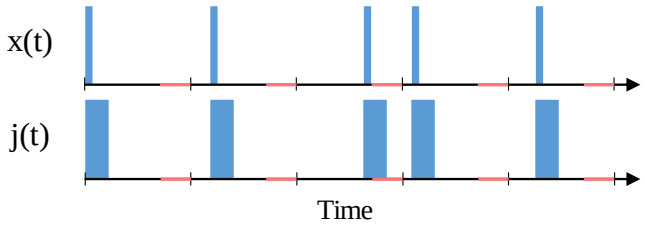

Figure 4.8: Example waveform when using PPM. Each tick represents a symbol period, with the red portion of each symbol representing the *ending-lag*. The actual signal and jamming signal have been aligned in absolute time (i.e., $d = 0$) for the sake of presentation.

Table 4.3: Modulation assignments for each jammer model.

Jammer Model	**Corresponding $j(t)$**	**Modulation Scheme**
DRFM	$\beta_A(k_e x + n_j)$	Default Scheme
Replace with Noise	$n_{jam}(t)$	FSK, chirp, OOK, PPM
Phase Flipping	$U(t)\beta(k_e x + n_j)$	Positive-ASK
Modulate Amplitude	$V(t)\beta(k_e x + n_j)$	PSK
Replace with CW	$cos(2\pi f t)$	FSK, chirp, OOK, PPM

Postive-ASK can be thought of as the ASK equivalent of single-polarity Pulse-Amplitude Modulation (PAM). Conversely, for the relatively unusual amplitude-only modulating jammer, a phase modulation approach (e.g., 8PSK) would work.

Because it represents an amplify-and-forward relay, the DRFM jammer model is the simplest to deal with, requiring no change in the original waveform.

A summary of the modulation scheme assignments for each jammer model is given in Table 4.3.

4.5.5 Combining Technique

Since the goal of this antifragile scheme is to relay some or all of the source node's data using the jammer, the destination is presented with two different signals carrying the same information. Thus, we have a similar situation to receive diversity in multiple-input multiple-output (MIMO) communications, with the exception that the two signals are likely received at different power levels. Within diversity combining there are three common techniques: selection combining (SC), maximal-ratio combining (MRC), and equal gain combining. In SC, the receiver simply selects the signal that is received with the highest SNR (or other channel quality indicator). MRC weights each received signal before combining, in such a way that the combined SNR is maximized (thus being the optimal scheme in terms of SNR).

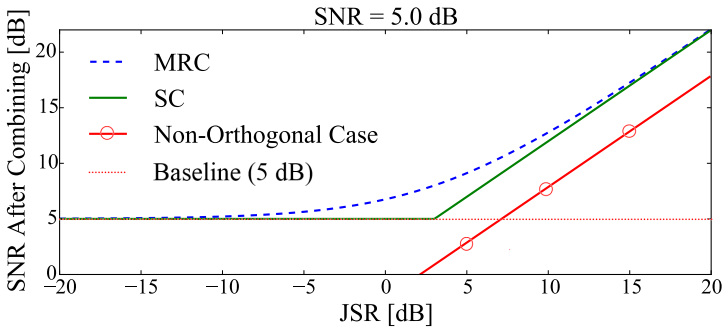

Figure 4.9: Average SNR after combining, for maximal-ratio and selection combining, when the main channel SNR is a constant 5 dB. Also included is the SNR when only decoding the jamming signal, treating the actual signal as noise, which would only make sense if orthogonality is not achievable and JSR is high.

It has been shown that when using MRC, the combined SNR is simply the summation of the individual SNRs when in linear form [86]. Equal gain combining is a special case of MRC in which the weights applied to each signal are equal (usually set to 1). Analysis of these combining techniques with unequal SNR is performed in [86]. Because the SNR of the actual signal and the effective SNR of the signal relayed through the jammer could vary greatly, equal gain combining does not make any sense for this application. MRC and SC are compared in Figure 4.9, under a constant main channel SNR of 5 dB and a varying JSR. To achieve an antifragile gain in this example, the SNR after combining must be above 5 dB. It is assumed that the channel noise at the jammer's receiver is the same as the noise at the destination node, and fading is not taken into account.

Also included in Figure 4.9 is a case where orthogonality is not achieved, so instead of combining the two signals, the destination node simply decodes the jamming signal and treats the actual signal as noise. In this case, the JSR must be at least 7.5 dB for the strategy to function.

4.6 Theoretical Channel Capacities

We will now find the theoretical channel capacity under each jammer model described in Section 4.4.2, when using the jammer piggybacking antifragile strategy. The channel capacities in this section are primarily provided to demonstrate the different parameters associated with the process of piggybacking off each jammer type, and how throughput performance is influenced.

The channel capacity when no jammer is present, which will act as a baseline to determine when the antifragile requirement is met, is given by Shannon's theorem [87]:

$$C = B \log_2 (1 + \gamma_s) \tag{4.10}$$

where C is channel capacity in bits per second and B is channel bandwidth in Hz (i.e., the passband bandwidth of the signal). In Equation 4.10, the signal is assumed to be a wide-sense stationary zero mean complex Gaussian random process.

For each jammer model, the channel capacities incorporate a factor that represents that fraction of orthogonality achieved, on a per-hop or per-frame basis (see Section 4.5.3 for more information on timing). We will denote this fraction as D_{orth}, where $D_{orth} = 1$ corresponds to zero overlap in the actual signal and jamming signal, and $D_{orth} = 0$ corresponds to full overlap. As such, this parameter appears as a multiplier in the channel capacity equations. Lastly, we assume MRC is the form of combining used.

The simplest jammer model to analyze is the **DRFM**, which resembles an amplify-and-forward relay. The effective relay channel (i.e., the channel through the jammer) is a combination of channels k_e and k_j. The SNR of this total link through the jammer, which we will denote as γ_{ej}, is based on the value of β_A. In relaying literature, it is assumed that the relay's amplification gain is adjusted in realtime to satisfy the relay's (or jammer's) maximum output power constraint, denoted as $|j|^2$ [88]. Under this condition, β_A will always equal:

$$\beta_A = \sqrt{\frac{|j|^2}{|x|^2 |k_e|^2}} \tag{4.11}$$

While this may not always be the case for a uncooperative jammer, it is not an unreasonable assumption, given that the jammer may also want to maximize its output power. Equation 4.11 allows us to approximate the SNR of the link through the jammer [89]:

$$\gamma_{ej} = \frac{\gamma_e \gamma_j}{\gamma_e + \gamma_j + 1} \tag{4.12}$$

When MRC is used, the overall SNR is simply $\gamma_s + \gamma_{ej}$. Therefore, the overall channel capacity for the DRFM case is:

$$C = D_{orth} B \log_2 \left(1 + \gamma_s + \frac{\gamma_e \gamma_j}{\gamma_e + \gamma_j + 1} \right) \tag{4.13}$$

When a **phase flipping** jammer is present, we transmit in a manner that does not involve the imaginary portion of the constellation, as discussed in Section 4.5.4. Therefore, we must use the channel capacity formula for a single dimension, instead of complex [90]. When using MRC, the resulting channel capacity becomes:

$$P_e = \frac{1}{k}\left[\sum_{i=1}^{M/2}\frac{w_i^{'}}{2\pi}\left(\int_0^{\pi-\frac{2i\pi-\pi}{M}}\exp\left(-\gamma\frac{\sin^2[(2i-1)\pi/M]}{\sin^2\theta}\right)d\theta - \int_0^{\pi-\frac{2i\pi+\pi}{M}}\exp\left(-\gamma\frac{\sin^2[(2i+1)\pi/M]}{\sin^2\theta}\right)d\theta\right)\right] \tag{4.15}$$

$$C = \frac{D_{orth}B}{2}\log_2\left(1+2\gamma_s+\frac{2\gamma_e\gamma_j}{\gamma_e+\gamma_j+1}\right) \tag{4.14}$$

Exploiting an **amplitude modulating** jammer involves using M-PSK, as discussed in Section 4.5.4. The bits per second when using M-PSK is simply $C = \log_2 M$. The maximum reliable value of M is a function of SNR, which can be approximated using the theoretical bit error rate for uncoded M-PSK [91], given in Equation 4.15. Finding an upper limit on C requires taking into account all combinations of M and channel coding schemes, which is not feasible. In Section 4.7, when feasibility regions are evaluated, we simply use Equation 4.15 and require the bit error rate to be below 10%.

The **replace with noise** and **replace with CW** jammers involve signal detection, and transmit an internally generated signal, $w(t)$, as the jamming waveform. Thus, the amount of information relayed through the jammer is partially based on the accuracy and speed of the jammer's detection process. If we assume the jammer has no a priori knowledge of the target waveform, then the only difference between observing a signal and observing noise is the statistical average energy they contain. Therefore, the optimum detector compares the average energy in an observed waveform to a threshold, also known as an energy detector or radiometer [92]. Probability of detection, P_D, of a Neyman-Pearson type energy detector is parameterized by eavesdrop SNR, γ_e, and number of samples, and given by [92]:

$$P_D = 1 - \Gamma\left(\frac{n}{2};\Gamma^{-1}\left(\frac{n}{2};1-P_{FA}\right)(1+\gamma_e)^{-1}\right) \tag{4.16}$$

where $\Gamma(x,y)$ is the incomplete gamma function, P_{FA} is the probability of false-alarm, and n is the number of samples taken from the observed waveform.

The amount of information relayed through the jammer is also based on the symbol rate, which must be decided on by the source node, and is a function of the jammer's *ending-lag* (so that symbols through the jammer do not overlap, see Section 4.5.1 for more information). If we assume the symbol period is equal to the *ending-lag* plus a safety margin factor, which we will denote as T_{guard}, the resulting symbol rate is $(T_{end-lag}+T_{guard})^{-1}$. This safety factor can be used to account for any jitter associated with the jammer's delay d. In a practical system, there would be digital signal processing required to verify the symbol rate is usable.

We formulate the bit rate based on using OOK with N number of multiplexed signals, where N is an integer greater than zero. We must assume that the jammer's detection process is very accurate, else we cannot formulate an equation for the rate at which information can be *reliably* transmitted. If this is the case, the capacity is:

$$C = N\left(T_{end-lag} + T_{guard}\right)^{-1} \quad \text{if} \quad \begin{matrix} P_D \approx 1 \\ P_{FA} \approx 0 \end{matrix} \tag{4.17}$$

Simply put, the data rate is equal to the number of multiplexed streams multiplied by the symbol rate, but only if the jammer is able to accurately perform signal detection. While no SNR appears in Equation 4.17, it should be noted that γ_e is a parameter in Equation 4.16, while the other two SNRs must simply be high enough for reliable transmission of OOK, which is around 0 dB [93]. Equation 4.17 is valid for both the *replace with noise* and *replace with CW* jammer models, because when using OOK in such a manner, the bandwidth of the jammer's signal is not a factor (as long as the signal can be demodulated at the destination node). Rather, the symbol rate is a function of the jammer's timing.

4.7 Numerical Results

In this section we provide numerical results showing the feasibility regions in which jammer piggybacking provides an antifragile gain. Feasibility is evaluated by comparing throughput with the baseline (jammer-free) case.

4.7.1 Simulation Scenario and Conditions

The simulation scenario matches Figure 4.2, with the jammer taking on each of the models described in Section 4.4 (although *replace with noise* and *replace with CW* are combined, because they provide the same results).

We vary JSR and main channel SNR, γ_s, instead of all three SNRs, for the sake of two-dimensional results. The only requirement is that we assume $n_j = n_d$, so that $JSR = \gamma_j/\gamma_s$ for a sensing jammer, and $JSR = \gamma_{ej}/\gamma_s$ for a repeating jammer. We have decided to use JSR as the x-axis and produce plots for a SNR of 3 dB and 10 dB.

Adaptive modulation and coding (AMC) is used, and we assume an additive white Gaussian noise (AWGN) channel. For a given set of JSR and SNR, the best modulation scheme, modulation order, and code rate is evaluated based on a pre-populated table (see Section 2.4 regarding this method). The only constraint is that the bit error rate after decoding must be below 10^{-3}. To provide numerical results that reflect a high-performance communication system, low-density parity-check (LDPC) codes with code rates spanning $1/6$ to $16/17$ are used. In terms of modulation schemes available to the AMC engine, *phase flipping* uses

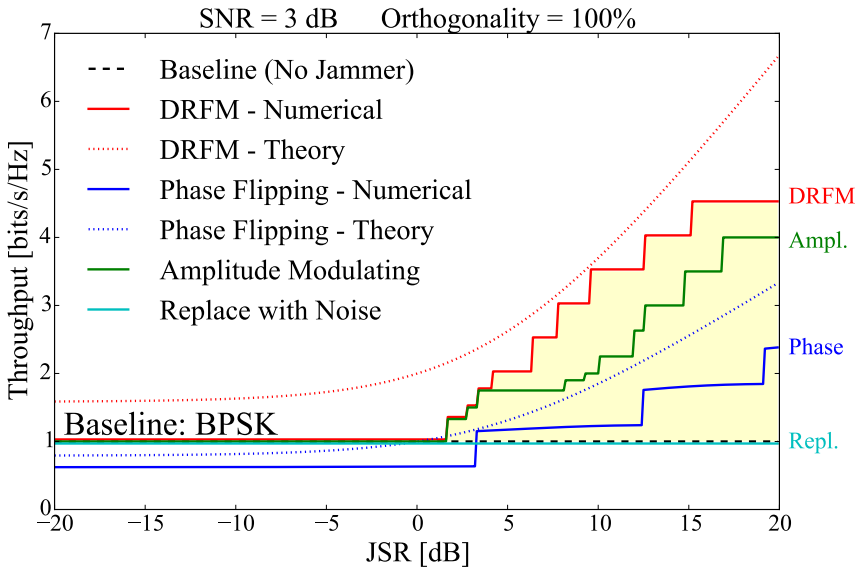

Figure 4.10: Feasibility region of reactive jammer piggybacking when the main channel SNR = 3 dB. Each sharp increase corresponds to the modulation and coding scheme being adapted for a higher quality channel. A curve above the baseline indicates an antifragile gain, as highlighted in yellow for the DRFM case.

Positive-ASK, *amplitude modulating* uses PSK, and *replace with noise* uses OOK. Both the baseline (no jammer present) and *DRFM* cases can choose from PSK, QAM, and ASK.

4.7.2 Simulation Results

Figures 4.10 and 4.11 show simulation results under full orthogonality, which can occur when the delay is greater than one hop, or delay is less than one hop and there is spatial orthogonality. Each point at which the modulation and coding scheme changes is identified by sharp transitions in the curve. We remind the reader that the baseline in each plot refers to the communications link when there is no jammer present; the baseline modulation and coding scheme is purely based on the main channel's SNR, and thus is constant in each plot. The theoretical channel capacities derived in Section 4.6 are shown for the DRFM and phase flipping case, representing the upper bound on throughput given the waveform constraints (the other two capacities cannot be represented without additional assumptions).

For both levels of SNR, the DRFM case has the best performance, which is expected considering the lack of constraints on the waveform under this case. A loss is associated with having to use PSK (in response to an *amplitude modulating* jammer) and Positive-ASK (in response to a *phase flipping* jammer), although the loss is not large enough to prevent an antifragile gain altogether when SNR = 3 dB. Being forced to use OOK (in response to a *replace with noise* jammer) leads to a lack of an antifragile gain across the entire region, for

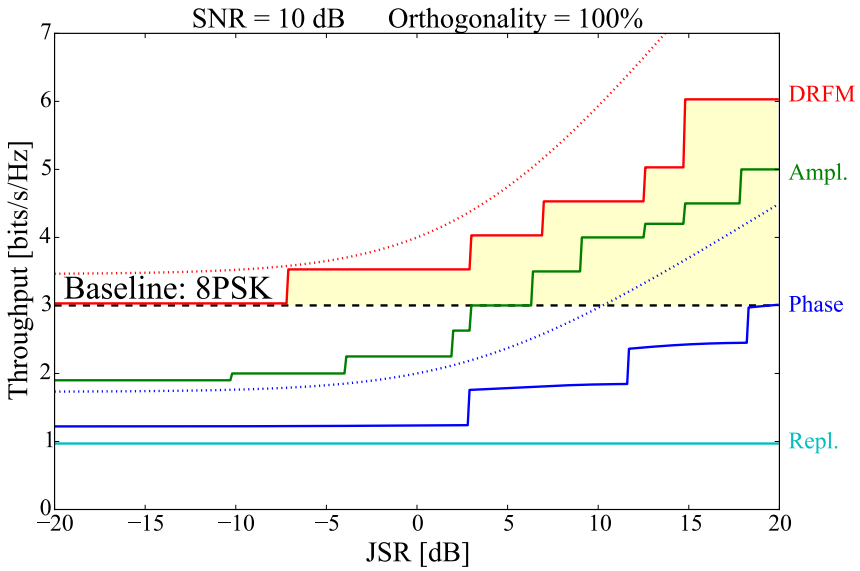

Figure 4.11: Feasibility region when the SNR = 10 dB.

all levels of SNR. However, these results reflect only having *one* data stream of OOK. As discussed before, it is likely possible (depending on the scenario) to transmit a large number of data streams in parallel using FDM.

It is expected that the antifragile gain mostly occurs in the right-hand portion of each plot, where the JSR is higher than zero. The antifragile gain when JSR is negative is largely due to the combining gain, while the gain when JSR is positive is due to the fact that the jamming signal reaches the destination at a much higher effective SNR than the main channel. Fortunately for this strategy, jammers typically intend to operate in the right hand portion of the plot, where JSR is above 0 dB (communications with adaptive modulation and coding tend to become denied at a SNR lower than 0 dB).

When full orthogonality cannot occur, such as when the delay d is between one symbol and one hop, a loss must be included to take into account the nulled symbols at the end of each hop. This loss simply shifts each curve down by a certain percent. Figure 4.12 shows an example scenario when SNR = 3 dB, and the actual signal and jamming signal overlap during 50% of each hop. The result is a 50% loss in throughput for each antifragile curve. It can be seen that the start of the antifragile gain shifts to the right, requiring higher jammer power compared to a fully orthogonal scenario.

In cases where an antifragile gain in the context of throughput is not feasible, it may still be possible to achieve antifragility by using the jammer to provide a low data rate control channel, as discussed at the end of Section 4.3.1.

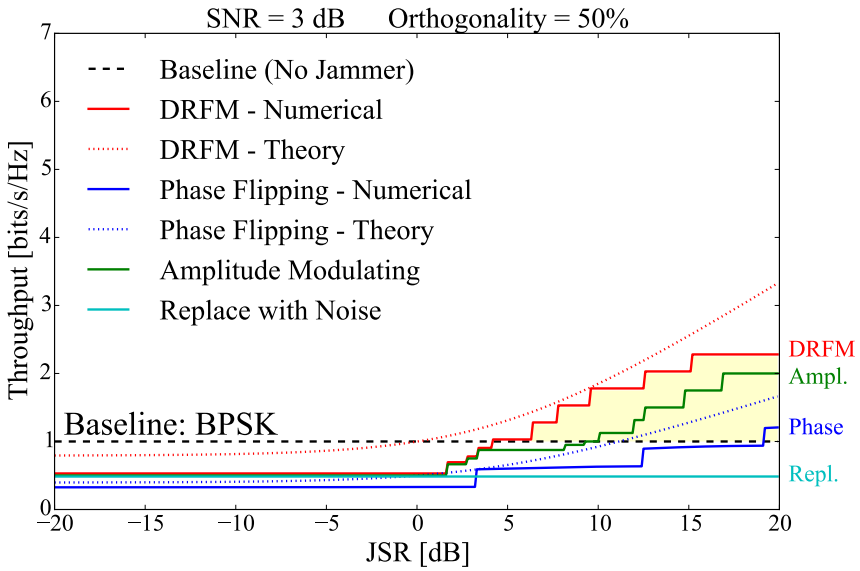

Figure 4.12: Feasibility region when the SNR = 3 dB and there is only 50% orthogonality.

4.8 Conclusion

In this chapter, we have introduced the concept of antifragile wireless communications, through a novel strategy that exploits a reactive communications jammer. Several reactive jamming models were described, including both repeater-based and sensing-based. We have outlined guidelines for realizing an *antifragile waveform*, and shown that an antifragile gain is possible under a wide variety of reactive-jamming scenarios.

As the sophistication of communications systems and jammers increases, reactive jamming will likely become a bigger threat in military and other mission-critical domains. Therefore, incorporating antifragility will not only improve protection of radios, but also bring about performance improvements in harsh conditions. In addition, a system that is able to achieve any level of antifragility disincentivizes an enemy jamming mission in its entirety, assuming that mission has nonzero cost. Even if the enemy is aware of its target's antifragile capability and prevents it from achieving an antifragile gain (e.g., by using basic barrage jamming), that antifragile capability indirectly leads to a less effective attack.

As part of future research we will continue to develop the concept of antifragile communications, including an investigation into how fading affects the performance of jammer piggybacking. There are several areas in which the concept can be expanded into future topics, including network-layer strategies, spatial signalling dimensions, reduction or elimination of required feedback signals, jammer herding, multi-source signalling, and interference alignment. In the case of the *replace with noise* jammer, the optimal symbol rate and number of frequency multiplexed signals is still an open question. Lastly, a deeper investigation into the antifragile schemes proposed in Section 4.3.2 could lead to additional applications.

Appendix: Bit Error Rate Equations

This appendix provides the Bit Error Rate (BER), P_e, for each uncoded modulation scheme through an AWGN channel, given the SNR, γ. These expressions were used in some of the simulation results presented in this chapter. We assume that the SNR γ is equal to the symbol energy E_s divided by the noise energy N_0.

Equation 4.18 shows the BER for **M-ASK** (also known as M-PAM) [94].

$$P_e = \frac{2}{M \log_2 M} \sum_{k=1}^{\log_2 M} \sum_{i=0}^{(1-2^{-k})M-1} \left[(-1)^{\left\lfloor \frac{i2^{k-1}}{M} \right\rfloor} \left(2^{k-1} - \left\lfloor \frac{i2^{k-1}}{M} + \frac{1}{2} \right\rfloor \right) Q\left((2i+1) \sqrt{\frac{6\gamma}{M^2-1}} \right) \right] \tag{4.18}$$

where $Q()$ is the Q-function equal to $\frac{1}{2} erfc(\frac{x}{\sqrt{2}})$, also known as the tail probability of the standard normal distribution.

The BER for **M-PSK** is given by [91]:

$$P_e = \frac{1}{k} \left[\sum_{i=1}^{M/2} \frac{w_i'}{2\pi} \left(\int_0^{\pi - \frac{2i\pi - \pi}{M}} \exp\left(-\gamma \frac{\sin^2[(2i-1)\pi/M]}{\sin^2 \theta} \right) d\theta - \int_0^{\pi - \frac{2i\pi + \pi}{M}} \exp\left(-\gamma \frac{\sin^2[(2i+1)\pi/M]}{\sin^2 \theta} \right) d\theta \right) \right] \tag{4.19}$$

The BER for **M-QAM** is given by [94]:

$$P_e = \frac{2}{\sqrt{M} \log_2 \sqrt{M}} \sum_{k=1}^{\log_2 \sqrt{M}} \sum_{i=0}^{(1-2^{-k})\sqrt{M}-1} \left[(-1)^{\left\lfloor \frac{i2^{k-1}}{\sqrt{M}} \right\rfloor} \left(2^{k-1} - \left\lfloor \frac{i2^{k-1}}{\sqrt{M}} + \frac{1}{2} \right\rfloor \right) Q\left((2i+1) \sqrt{\frac{6\gamma}{2(M-1)}} \right) \right] \tag{4.20}$$

The BER for **binary phase-shift keying (BPSK)** and **On-Off Keying** is given by [93]:

$$P_e = Q\left(\sqrt{2\gamma} \right) \tag{4.21}$$

The BER for **quadrature phase-shift keying (QPSK)** is given by:

$$P_e = Q\left(\sqrt{\gamma} \right) \tag{4.22}$$

Chapter 5

Energy-Based Antifragile Gain

5.1 Introduction

In the previous chapter we investigated the potential for a throughput-based antifragile gain, under a variety of reactive jamming behaviors. However, the analysis was primarily system-level, intended to apply to many different jamming behaviors. We will now focus solely on the sense-and-transmit type jammer model, which monitors a large number of subchannels and reactively jams each one independently. In other words, the jammer only jams a specific subchannel when it senses energy or a signal on that subchannel. By doing so, the jammer greatly reduces its power consumption, and remains difficult to detect due to being hidden under or adjacent to another signal. Section 3.3.1 introduced this type of jammer, and Chapter 9 is devoted to the design of an orthogonal frequency-division multiplexing (OFDM)-based version of the jammer. Achieving a throughput gain for this type of jammer is impractical due to the long response time associated with the jammer's behavior. As such, we focus on achieving an energy-based gain, where we seek to harness the jammer's signal power and use it for our own communications. This can lead to significantly lower average transmission power for the communication system, and thus, lead to energy savings with respect to normal operation. An energy-based antifragile waveform is well-suited for energy-constrained nodes, such as those in a wireless sensor network. It could also be useful for a radio that wants to minimize its transmitted power, to reduce the probability of being detected/intercepted.

In this chapter we describe a frequency-shift keying (FSK) type waveform that can be used to exploit a multichannel sense-and-transmit jammer, turning it into an unwitting relay. We develop the theoretical capacities, cutoff rates, and energy savings associated with the waveform. Using numerical results, we show under what conditions an antifragile gain can be achieved.

The remainder of this chapter is organized as follows. We first introduce the jammer exploitation strategy and discuss how the waveform parameters can be found. We then describe the system model and assumptions made in analyzing the waveform performance. In Section 5.4 we derive the theoretical and practical channel capacity under successful piggybacking using the cutoff rate. Numerical results under a variety of scenarios are provided in Section 5.5. In Section 5.6 we discuss additional design guidelines to be considered during the design and implementation of such a waveform. Section 5.7 concludes.

5.2 Jammer Exploitation Strategy

The jammer exploitation strategy described in this chapter can be thought of as an M-ary FSK waveform, tuned to the jammer's timing and frequency configuration. We seek to cause the jammer to effectively relay this FSK signal, such that it combines with the source node's signal at the destination node. Thus, we are assuming the jammer has a fairly quick response

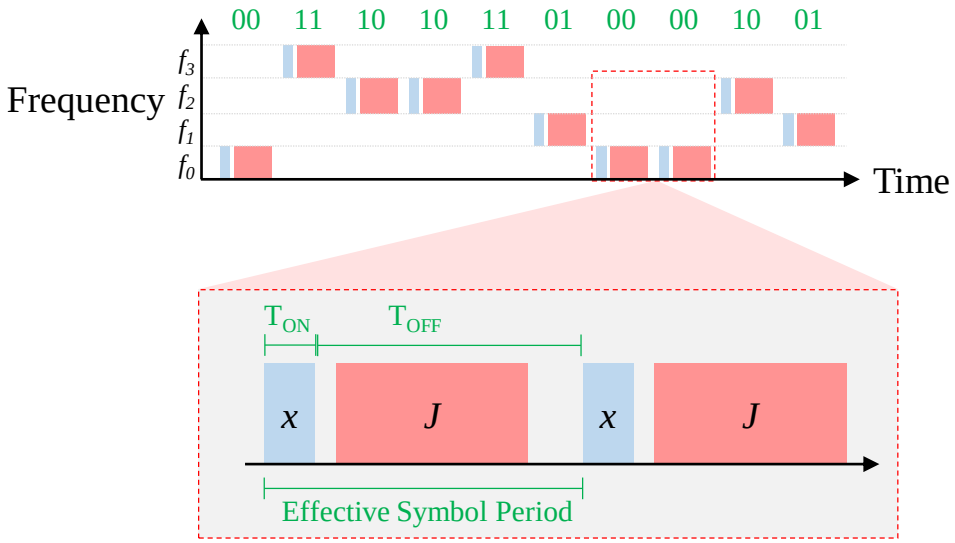

Figure 5.1: Example of the jammer exploitation strategy that seeks to minimize transmitted energy per bit. By piggybacking off the jammer using appropriately timed M-FSK ($M = 4$ in this example), the jammer effectively extends each symbol, increasing the energy per bit at the destination node.

time (on the order of microseconds to milliseconds), in order for this strategy to have any practicality. A reactive jammer that has a slow reaction time would translate into our bit rate being capped at an excessively low value. Under this waveform, the source node generates one of M orthogonal signals. In order to exploit the sense-and-transmit jammer in an antifragile manner, the source node will intentionally insert gaps between symbols, in such a way that the jammer's waveform *extends* each symbol, effectively adding energy to it. The destination node uses non-coherent detection, because achieving phase coherency between the two signals at the destination node is not practical. This form of jammer piggybacking requires finely tuned waveform parameters, so that the jammer's signal does not overlap with the next FSK symbol at the destination node. Let us treat each symbol transmitted by the source node as a pulse. We will seek to minimize the pulse duration, as well as the gap between pulses, which combined form the symbol period. Figure 5.1 shows an example of this strategy, with $M = 4$ (although in many scenarios, much greater values of M are possible).

The transmitted signal, $x(t)$, has two important parameters: the pulse length, and the time between pulses. We will denote these parameters as T_{ON} and T_{OFF} respectively. Note that T_{ON} plus T_{OFF} provides the symbol period. Our goal is to reduce the duration of T_{ON}, so as to minimize the transmitted energy per bit, as well as tune T_{OFF} to maximize the channel capacity. This represents a trade-off; because we are using non-coherent detection, we want each symbol to contain as much energy as possible, but extending the symbol leads to fewer symbols per second. In addition, appropriate values of T_{ON} and T_{OFF} are a function of the jammer, which we must probe as part of the proposed strategy. In the next subsection we

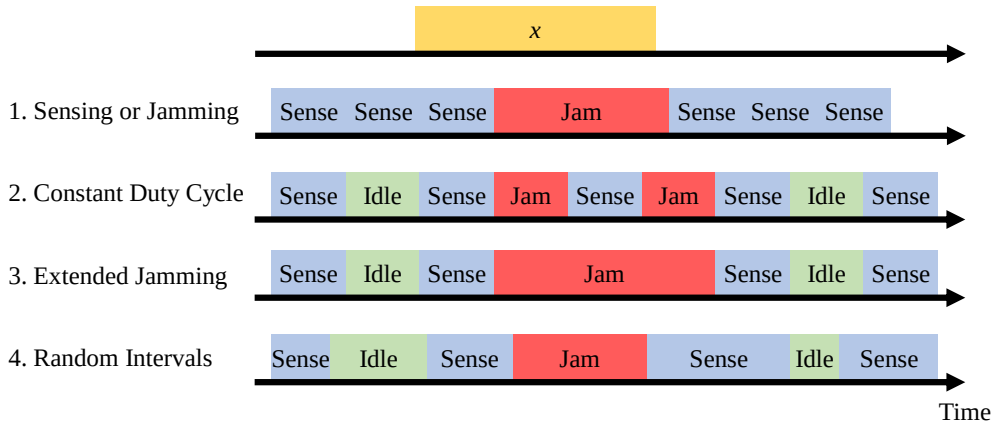

Figure 5.2: Example of four different approaches to sense-and-transmit type reactive jamming. In this example, the constant duty cycle type is using a sensing duty cycle of 50%, which was chosen arbitrarily.

discuss how these two parameters can be determined, given a jammer's behavior.

5.2.1 Determining FSK Parameters

Instead of thinking about $x(t)$ as part of an FSK symbol, with the purpose of communicating information, we can think about it as a signal designed to trigger one subchannel of a reactive jammer (it actually serves both purposes simultaneously). Thus, finding the minimum value of T_{ON} is equivalent to finding the minimum pulse length that will trigger the jammer's detection process. For this problem, we need to consider the different sense-and-transmit schemes the jammer may use, as well as the percent of time signal $x(t)$ is received within the sensing window.

With a sense-and-transmit type reactive jammer, it is assumed that the jammer does not transmit during the sense duration (i.e., the jammer is half-duplex). There are four types of sense-and-transmit patterns the reactive jammer could use, listed below and depicted in Figure 5.2.

1. Always sensing or jamming (never idle)
2. Constant duty cycle, switching between sensing and jamming/idle
3. Extended jamming duration, where the jamming time is longer than the sense time
4. Randomized sensing and jamming/idle intervals

The first type is the typical behavior of a reactive jammer that only jams one channel and

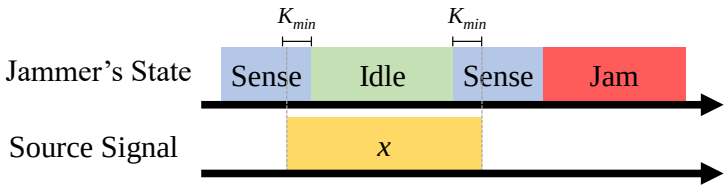

Figure 5.3: Visual representation showing the minimum pulse length required to guarantee detection triggering, with $K_{min} = 0.3$.

is, therefore, not considered in the remainder of this chapter. As discussed before, reactive jammers often channelize the spectrum so they only have to jam the active frequencies. In this case they must have silent periods to perform sensing, so that the jammers transmissions do not interfere with the sensing, or because the jammer is not designed to transmit and receive simultaneously. The last three types reflect this behavior.

Let T_{sense} and T_{idle} represent the duration of the jammer's sense and idle periods respectively. We assume there is no synchronization between the jammer and transmitted signal. Next, let us assume that to reliably detect the signal, the signal x must overlap with at least K_{min} of the jammer's sense period, where K_{min} is a fraction between zero and one. Figure 5.3 depicts an example showing the worst-case scenario, where the source node's transmission starts a little too late within the first sense window, and must be detected during the second sense window. In this example, K_{min} was arbitrarily set to 0.3. Now let us switch to the perspective of the antifragile radio, trying to trigger the jammer's detection process. For the first three types listed above, to guarantee overlap with K_{min} of at least one sensing window, the pulse length must be at least:

$$T_{ON} \geq 2K_{min}T_{sense} + T_{idle} \tag{5.1}$$

where $T_{idle} = 0$ for the first type. The expression associated with the randomized pattern requires a more complex model for the jammer, and will be left for future work.

In a practical scenario, the source node won't actually know the values of T_{sense}, K_{min}, and T_{idle} a priori. Instead of trying to estimate these values, we propose that the source node begins by transmitting its FSK signal with a very long T_{OFF} (e.g., 10 ms), and very short T_{ON} (e.g., 10 μs). It can then increase T_{ON} each pulse until feedback from the destination node is received, indicating a functioning link.

The value of T_{OFF} simply determines how long before the next pulse is sent. We seek to have the next pulse arrive at the destination node after the previous jamming signal has ended, so that there is no intersymbol interference. This is complicated by the fact that we do not know when in the sense-transmit cycle our pulse landed (or even the cycle itself, in many cases). To be sure there is no intersymbol interference, we will investigate the worst-case scenario, depicted in Figure 5.4. This is the case when the pulse ends at the same time the

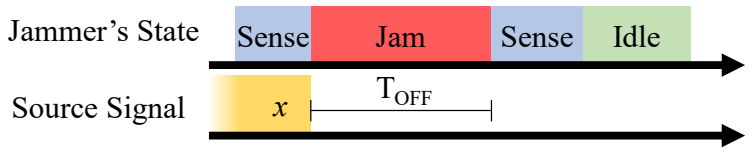

Figure 5.4: Visual representation showing why the minimum T_{OFF} must be at least T_{jam}.

jammer's sense window ends, resulting in the longest gap between pulses. As shown in the figure, the value of T_{OFF} must be at least as long as the jamming duration, which we will denote as T_{jam}.

To determine an effective value for T_{OFF}, the source node uses the strategy discussed in the previous subsection, but once it finds T_{ON}, it slowly reduces the length of T_{OFF} until the connection is degraded.

At this point, it should be clear that the data rate associated with this jammer piggybacking strategy is fairly low, as the symbol periods are on the order of the jammer's sense-transmit cycle, which could be tens of ms. However, this strategy emphasizes an energy-savings gain, not a throughput gain, relative to a jammer-free case.

5.3 System Model and Assumptions

To analyze the performance of the proposed strategy, we examine the system under Rician fading between the source and destination nodes, and an additive white Gaussian noise (AWGN) channel between the jammer and destination node. The Rician channel was chosen to model a moving source node in a multipath environment. While it is quite possible for the jammer and destination node to be moving as well, an AWGN channel was chosen to reduce the complexity of the analysis, as well as the parameter space.

The theoretical capacity is determined as the maximum data rate in which reliable information can be communicated, which implies the error probably is arbitrarily small. Naturally, this assumes a form of channel coding (a.k.a. forward error correction) is used. In order to remain agnostic of the specific form of channel coding, while still taking into account the presence of powerful channel coding, we assume optimal coding is used. Two decades ago this may have been a strong assumption, but it is not anymore given the powerful turbo and low-density parity-check (LDPC) codes that exist and are realizable in low-cost hardware.

We assume an M-FSK signal, where each symbol is made up of a transmission by the source node, denoted as $x(t)$, and one by the jammer, denoted as $J(t)$. Figure 5.5 shows the system diagram, including all signals and channels involved. These two signals are combined to form one total FSK symbol. As in conventional FSK, there are $\log_2 M$ bits per symbol. We assume an unknown phase of both signals at the destination node, and thus are only

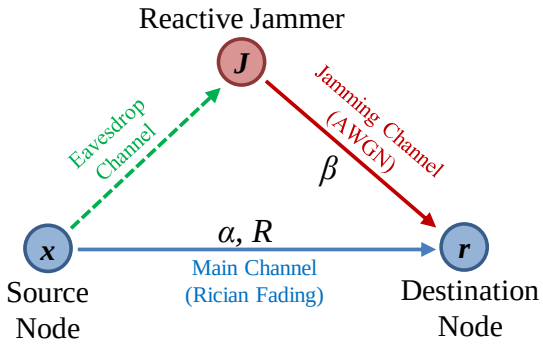

Figure 5.5: The system diagram, showing all signals and channels involved, as well as channel type used for analysis.

considering non-coherent detection. In addition, we assume soft decision detection. Because of our channel models, the received signal at the destination node consists of the following components:

1. The line-of-sight (LOS) component of the source node signal $x(t)$ with a constant attenuation α (primarily based on path loss) and arbitrary phase shift.
2. The scattering component of the source node signal $x(t)$ with Rayleigh distributed attenuation R and phase uniformly distributed between $[0, 2\pi]$, due to fading.
3. The jamming signal $J(t)$ with constant attenuation β and arbitrary phase shift.
4. White Gaussian noise component with two-sided spectral density $N_0/2$.

Table 5.1 summarizes all of the symbols used in this chapter.

We assume a memoryless channel, and that the Rayleigh distributed attenuation R is constant over one symbol, which for $x(t)$ means over the interval $[0, T_{ON}]$. This also means the values of R are assumed to be independent, which is typically not the case in an actual system, but what the memoryless assumption does is provide a lower bound on capacity [95].

Because we are using M-FSK, the alphabet consists of M different symbols, which we denote as $\{0, 1, 2, ..., M-1\}$. Let $X = i$ correspond to transmitting on frequency f_i with power P_T^x for T_{ON} seconds. As such, the transmitted signal is given by:

$$x_i(t) = \sqrt{2P_T^x}\cos(2\pi f_i t + \theta_i^x) \qquad 0 \leq t \leq T_{ON} \tag{5.2}$$

At the destination node, the signals $x(t)$ and $J(t)$ can be received consecutively, with a gap, or with overlap, as shown in Figure 5.6. The third case, when there is overlap, is unlikely to occur if proper T_{ON} and T_{OFF} values are found. Thus, analysis will be limited to the first

Table 5.1: Summary of symbols used in this chapter.

Symbol	Description
$x(t)$	signal transmitted by the source node
$J(t)$	signal transmitted by the jammer
$r(t)$	received signal at the destination node
$n(t)$	noise component of the received signal at the destination node
M	number of frequencies used in FSK, i.e., order of the modulation scheme
T_{ON}	duration of the pulse transmitted by the source node, for each symbol
T_{OFF}	time between pulses, when the source node is not transmitting
T_{sense}	sense interval of the jammer, based on how often it makes a detection decision
T_{idle}	the jammer's idle interval, which is time it is not sensing or jamming
T_{jam}	duration of the pulse transmitted by the jammer, for each symbol
T_{gap}	time between when signal $x(t)$ stops, and $J(t)$ starts, for each symbol
K_{min}	fraction of the sense interval which must overlap with $x(t)$ for reliable detection
α	attenuation associated with LOS component of the source node's signal $x(t)$
R	random variable: attenuation associated with scattering component of $x(t)$
β	attenuation associated with the jamming signal $J(t)$
E_b	energy per information bit, which takes into account overhead associated with r
N_0	noise spectral density of the two-sided white Gaussian noise process $n(t)$
f_i, i, $x_i(t)$	current frequency, current frequency index, $x(t)$ at that frequency
P_T^x	source node's transmitted power (while it's transmitting, not averaged)
P_T^J	jammer's transmitted power (while it's transmitting, not averaged)
θ_i^x, θ_i^J, ϕ	random variables: random and independent phase shifts on the interval $[0, 2\pi]$
σ	parameter of the Rayleigh distribution; R has mean square equal to $2\sigma^2$
E_s^x, E_s^J, E_s	received energy per symbol at the destination node, for $x(t)$, $J(t)$, and combined
JSR_{avg}	average jammer-to-signal ratio (taking into account many symbols)
SNR	signal-to-noise ratio, not taking into account the jamming signal
Y_k	random variable: decision metric of the energy detector
$Y_{k,i}$, $Y_{k,q}$	random variables: real and imaginary portions of Y_k
C	theoretical channel capacity for reliable communications, i.e., the Shannon limit
R_0	cutoff rate, the channel capacity for reliable and practical communications
r	code rate associated with the forward error correction, on the interval $[0, 1]$

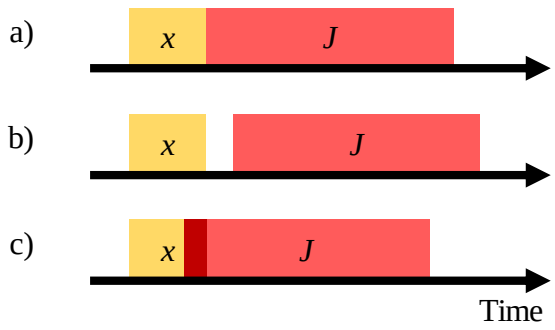

Figure 5.6: At the destination node, the two signals can be received: (a) consecutively, (b) with a gap, or (c) with overlap.

two cases. Let T_{gap} denote the time between when signal $x(t)$ ends and when signal $J(t)$ begins. Case (a) in Figure 5.6 corresponds to $T_{gap} = 0$.

Assuming the jammer is triggered by $x(t)$ and immediately transmits "noise" on subchannel f_i, the resulting jamming signal is given by:

$$J_i(t) = \sqrt{2P_T^J}\cos(2\pi f_i t + \theta_i^J) \qquad (T_{ON} + T_{gap}) \leq t \leq (T_{ON} + T_{gap} + T_{jam}) \tag{5.3}$$

where P_T^J is the jammer's transmit power. Note that this signal representation closely resembles what the OFDM-based reactive jammer, discussed in Chapter 9, would transmit (i.e., a sinusoid). However, the jammer may also transmit a signal that resembles Gaussian noise spanning the bandwidth of the subchannel, as discussed in Chapter 4. We will leave this latter case for future analysis.

The received signal at the destination node is given by:

$$r(t) = \alpha x_i(t) + R\sqrt{2P_T^x}\cos(2\pi f_i t + \theta_i^x + \phi) + \beta J_i(t) + n(t) \tag{5.4}$$

where α is the strength of the LOS component, ϕ is uniformly distributed between $[0, 2\pi]$, and R is the Rayleigh random variable with distribution:

$$p_R(r) = \begin{cases} \frac{r}{\sigma^2}e^{-r^2/2\sigma^2} & : \quad r \geq 0 \\ 0 & : \quad r < 0 \end{cases}$$

The noise term $n(t)$ is a white Gaussian noise process with two-sided spectral density $N_0/2$ (we will later discuss how to find N_0 for this unconventional system). Note that we can evaluate a Rayleigh and AWGN channel between the source and destination nodes by setting $\alpha^2 = 0$ and $R = 0$ respectively. All other configurations represent a Rician fading channel. Likewise, we can remove the jamming component (e.g., to investigate the baseline scenario to determine antifragile gain) by setting $\beta^2 = 0$.

The average received energy per symbol from the source node only is given by:

$$E_s^x = (\alpha^2 + 2\sigma^2)P_T^x T_{ON} \tag{5.5}$$

while the average received energy per symbol from the jammer is:

$$E_s^J = \beta^2 P_T^J T_{jam} \tag{5.6}$$

making the total received energy per symbol:

$$E_s = (\alpha^2 + 2\sigma^2)P_T^x T_{ON} + \beta^2 P_T^J T_{jam} \tag{5.7}$$

The average jammer-to-signal ratio (JSR) measured at the destination node is expressed as

$$\text{JSR}_{\text{avg}} = \frac{E_s^J/\text{symbol}}{E_s^x/\text{symbol}} = \frac{E_s^J}{E_s^x} = \frac{\beta^2 P_T^J T_{jam}}{(\alpha^2 + 2\sigma^2)P_T^x T_{ON}} \tag{5.8}$$

The notion of JSR is not that clear in this framework, because many scenarios, such in the one depicted in Fig 5.1, involve the source signal and jamming signal never being received at the same time. Typically, power ratios such as these are not used when signals are not received simultaneously and co-channel. However, we believe the JSR convention lends itself to more easily interpreted results (as opposed to using only signal-to-noise ratio (SNR) and jammer-to-noise ratio (JNR), for example).

While not related to the remainder of this analysis, we note that the value $(E_s^x + E_s^J)/E_s^x$ represents the antifragile gain in terms of energy out versus energy in, when taking into account the channels (this is further explored in the analysis of results).

The FSK receiver evaluates an M-dimensional array of decision metrics, as shown in Figure 5.7, with metric for frequency f_i denoted as Y_i. Because we are using an energy detector, each decision metric Y_i is given by:

$$Y_i = Y_{i,c}^2 + Y_{i,s}^2 \tag{5.9}$$

$$Y_{i,c} = \frac{2}{\sqrt{N_0(T_{ON} + T_{gap} + T_{jam})}} \int_0^{T_{ON}+T_{gap}+T_{jam}} r(t)\cos(2\pi f_i t)dt \tag{5.10}$$

$$Y_{i,s} = \frac{2}{\sqrt{N_0(T_{ON} + T_{gap} + T_{jam})}} \int_0^{T_{ON}+T_{gap}+T_{jam}} r(t)\sin(2\pi f_i t)dt \tag{5.11}$$

For hard decision demodulating, the highest Y_i is simply chosen. For soft decision, the values are passed into the channel decoding module, conserving the information. Note that the term outside the integral can be left out if the noise power on each carrier is the same, or if the noise powers are unknown.

To determine theoretical channel capacity and cutoff rate, we must find the probability density functions of Y_k conditioned on $X = i$, for when $k = i$ and $k \neq i$. The following

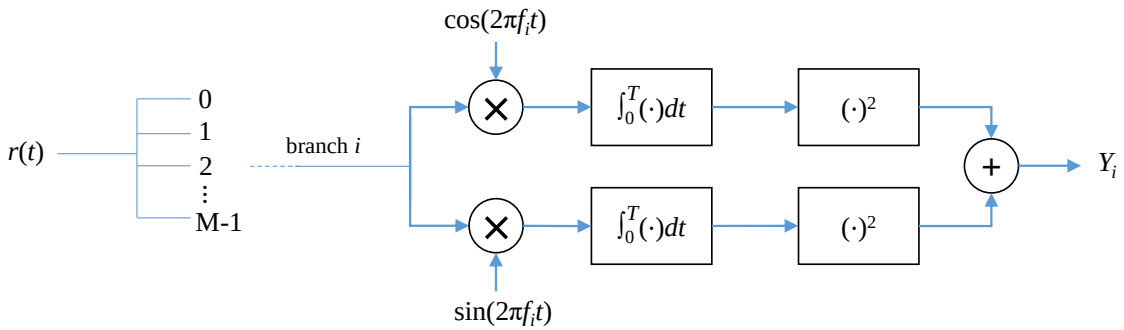

Figure 5.7: M-FSK receiver, zooming into a single branch.

analysis assumes there is no adjacent channel interference. Following the analysis in [95], but taking into account the extra signal $J(t)$, we can show that these random variables are Gaussian with means and variance:

$$E[Y_{k,c}|X=i] = \begin{cases} \alpha\sqrt{2P_T^x T_{ON}/N_0}\cos(\theta_k^x) + \beta\sqrt{2P_T^J T_{jam}/N_0}\cos(\theta_k^J) & k=i \\ 0 & k\neq i \end{cases}$$

$$E[Y_{k,s}|X=i] = \begin{cases} \alpha\sqrt{2P_T^x T_{ON}/N_0}\sin(\theta_k^x) + \beta\sqrt{2P_T^J T_{jam}/N_0}\sin(\theta_k^J) & k=i \\ 0 & k\neq i \end{cases}$$

$$\mathrm{VAR}[Y_{k,c}|X=i] = \mathrm{VAR}[Y_{k,s}|X=i] = \begin{cases} 2\sigma^2 P_T^x T_{ON}/N_0 + 1 & k=i \\ 1 & k\neq i \end{cases}$$

Note that the jamming signal does not influence the variance, because we are assuming the attenuation β^2 is constant.

When $X \neq i$ (i.e., only noise is received), Y_k is the sum of the squares of two standard normal random variables, so Y_k takes on the chi-squared distribution with $k=2$:

$$p(y_k|X\neq i) = \frac{1}{2}\exp\left[-\frac{1}{2}y_k\right] \qquad k\neq i \tag{5.12}$$

when $y_k \geq 0$, and zero otherwise.

When $X=i$, the signals are received, and the probability density function of Y_k is expressed as:

$$p(y_k|X=i) = \frac{1}{2+\frac{4\sigma^2 P_T^x T_{ON}}{N_0}}\exp\left[-\frac{1}{2}\frac{y_k N_0 + 2\alpha^2 P_T^x T_{ON} - 2\beta^2 P_T^J T_{jam}}{N_0 + 2\sigma^2 P_T^x T_{ON}}\right] \cdot I_0\left[\frac{\sqrt{2\alpha^2 P_T^x T_{ON}(N_0 y_k - 2\beta^2 P_T^J T_{jam})}}{N_0 + 2\sigma^2 P_T^x T_{ON}}\right] \qquad k=i \tag{5.13}$$

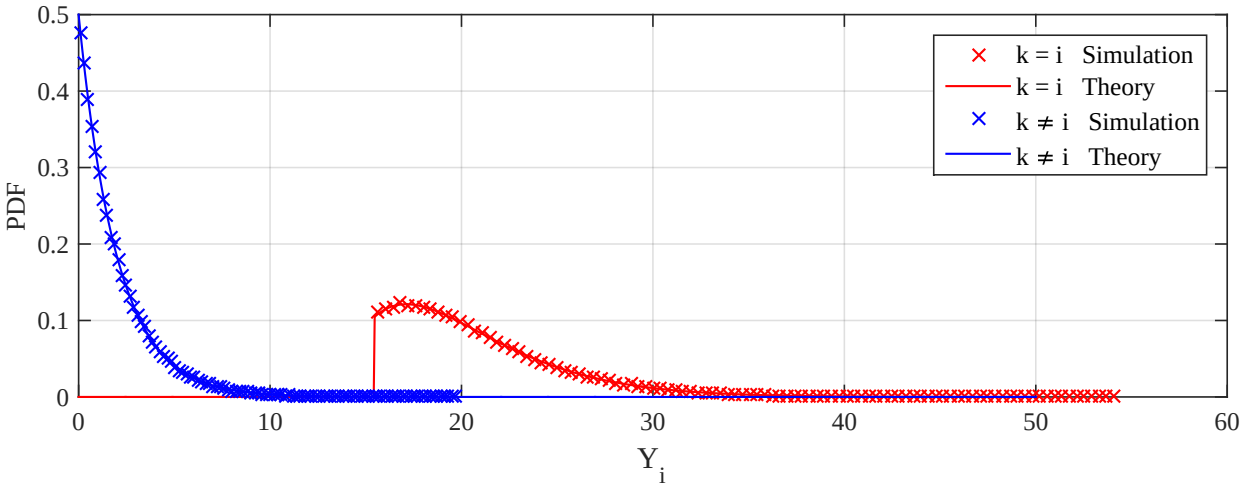

Figure 5.8: Validation of Equation 5.13 using an arbitrary set of parameters.

when $y_k \geq 0$, and zero otherwise. I_0 is the modified Bessel function of the zero'th order. Equation 5.13 is similar to the one derived in [95], except the jamming signal adds additional energy and shifts Y_i accordingly. Note that $\cos(\cdot)$ and $\sin(\cdot)$ disappeared, because θ_k^x and θ_k^J are uniformly distributed between zero and 2π. This is more easily understood when considering how a phase shift does not change the received energy (recall that this is what Y_k represents).

In order to validate Equation 5.13, we have simulated a baseband version of the scenario (depicted in Figure 5.5). Figure 5.8 shows the probability density function of Y_i under an arbitrary set of parameter values, using both simulation and theory, and for both $k = i$ and $k \neq i$. The specific values chosen for simulation were: $\alpha = 0.92$, $\beta = 0.9$, $\sigma = 0.158$, $P_T^x = 5.2323$, $P_T^J = 10.323$, $N_0 = 0.89$, $T_{ON} = 0.324$, and $T_{jam} = 0.823$.

5.4 Channel Capacity Using Cutoff Rate

We will now evaluate the channel capacity of the proposed waveform, under the same memoryless Rician channel (between the source and destination node), non-coherent detection, and optimal decoding (equivalent to maximum likelihood decoding). Instead of using the Shannon limit, which provides the maximum rate at which information can be reliable communicated, we make use of the cutoff rate. The cutoff rate provides an expression for channel capacity when taking into account the practical limits of convolutional codes.

The cutoff rate is defined as the rate at which the number of steps per sequentially-decoded digit becomes infinite [96]. In other words, when information is sent at a rate above the cutoff rate, then the average computational complexity of the decoding algorithm, when using convolutional codes, becomes unbounded. This applies to convolutional codes, and although

they form the building blocks for modern forward error correction schemes, the cutoff rate does not directly apply to schemes like turbo codes and LDPC codes. In fact, decades ago, when the cutoff rate was widely used in analysis, it was thought that there were no codes that could achieve higher than the cutoff rate. We now know that turbo and LDPC codes have the ability to perform between the cutoff rate and the Shannon limit (the irrefutably maximum rate). As such, the term cutoff rate (and M-FSK in general) rarely appears in modern literature. In this analysis, instead of using cutoff rate purely as a way to produce more practical results (which was the case during the 80s and 90s), we take advantage of its ability to simplify the mathematical expressions when calculating channel capacity.

Under M orthogonal signals and optimal decoding, the cutoff rate, denoted as R_0, is given by [97]

$$R_0 = 1 - \log_M \left[1 + (M-1) \left(\int_{-\infty}^{\infty} \sqrt{p(y|X=0)p(y|X=1)} dy \right)^2 \right] \tag{5.14}$$

By substituting in Equation 5.13 we find the cutoff rate for the proposed waveform under optimal decoding and soft detection:

$$R_0 = 1 - \log_M \left[1 + \frac{M-1}{4 + \frac{8\sigma^2 P_T^x T_{ON}}{N_0}} \left\{ \int_0^{\infty} \exp\left(-\frac{1}{4} \frac{y_k N_0 + 2\alpha^2 P_T^x T_{ON} - 2\beta^2 P_T^J T_{jam}}{N_0 + 2\sigma^2 P_T^x T_{ON}} - \frac{1}{4} y_k \right) \right. \right.$$
$$\left. \left. \cdot \sqrt{I_0 \left(\frac{\sqrt{2\alpha^2 P_T^x T_{ON} (N_0 y_k - 2\beta^2 P_T^J T_{jam})}}{N_0 + 2\sigma^2 P_T^x T_{ON}} \right)} dy \right\}^2 \right] \tag{5.15}$$

Equation 5.15 is similar to the cutoff rate derived in [95], except the jamming component adds some additional terms, and we have inserted the full expression for received power associated with each component (due to the pulsed nature of the signals). Because Equation 5.15 can be evaluated relatively easily, it lets us gain quick insight into the impact of each parameter, without having to run an entire simulation of the waveform for each set of parameters.

In this case, the cutoff rate is expressed in terms of information bits per symbol (i.e. bits per complex dimension) instead of bits/s/Hz. This is often done to avoid the dependence on an arbitrary definition of bandwidth. In our case, it is because the bandwidth occupied by our FSK signal depends on the subchannel configuration of the jammer, not the transmitted signal itself. In fact, the occupied bandwidth could be on the order of 10 dB or higher compared to the bandwidth of $x(t)$. This is not necessarily a problem, as long as the source and destination nodes have an RF chain that can support the bandwidth. In order to keep our analysis agnostic of bandwidth, we have only discussed energy per bit or symbol, noise spectral density N_0 (in units of joules), and non-coherent detection using an energy detector. In Section 5.6.2 we discuss the power spectral density of $x(t)$ and how it must be narrower than the jammer's subchannel width.

5.5 Numerical Results

In this section we present numerical results, showing the E_b/N_0, SNR, and JSR necessary for reliable (arbitrarily small error rate) and practical (taking into account the cutoff rate) communications. We also provide example throughput figures for scenarios based on exploiting a reactive jammer designed to jam IEEE 802.11.

5.5.1 Bits per Symbol

To reduce the parameter space for the sake of visualizing results, and focus on a given SNR and JSR, we investigate the scenario in which the source node and jammer have the same transmit power, the LOS and Rayleigh components have the same received power, and the channel attenuations are the same, denoted as $P_T = P_T^x = P_T^J$ and $\alpha^2 = 2\sigma^2 = \beta^2/2$. Under these assumptions, the average JSR is simply $\text{JSR}_{\text{avg}} = T_{jam}/T_{ON}$, allowing the absolute time durations to be removed from the equations. These assumptions allow us to rewrite Equation 5.15 in terms of SNR and JSR_{avg}:

$$R_0 = 1 - \log_M \left[1 + \frac{M-1}{4 + \frac{2\cdot\text{SNR}}{1+\text{JSR}_{\text{avg}}}} \left\{ \int_0^\infty \exp\left(-\frac{1}{4} y \frac{\frac{1+\text{JSR}_{\text{avg}}}{\text{SNR}} + 1 - 2\cdot\text{JSR}_{\text{avg}}}{\frac{1+\text{JSR}_{\text{avg}}}{\text{SNR}} + \frac{1}{2}} - \frac{1}{4} y \right) \cdot \sqrt{I_0\left(\frac{\sqrt{y\frac{1+\text{JSR}_{\text{avg}}}{\text{SNR}} - 2\cdot\text{JSR}_{\text{avg}}}}{\frac{1+\text{JSR}_{\text{avg}}}{\text{SNR}} + \frac{1}{2}} \right)} dy \right\}^2 \right] \quad (5.16)$$

Numerical results are shown in Figures 5.9 and 5.10 when SNR is set to a constant 3 dB and 10 dB respectively, and M is varied. The y-axis indicates the number of information bits per symbol, which can be converted to throughput by multiplying by the symbol rate. The choice to display bits per symbol was made because the symbol rate is heavily based on the response time of the jammer. The performance without a jammer present, i.e., using FSK at the same rate except without jammer piggybacking, is simply the level at which each curve flattens out on the left-hand side. This *non-jammed case* is worth noting for the sake of visualizing the antifragile gain in terms of energy savings. However, it does not make a lot of sense to use the proposed approach without the jammer piggybacking aspect to it, as the symbol rate would be impractically low.

It is clear that a significant antifragile gain only occurs when M is fairly large, and is most apparent when the SNR is low. This makes sense, because if the SNR were high, then the source node does not need any "help" from the jammer in order to transmit reliably at a medium-to-high data rate, for a given amount of energy. Instead of simply using more energy,

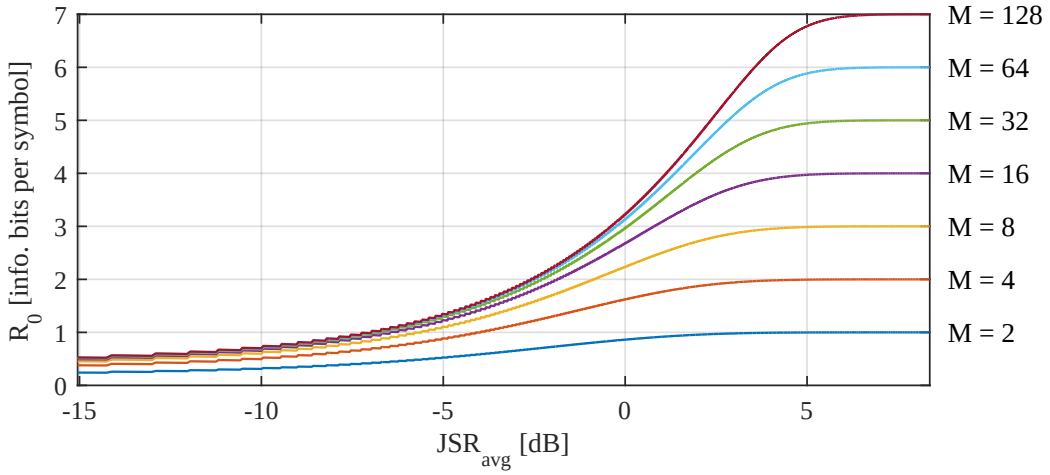

Figure 5.9: Practical channel capacity (based on the cutoff rate) of the proposed FSK antifragile waveform when SNR $= 3$ dB, under a Rician source-destination channel and AWGN jammer-destination channel. In this example, $\alpha^2 = 2\sigma^2$, meaning the LOS and scattering components have equal received power. The number of subchannels being piggybacked on, M, is varied. For each value of M, the *non-jammed* case is simply when $\text{JSR}_{\text{avg}} = -\infty$, or roughly where each curve hits the y-axis.

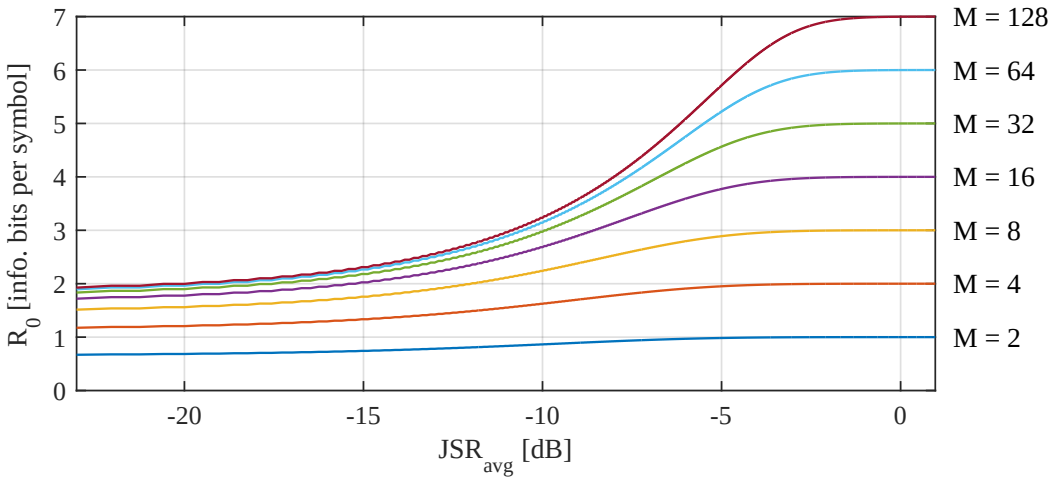

Figure 5.10: Practical channel capacity (based on the cutoff rate) of the proposed FSK antifragile waveform when SNR $= 10$ dB.

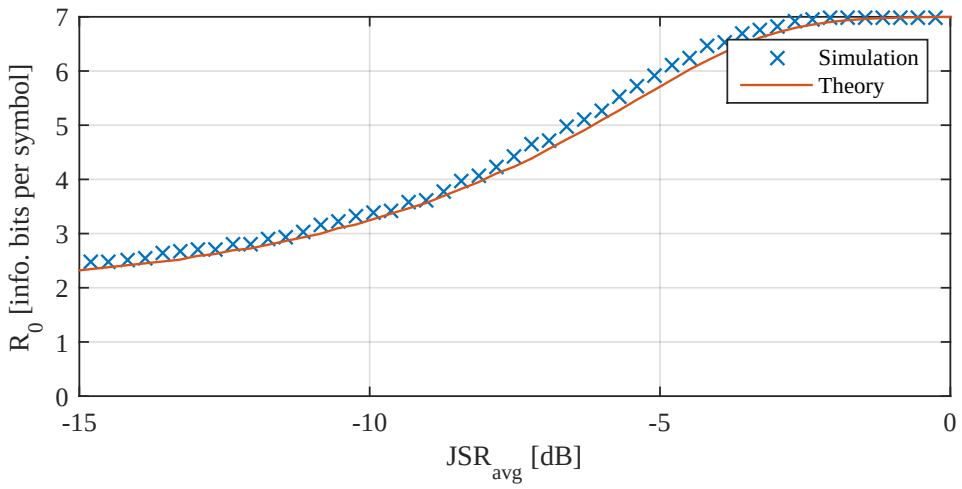

Figure 5.11: Validation of Equation 5.15 using the same parameters as discussed above, and with $M = 128$.

and increasing the effective SNR, we are taking advantage of the enemy jammer to help us communicate, and hence the strategy takes on an antifragile property. In both figures, the throughput for each curve levels out at $\log_2(M)$.

For the sake of validation, we have taken the same simulation framework used to validate Equation 5.13 (shown in Figure 5.8) and found the cutoff rate when $M = 128$. The theoretical and simulated curved are shown in Figure 5.11. The slight offset is due to the nature of numerical integration; the offset can be reduced by using smaller bins and increasing the number of trials when finding the histogram of Y_i.

5.5.2 Throughput

The results presented thus far show what SNR and JSR levels are needed to achieve a certain number of bits per symbol. However, as discussed earlier, the number of FSK frequencies (which determines the achievable bits per symbol) is based on the jammer's subchannel configuration. Meanwhile, throughput is heavily based on the response time on the jammer, and how well T_{ON} and T_{OFF} are tuned. Because of these factors, it makes sense to investigate the channel throughput under a variety of scenarios, while assuming that the SNR and JSR are "high enough" for FSK under a certain code rate r. We must also take into account the detection engine within the jammer, which will not always trigger if T_{ON} is too small.

When there are no block errors, throughput performance in terms of bits per second is simply $\frac{1}{T_{ON}+T_{OFF}} r \log_2 M$. In the next subsection we discuss the nature behind M. The question now is: what values of T_{ON} and T_{OFF} are likely achievable in a practical scenario? Let us consider a reactive jammer that intends to jam Wi-Fi type packets. IEEE 802.11b uses a

Table 5.2: Potential Throughput in Bits per Second Under Various Assumptions.

Packet Length	Overlap	T_{sense}	T_{idle}	T_{jam}	M	**Throughput**
96 μs	75%	10 μs	10 μs	100 μs	64	40 Kbps
96 μs	50%	34 μs	10 μs	100 μs	64	35 Kbps
192 μs	75%	34 μs	20 μs	200 μs	64	20 Kbps
96 μs	75%	10 μs	10 μs	100 μs	256	53 Kbps
96 μs	50%	34 μs	10 μs	100 μs	256	46 Kbps
192 μs	75%	34 μs	20 μs	200 μs	256	27 Kbps

packet length between 96 and 192 μs for control packets (with no data payload). In 802.11n and 802.11ac, the maximum packet length is roughly 5.5 ms [98]. So if the jammer intends to jam all packets, and overlap with a significant portion of each packet (let's say 75%) after taking into account propagation delay (let's say 14 μs, equating to 4 km), then it should be able to detect the signal within the first 10 μs or so. The 802.11n symbol duration is 4 μs in most configurations [98], so even though this only equates to 2.5 symbols, these are OFDM symbols so they are much longer than the equivalent single-carrier signal's symbol length.

As an example, if we have a reactive jammer with a $T_{sense} = 10\ \mu s$, $T_{idle} = 10\ \mu s$, and $T_{jam} = 100\ \mu s$, then T_{ON} must be at least 16 μs and T_{OFF} must be at least 100 μs (recall Equation 5.1). We have assumed $K_{min} = 0.3$, as before. There is no need for the jammer to take into account propagation delay, because it is already taken into account when determining how quick the jammer must detect the signal. In this example, if we assume that $M = 64$ and a code rate $r = 0.75$ is used, then the throughput would be roughly 40 Kbps. This example reflects a fairly quick reactive jammer, aiming to jam Wi-Fi control packets. A jammer only intending to jam the much longer data packets would require a lower reaction time, and thus lead to an antifragile waveform with lower throughput. Additional examples are provided in Table 5.2, all assuming a code rate of 0.75 and propagation delay of 14 μs. Note that the sense time is based on the packet length, overlap desired, and propagation delay. The jam durations chosen are based on the packet length. The value M is chosen arbitrarily, and throughput is calculated based on the three T values and M. Further information regarding the timing constrains associated with sensing-based reactive jamming can be found in [53] and [99], both of which discuss implementations of reactive jammers.

To gain further insight, we have simulated the jammer's detection engine using $K_{min} = 0.3$ and $T_{sense} = 10\ \mu s$. The jamming duration T_{jam} is varied. We then calculated throughput assuming a code rate of 0.75 and 256 subchannels. It is assumed that both SNR and JSR are high enough for reliable FSK at this code rate. The parameter T_{ON} is adjusted, to show the trade-offs between it being too large and too small. Results are shown in Figure 5.12. The throughput hits zero when T_{ON} is simply too short to trigger the jammer's detection engine. As T_{ON} gets larger, the symbol period gets longer and throughput is reduced. Note that spectral efficiency is not shown because the bandwidth occupied by the FSK signal depends

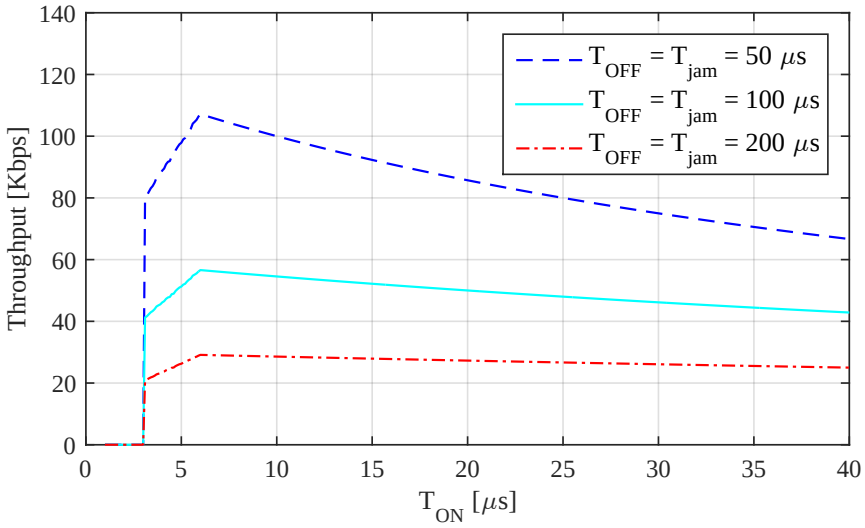

Figure 5.12: Example throughput assuming $K_{min} = 0.3$, $T_{sense} = 10\ \mu s$, $r = 0.75$, and $M = 256$. The throughput hits zero when T_{ON} is simply too short to trigger the jammer's detection engine. As T_{ON} gets larger, the symbol period gets longer and throughput is reduced.

on the jammer's subchannel configuration, and could vary greatly (this is further explored in Section 5.6.2).

5.6 Additional Design Guidelines

In this section we discuss additional design guidelines that should be considered when designing and implementing the described antifragile waveform.

5.6.1 Number of Parallel Channels and the Probing Signal

The number of parallel channels is based on the reactive jammer's channelizer or FFT configuration, which can be parameterized as how many subchannels it uses and the width of each subchannel. It is also based on the ability to achieve orthogonality between subchannels, which is discussed in Section 5.6.2. In order to maximize the number of frequencies used, M, the source and destination nodes must probe the reactive jammer to learn its subchannel configuration. This includes the subchannel spacing, the center frequency of each subchannel, and the frequency range that the jammer covers. An example probing signal is a very slow frequency sweep of an unmodulated carrier, with an occasional pause or even time reversal. The sweep would allow the destination node to detect the minimum and maximum frequencies covered by the jammer, as well as the subchannel spacing. The pause would

give the destination node time to observe only a single subchannel of the jammer being active, which it could use to determine the center frequency of that specific subchannel (and thus every subchannel, if they are equally spaced). The time reversal would give the destination node insight about the ending-lag of the jammer, allowing for a well-tuned T_{OFF}. The destination node can then share this information with the source node, and together they can decide on which subchannels to target to create orthogonal FSK signals. If the jammer is OFDM-based, as discussed in Chapter 9, then the radios may not want to use every subchannel, as OFDM does not provide orthogonality between subcarriers when it is non-coherently detected. This would result in significant energy smearing between adjacent subchannels, making FSK detection difficult in low SNR conditions.

5.6.2 Signal and Noise Bandwidth

As explained in Section 2.3, the noise bandwidth is an important factor in any link budget. The received noise power, P_{noise}, is usually expressed in terms of kTB, where k is the Boltzmann constant, $k \approx 1.38e23\ [J/K]$, T is the noise temperature in Kelvin, and B is the noise bandwidth. Note that noise spectral density is equal to $N_0 = kT$, which is agnostic of bandwidth.

The noise bandwidth is usually equal to the most selective or narrow filter within the receive chain, which may occur at intermediate frequency (IF), or in the digital baseband in the case of a software-defined radio (SDR). Furthermore, the bandwidth of this selective filter is usually based on the bandwidth of the signal being received, so that as much noise and spurious signals can be rejected as possible. Thus, what we are really concerned with finding is the signal bandwidth. Note that there are two different signals being received: $x(t)$ and $J(t)$. For now we will focus on $x(t)$, as we can only speculate about the bandwidth of $J(t)$.

In terms of signal bandwidth and pulse shaping, the Gaussian filter is commonly used in FSK and minimum shift keying (MSK) modulation, as it does not involve zero crossing points. Given a rolloff factor α, the impulse response of the Gaussian filter is given by [100]

$$h(t) = \frac{\sqrt{\pi}}{\alpha} \exp\left(-\frac{\pi^2}{\alpha^2} t^2\right) \tag{5.17}$$

and the 3 dB bandwidth is equal to $B = \frac{1}{T_{sym}}(\alpha + 1)$. The overall symbol period T_{sym} is given by $T_{ON} + T_{gap} + T_{jam}$ or $T_{ON} + T_{OFF}$, although it is worth noting that there are additional periodicities based on the duration of T_{ON}, T_{gap}, and T_{jam}, which are evident as spikes in the spectral density. Using the example in the previous subsection, let us say we have a $T_{ON} = 30\ \mu s$, $T_{gap} = 5\ \mu s$, and $T_{jam} = 100\ \mu s$. Figure 5.13 shows the power spectral density of such a signal. Note that this is only the power spectral density of one FSK subchannel, and thus the subchannels must be spaced by at least roughly 100 kHz in this example. If the jammer is already using a subchannel spacing this large, then M can simply be equal to the

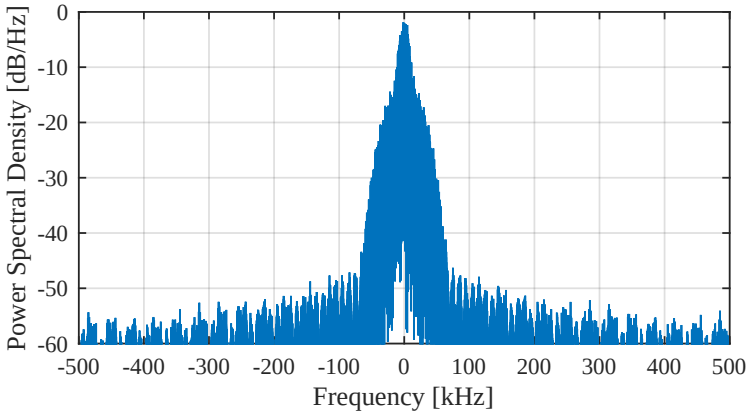

Figure 5.13: Example power spectral density at the received node, when $T_{ON} = 30\mu s$, $T_{gap} = 5\mu s$, and $T_{jam} = 100\mu s$.

number of the jammer's subchannels (as long as the jammer's signals don't have too much adjacent subchannel overlap). If the jammer's subchannel spacing is a fraction of this 100 kHz figure, subchannels must be left unused for the sake of creating guard bands between the orthogonal FSK signals. Alternatively, T_{OFF} can be increased, to increase the symbol period and ensure that each and every symbol transmitted by the source node gets relayed by the jammer.

Lastly, we note that a sinusoid may not be the best signal for triggering a jammer's detection engine (recall that FSK is essentially transmitting sinusoids). If the jammer is using a wide subchannel width, then it may be better to use a (relatively) wide bandwidth signal for $x(t)$, in an FSK manner. For example, each FSK symbol could use a $x(t)$ made up of a dozen quadrature phase-shift keying (QPSK) symbols. This would increase the bandwidth of $x(t)$, allowing it to better fit the jammer's subchannel configuration. Further exploration into this concept of synthetically increasing the bandwidth of $x(t)$ is left for future work.

5.6.3 Hybrid Approach

We remind the reader that the strategy proposed in this section should only be initiated once there is an indication that jamming is present. However, even under a reactive jamming attack, there may be situations in which the proposed strategy would cause poor performance (e.g., high SNR and low JSR scenarios). Given the radio's hardware, it may be possible to communicate using a hybrid approach, whereby the proposed antifragile waveform is combined with higher order symbols. Under this hybrid approach, the source and destination can communicate at a high data rate when the channel is favorable and there is only a low-to-moderate level of jamming. The mechanism for jammer piggybacking would still be included; the destination node would simply discard the jamming portion of $r(t)$ and use as much of $x(t)$ as is recoverable. Under severe reactive jamming and poor channel

conditions, the low data rate with piggybacking would be used, and the fact that the source node is transmitting higher order symbols would be ignored. For example, when using the hybrid approach, the source node could simply use M-PSK augmented FSK (instead of just a sinusoid transmitted on frequency f_i), so that when coherent detection is possible, a higher data rate is achieved. Note that this hybrid approach requires the usual overhead associated with coherent detection, such as channel estimation using pilots. In addition, this hybrid approach assumes that the source and destination nodes are able to do light reconfiguring of the waveform in the field, implying they are SDRs.

5.6.4 Counter Strategy

With every countermeasure, or exploitation in this case, comes a counter-countermeasure. In fact, anti-jamming as a discipline used to be called electronic counter-countermeasures (ECCM), as jamming attacks were considered electronic countermeasures (ECM). Using this old electronic warfare terminology, what we are interested in this subsection is the Electronic Counter-Counter-Countermeasure (ECCCM), or EC^{3}M™.

From the perspective of the jammer, the antifragile waveform proposed in this chapter can be easily countered by the occasional jamming of subchannels that are observed to be idle (remember, we are under the assumption that the jammer will jam all subchannels that are active). However, this would cause the jammer to use more power, and be easier to detect, because it would no longer be as hidden under the target's transmissions. Alternatively, the jammer could intentionally extend its ending-lag, which would reduce the achievable data rate. However, given the obscure nature of these antifragile techniques, it is unlikely that a modern jammer would be programmed with this type of EC^{3}M™ in mind.

5.7 Conclusion and Future Work

In this chapter we introduced a FSK-based jammer piggybacking waveform, to be used against a sense-and-transmit type reactive jammer. This waveform has been shown to have antifragile properties, and an antifragile gain is achievable under low SNR, high JSR, and high M scenarios. We derived the channel capacity and cutoff rate when using the FSK waveform, and provided numerical results via simulation to gain further insight. A variety of design guidelines were provided to assist in implementation of the proposed strategy.

Future work includes the implementation of such a strategy, to really put it to the test (note that this may only be worth doing once a sense-and-transmit type reactive jammer is acquired). In addition, it may be possible to use each subchannel as an independent on-off keying (OOK) stream, so that more bits can be sent through (M instead of $\log_2 M$ bits per symbol). Unlike FSK, this would require the destination node to keep track of the noise power for each subchannel so it can calculate the optimal OOK threshold.

Chapter 6

Network Level Antifragile Gain

6.1 Introduction

In this chapter we investigate two antifragile strategies that function on the network layer (NET), as opposed to the physical layer (PHY) and media access control layer (MAC) which was the focus of Chapters 4 and 5. Specifically, we seek to exploit an enemy jammer to provide a gain (relative to a non-jammed case) in network connectivity, and in some cases covertness. This chapter is split into two sections, each on one of the network-level antifragile strategies.

6.2 Coarse Timing Synchronization and Rendezvous using Enemy Pulsed Jammer or Radar

We will now take a brief departure from jammer piggybacking, to investigate an antifragile strategy that does not require a reactive-type jammer to be present.

6.2.1 Motivation

In a distributed ad-hoc wireless network, there is no simple solution for coordinating transmissions among nodes. This is why contention-based MAC protocols, ones in which nodes transmit at any time using a first come-first served basis, are so common. IEEE 802.11, for example, uses carrier sense multiple access with collision avoidance (CSMA/CA), which is a contention-based MAC protocol where the nodes sense the wireless channel to check if it's idle before transmitting. However, even with this check in place, collisions are still possible, and common when there is high network load. The alternative to a contention-based MAC protocol is to schedule node's transmissions, which is the method used in all cellular technologies. When a centralized base station is used to arbitrate medium access, the base station can transmit a synchronization sequence which defines the time and/or frequency slots. Unfortunately a centralized architecture has many issues, including: requirement of nodes to be within range of base station, overhead associated with state information being shared with base station, and reliability issues associated with a central point of failure. In a purely distributed wireless network, scheduling is not so easy, making it an important area of wireless communications research (for more information we refer the reader to [101] and [102]). In this section we will discuss a strategy that involves a distributed wireless networking using a signal of opportunity as a solution for scheduling and synchronization.

6.2.2 Problem Formulation of Antifragile Strategy

Continuing with the theme of antifragility, in this section we discuss the process of using an enemy signal as the source for a synchronization signal that nodes in a given area can use to coordinate transmissions. Our analysis is within the context of a time division multiple access (TDMA) system, one in which nodes share a channel by taking turns transmitting. This includes TDMA-OFDM systems, although not orthogonal frequency-division multiple access (OFDMA) systems. As stated in [103], "The main problem with TDMA lies in achieving synchronization between the different stations." In the proposed strategy, we aim to overcome this problem by exploiting an enemy signal. Time slots are divided among nodes, although in this analysis we are not concerned with how this division takes place (in terms of how resources are split between nodes and assigned). Each falling-edge of the signal of opportunity defines when the next time-slot begins. The main requirement of this signal of interest is that it is pulsed, in a predictable manner (i.e., with a fairly constant duty-cycle and minimal jitter). Example signals include enemy radars, and pulsed enemy jammers. It has been shown that using a pulsed signal is optimal jamming waveform against most modern digital modulation schemes [104]. Thus, it is worthwhile to consider a strategy that exploits this type of jammer. The pulse rate is typically on the order of several symbols, similar to the distribution of pilot symbols. In an automatic-gain control type jamming attack, it could be on the order of hundreds of symbols. This strategy can be extended to using a friendly signal of opportunity, although it would no longer hold the antifragile property. It is envisioned that this synchronization strategy is used with wireless nodes that transmit at a much lower power compared to that of the enemy radar or pulsed jammer, such that the wireless nodes may not all see each other, but they all can detect the signal of opportunity.

To understand the benefits of this strategy, it helps to first consider an alternative approach, where there is no signal of opportunity. The wireless nodes simply transmit for a fixed duration, and the time-slot boundaries are determined by the end of each node's transmission. One problem with this approach stems from how the first transmission occurs, within each string of transmissions. The first packet could be sent by the node that was turned on first, or has data to communicate first, and all subsequent packets would be based on that one. However, when there is a gap in communications, or the hidden node problem surfaces, time synchronization might be lost. By using an external signal of opportunity that can be detected by all nodes in the network, and that is assumed to always be on, this problem is avoided. The node which transmits first is simply the first one to discover the signal of opportunity

6.2.3 Strategy Variations

The antifragile strategy proposed in this section has two variations, one in which the nodes are co-channel with the enemy signal, and one in which they are offset in frequency from the signal. The latter approach results in a more complicated RF front-end, because the radios have

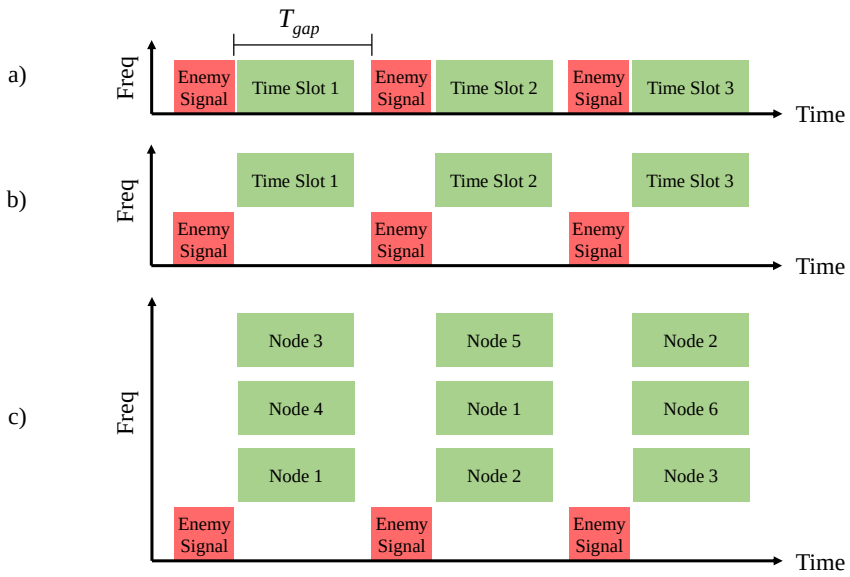

Figure 6.1: TDMA-based antifragile coarse timing synchronization strategy, where the enemy signal of opportunity and wireless nodes are a) co-channel, b) on separate channels, and c) on many separate channels.

to keep track of the enemy signal while receiving on a separate channel. However, being on separate channels allows this strategy to work with a TDMA system that involves more than one channel. In addition, even if the enemy signal is not predicable, it could still provide slot timing as long as the nodes support transmitting and receiving simultaneously (on different frequencies). If the radios are co-channel with the enemy signal, then they have to all refrain from transmitting while the enemy signal is active (unless interference cancellation tricks are used). Figure 6.1 illustrates these variations, as well as the TDMA-based synchronization strategy altogether. We define T_{gap} as the time between enemy transmissions, which for now will be treated as constant.

6.2.4 Side Benefits

In addition to providing a synchronization source, this strategy provides two more advantages. First, if the wireless nodes are co-channel with the enemy signal, then they are partially covered by the enemy signal. This could be of value if the enemy electronic warfare (EW) capabilities involve blacklisting their own radar and jammer signals and jamming everything else in a wide block of spectrum (e.g., 100 MHz to 10 GHz). This approach is known as "DC to Daylight Jamming", and it is only practical to do it with reactive jammers. Clearly this benefit only applies to scenarios in which the signal of opportunity has a bandwidth comparable to the friendly signals.

Secondly, the signal of opportunity can also provide a method for rendezvous between the

wireless nodes (i.e., a frequency in which to "meet" and begin communicating on). In the case that the radios are not co-channel, it is assumed that there is a preset fixed frequency offset that the nodes use, in order to rendezvous off the signal of interest. This second advantage is especially valuable for frequency-agile radios, who want to avoid transmitting a beacon signal. For example, let us say the wireless nodes are deployed in enemy territory, and thus nearly all signals in the area can be treated as enemy signals. The set of available channels is simply the set of observable enemy signals that fit the pulsed criterion. To avoid confusion associated with many signals of opportunity, the nodes can use a predefined priority scheme where they choose the signal that is closest to a certain frequency (previously agreed on), such as the center of their operational range.

In the remainder of this section we discuss the overhead associated with using this approach. We do not discuss throughput associated with this approach, because it depends too much on the specific waveform being used alongside this TDMA-based antifragile strategy.

6.2.5 Overhead Associated with Approach

One downside to this synchronization strategy is that there may be significant overhead associated with using a signal of interest. We define overhead as the fraction of time in which the nodes are not allowed to transmit, which lets us remain agnostic of the specific waveform being used. The following four factors contribute to this overhead:

1. Duty-cycle associated with enemy signal, in the case that the radios are co-channel with the enemy signal
2. Guard time associated with propagation delay, for the given radius of coverage
3. Lost time associated with delay in the RF chain
4. Guard time associated with stochastic nature of the signal's falling-edge, modeled as Gaussian

The first item is straightforward; in the co-channel case, nodes must refrain from overlapping with the enemy signal, to avoid being jammed by it. We denote the duty-cycle of the signal of opportunity as D_{cycl}.

To gain more insight about the second item, let us say we desire a coverage radius of d, as illustrated in Figure 6.2. A node in the same location as the enemy transmitter would see zero propagation delay, $T_{prop} = 0$. Meanwhile, a node located distance d from the enemy transmitter would see a delay $T_{prop} = \frac{d}{c}$ where c is the speed of light. However, we do not assume that the nodes know how far they are from the transmitter. When a node is assigned to a certain slot, it begins its transmission immediately after detecting the falling-edge of the signal of opportunity. However, it must only transmit for $T_{tx} = T_{gap} - \frac{d}{c}$, to ensure it does not have a collision with the signal of opportunity. It is for this reason that the nodes must

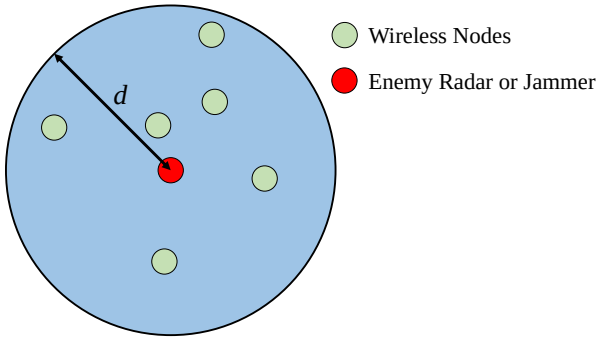

Figure 6.2: Given a coverage radius d around the enemy signal, we can determine the guard-interval to take into account propagation delay differences between nodes.

maintain an estimate of T_{gap}, and be preprogrammed with a maximum anticipated value of d, which can be proportional to the transmit power and frequency of the nodes. Thus, the guard time associated with propagation delay is $T_{guard} = \frac{d}{c}$. If the radios are not co-channel with the enemy signal, then they must avoid transmitting during the wrong time slot. To take into account the case of two nodes that are on opposite sides of the enemy transmitter, at a distance d, the guard interval should be set to $T_{guard} = \frac{2d}{c}$. Multipath should also be taken into account, if the delay spread is not insignificant compared to the propagation delay. The maximum anticipated delay spread can simply be added to T_{guard}.

Item three, delay associated with the RF chain, is a factor because the nodes only transmit after detecting the falling-edge of the signal of opportunity, so any delay between that detection and when a signal is actual transmitted is overhead. We denote this delay as T_{RF}.

Until this point, we have treated T_{gap} as constant. However, in a practical scenario, it may have some jitter associated with it. For example, a jammer may use a pulsed waveform to increase jamming efficiency, but incorporate randomness so that the target radios cannot as easily mitigate the jamming. For this analysis we model the jitter of the time associated with the rising-edge of the enemy signal as a Gaussian random variable, denoted T, with variance σ_s^2 and mean T_{gap}, as depicted in Figure 6.3 (the variance with respect to T_{gap} is exaggerated for the sake of demonstration). This timing is with reference to the falling-edge of the previous pulse's falling edge, so any jitter in the previous pulse's falling-edge, or randomness in the enemy signal's pulse duration, can be seen as jitter in the rising-edge of the next pulse. To take this randomness into account, the wireless nodes must cut their transmissions short, in the same way that they take into account propagation delay. In order to not overlap with the signal of opportunity c_s fraction of the time, where c_s is on the interval $[0.5, 1]$, we can find T_s by evaluating the expression $P(T < -T_s) = 1 - c_s$ or $F_T(-T_s) = 1 - c_s$ where $F(\cdot)$ is the cumulative distribution function of T. For a Gaussian distribution, T_s is found using the error function:

$$T_s = -F_T^{-1}(1 - c_s) \tag{6.1}$$

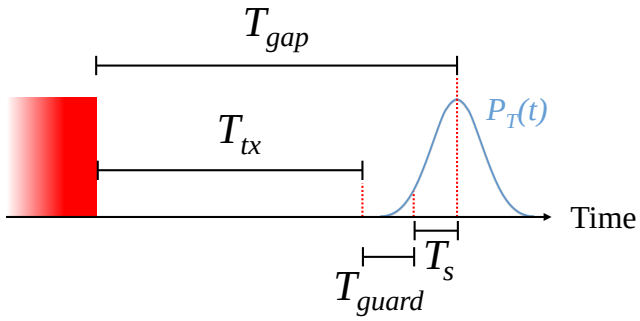

Figure 6.3: One time slot, showing how the rising-edge of the enemy signal is modeled as a Gaussian random variable. The wireless nodes must cut their transmissions short so that it has a low likelihood of overlapping with the enemy signal.

$$T_s = -\left[\mu + \sigma\sqrt{2}\,\mathrm{erf}^{-1}(1 - 2c_s)\right] \tag{6.2}$$

$$T_s = -\sqrt{2\sigma_s^2}\,\mathrm{erf}^{-1}(1 - 2c_s) \tag{6.3}$$

The total overhead is defined as the fraction of time the wireless nodes are not able to use the wireless medium. We can express this overhead as a function of D_{cycl}, σ_s^2, c_s, T_{gap}, T_{RF}, and d:

$$\text{overhead} = \frac{T_{gap}\frac{D_{cycl}}{1-D_{cycl}} + \frac{d}{c} + T_{RF} - \sqrt{2\sigma_s^2}\,\mathrm{erf}^{-1}(1 - 2c_s)}{\frac{T_{gap}}{1-D_{cycl}}} \tag{6.4}$$

for the co-channel case. A multiplier of two must be included in the $\frac{d}{c}$ term for the non-co-channel case. In Figure 6.4 we have plotted overhead while varying distance d and overlap parameter c_s, for a scenario involving a pulsed jammer with a 20% duty-cycle. The gap duration was set to 100 microseconds, and the standard deviation of the jammer's jitter was set to 5 microsecond. We assumed T_{RF} was negligible.

For a given waveform, peak throughput can be multiplied by (1 − overhead) to find the data rate when using this strategy.

6.2.6 Antifragile Slotted ALOHA

So far we have only discussed a TDMA-based approach. It may also be possible to use the exploited timing synchronization source as a way to improve a contention-based MAC protocol. In this subsection we investigate using the proposed antifragile strategy as part of a waveform using the Slotted ALOHA MAC protocol, and compare it with a CSMA-based waveform that operates in open spectrum.

ALOHA is the earliest known random access method, developed in 1970 at the University of Hawaii [103]. It works by nodes transmitting a frame whenever they have data to send,

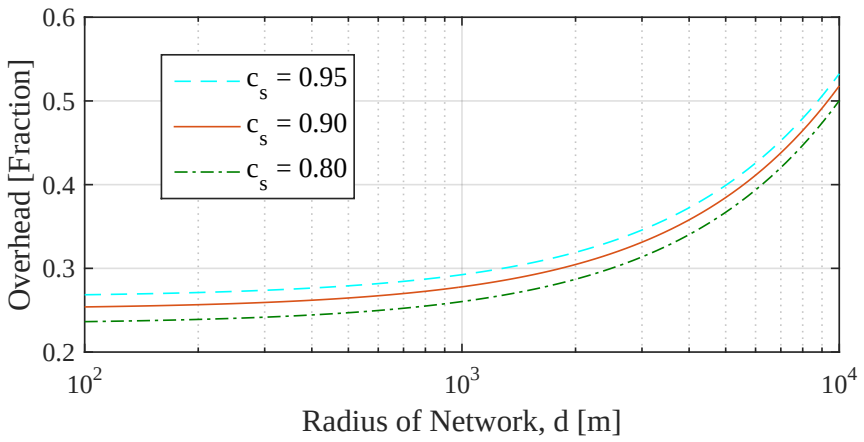

Figure 6.4: Plot of overhead while varying distance d and overlap parameter c_s, for a scenario involving a pulsed jammer with a 20% duty-cycle.

and if no acknowledgment of a frame is received after a predefined period of time, the node attempts a retransmission after waiting a random amount of time. This random back-off period is usually based on the number of failed attempts to send the frame. An improvement to ALOHA was created in the form of Slotted ALOHA, which involves dividing time into slots, implying that the nodes are globally synchronized in time (such as when using our antifragile time-sync strategy). By forcing each packet into a discrete slot of time, it reduces the probability of a collision, increasing performance compared to pure ALOHA [103]. ALOHA has been largely rendered obsolete due to protocols based on CSMA, which involves each node sensing the wireless medium before attempting a transmission, something that ALOHA does not require. However, there are still situations in which CSMA is undesirable, such as when there is a low-signal-to-noise ratio (SNR) at each node, and many interferers in the area, or when propagation delay is large. In these types of harsh wireless scenarios, the coarse timing synchronization provided by the proposed antifragile strategy can enable the gains associated with Slotted ALOHA (in addition to the gains discussed in Section 6.2.4).

6.2.7 Performance Comparison

It is well known that 802.11 is based on the distributed coordination function (DCF) technique, which is a form of CSMA/CA. Fundamentally, CSMA can be split up into persistent and non-persistent schemes. Under a persistent scheme, after there is a collision, a node attempts a retransmission immediately after the channel is idle again with a certain probability p, and waits a random back-off period with probability $1 - p$. Non-persistent CSMA, on the other hand, involves nodes always waiting a certain back-off period before attempting a retransmission. Modern communications protocols that use CSMA/CA use a scheme that is a combination of these two forms, called P-Persistent CSMA. The parameter p is typically somewhere between zero and one, and depends on the maximum number of nodes that can

be active simultaneously. In this subsection we compare our Antifragile Slotted ALOHA scheme with non-persistent and 1-persistent CSMA, representing the upper and lower bounds of CSMA performance. We also compare our scheme to Pure ALOHA.

In order to compare performance we assume a traffic model consisting of a number of users that collectively form a Poisson source with aggregate packet generation of λ packets per second on average. The inter-arrival times of each packet are independent. Let each packet take T seconds to transmit. We denote S as the average number of packets generated during a period of time T, which can be thought of as the channel throughput rate. Under our traffic model, $S = \lambda T$. At most, S can be 1 (e.g. if there was only one user with infinite data to send and no overhead was required), meaning S also acts as the utilization rate of our channel. Because failed transmissions result in retransmissions, we define G as the "offered" rate of packets. As such, $G \geq S$. The following analysis also assumes that the average retransmission delay is much greater than the transmission time T. The set of models and assumptions used are further discussed and validated in [105].

The throughput of pure ALOHA under the Poisson traffic model was first derived in [106], and is given by:

$$S_{ALOHA} = Ge^{-2G} \tag{6.5}$$

This is maximized when $G = 0.5$, causing a throughput of only 0.184. As discussed in the previous section, Slotted ALOHA achieves a performance increase over Pure ALOHA, because packets no longer only partially overlap. Under the Poisson traffic model, the throughput of Slotted ALOHA is [105]:

$$S_{S-ALOHA} = Ge^{-G} \tag{6.6}$$

which also has its maximum at $G = 0.5$, but instead reaches a throughput double that of ALOHA. Throughput of Antifragile Slotted ALOHA is simply (overhead $\cdot\, S_{S-ALOHA}$).

In order to determine the throughput of CSMA schemes, we must take into account propagation delay. For convenience, we introduce parameter a, representing the ratio of propagation delay to packet transmission duration. The authors of [105] were the first to derive the throughput of Non-Persistent CSMA (NP-CSMA) under the assumptions listed above, which is given by:

$$S_{NP-CSMA} = \frac{Ge^{-aG}}{G(1+2a) + e^{-aG}} \tag{6.7}$$

Similarly, the throughput for 1-Persistent CSMA (1P-CSMA) is given by [105]:

$$S_{1P-CSMA} = \frac{G\left[1 + G + aG(1 + G + aG/2)\right] e^{-G(1+2a)}}{G(1+2a) - (1 - e^{-aG}) + (1 + aG)e^{-G(1+a)}} \tag{6.8}$$

While Non-Persistent and 1-Persistent represent the two extremes of CSMA, we also include Slotted 1-Persistent CSMA (S1P-CSMA), which is simply 1-Persistent CSMA where nodes are synchronized with each other, and only transmit packets on a time-slot basis. It makes sense to include this variation in the comparison with the Antifragile Slotted ALOHA scheme,

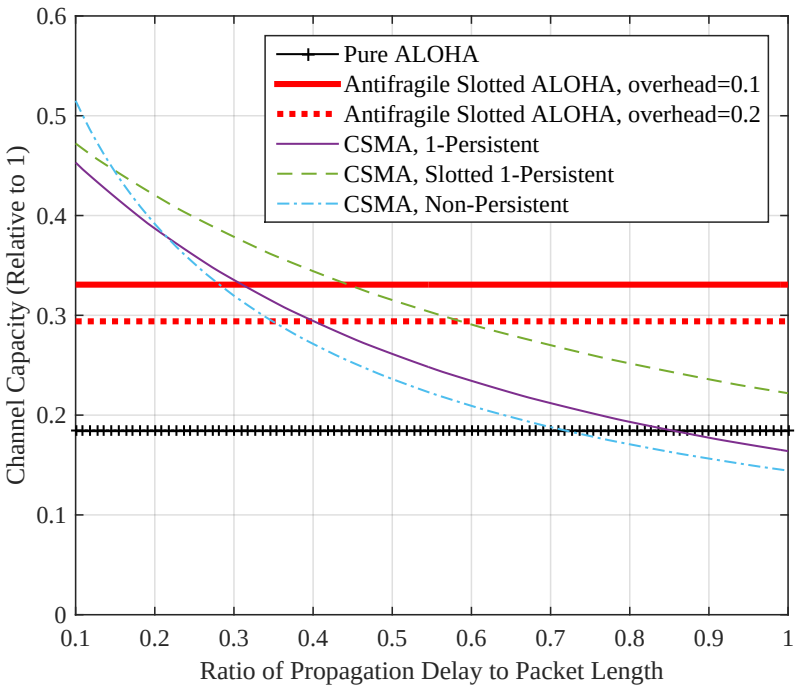

Figure 6.5: Comparison of MAC-layer channel capacity for various ALOHA and CSMA-based schemes. The ratio of propagation delay to packet length is varied, to take into account a range of tactical scenarios. For the antifragile scheme, overhead is set to a constant 0.1 and 0.2, representing practical values for a pulsed jammer.

even though the Slotted CSMA also requires a global synchronization source. The throughput of Slotted 1-Persistent CSMA is [105]:

$$S_{S1P-CSMA} = \frac{Ge^{-G(1+a)}\left[1+a-e^{-aG}\right]}{(1+a)(1-e^{-aG})+ae^{-G(1+a)}} \tag{6.9}$$

While plotting S vs. G can be useful, finding the maximum S over all values of G provides us with the MAC-layer form of channel capacity, where 1 is the maximum. In Figure 6.5 we compare MAC-layer channel capacity for each of the MAC schemes discussed above. Parameter a is varied, in order to take into account a range of tactical scenarios (i.e. various distances between nodes). For the Antifragile Slotted ALOHA scheme, overhead (defined in Section 6.2.5) is set to a constant 0.1 and 0.2, representing practical values for a pulsed jammer that is co-channel with the friendly nodes.

We remind the reader that curve associated with the Pure ALOHA scheme and all CSMA curves assume the radios are communicating in a clear portion of spectrum. Meanwhile, the Antifragile Slotted ALOHA scheme is using the antifragile technique described in this section, along with the Slotted ALOHA MAC technique (the synchronization provided by the jammer enables the use of Slotted ALOHA over Pure ALOHA).

Based on the results in Figure 6.5, we can see that as a increases, the performance of CSMA in general decreases, while ALOHA remains constant. This is because as a increases, so does the vulnerable period of a packet, as nodes far away cannot detect activity on the wireless channel until after it has started (remember that ALOHA does not involve sensing the medium). This is one reason why CSMA/CA schemes are widely used in Local Area Networks (LANs) that have a typical radius on the order of the size of a home or floor in an office building, but less common in longer range tactical communications. While not shown, we note that P-Persistent CSMA schemes will fall in between the CSMA curves shown in Figure 6.5. Secondly, we see that the Slotted ALOHA, even with overhead taken into account, is able to outperform Pure ALOHA. This is an indication that the Antifragile Slotted ALOHA scheme has value, at least compared to the other contention-based MAC techniques we investigated. For a non-contention-based MAC, we recommend TDMA, as discussed at the beginning of the section.

6.2.8 Conclusion and Further Work

In this section we have proposed a method for using a signal of opportunity as the source for a coarse synchronization signal that nodes in a given area can use to coordinate transmissions and perform rendezvous. It has been shown to have the antifragile property as long as the signal of opportunity originates from the enemy forces, such as a pulsed jammer. Further work includes investigating:

1. Distinguishing between the signal of opportunity and other traffic or interference on the channel.
2. Using the proposed strategy purely for rendezvous, using another channel for data transmission.
3. Implications of having the enemy (unwittingly) define the coverage region.
4. How to tie the proposed strategy into the scheduling process, so that nodes with more information to send have more resources.

6.3 Control Channel Through Jammer Piggybacking

A discussed briefly in Chapter 4, one application of jammer piggybacking is to create a network-wide control channel that all wireless nodes in a given area can receive. In this section, we expand on this antifragile strategy. The basic concept is for one node in the network, which we will refer to as the master node, to inject a control channel signal into the jammer being exploited. If the jammer piggybacking was successful, and the jammer is transmitting at a high enough power (which is not under our control), then the result is a high power broadcasted control channel that all the nodes in the vicinity of the jammer

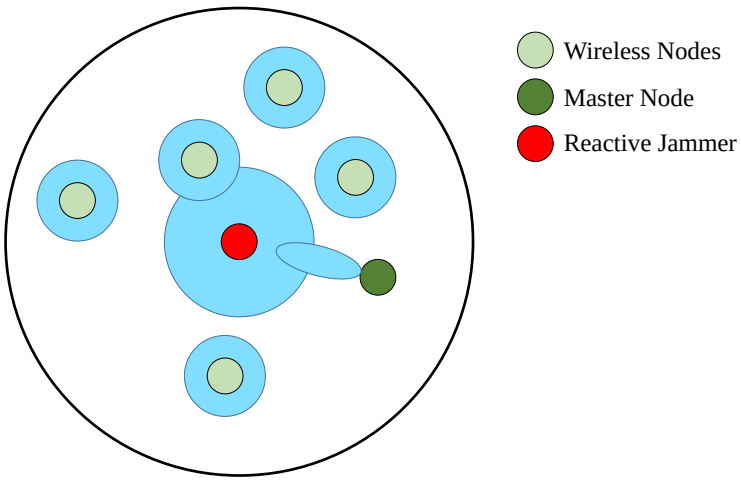

Figure 6.6: Visualization of the control channel piggybacking strategy, where one node in the network is tasked with injecting a control channel signal through the jammer being exploited, preferable with a directional antenna. Each node's antenna radiation pattern is shown in blue, and is proportional to the transmit power.

can receive. Thus, for a multi-hop network, the jammer's transmit power and antenna radiation pattern define the network coverage area. Under the ideal scenario, the master node has a directional antenna that can be pointed towards the jammer, to 1) maximize SNR at the jammer's receiver and 2) minimize interference with other nodes in the network. Figure 6.6 illustrates this idea, showing each node's antenna radiation pattern in blue. As demonstrated in Chapters 4 and 5, jammer piggybacking often does not lead to a high throughput relayed signal. However, a control channel is typically of a much lower data rate than the communications carrying data, and carries information such as resource allocations, acknowledgments, or routing tables. It is for this reason that control signalling is a great candidate for applying jammer piggybacking in practical scenarios.

Under this strategy, the antifragile gain comes in the form of increased network connectivity, assuming the jammer has a much higher limit of its transmit power than the wireless nodes. Especially in a centralized network, the increased transmit power from the jammer can allow a base station or wireless access point to synthetically increase its range, which may lead to it reaching more nodes. In addition, if a directional antenna is used, it is possible to achieve an energy-based gain (similar to Chapter 5) because the master node will not have to use as much power to transmit the control channel to all wireless nodes in the area.

In addition to a control channel, the master node could also relay a control signal (e.g., a synchronization signal) through the jammer. However, because of the non-cooperative nature of the jammer, this control signal will have limited precision as a timing synchronization signal.

Chapter 7

Analysis of Reactive Jamming against Satellite Communications

A portion of the material in this chapter has been previously published in [5], and is being reproduced in this dissertation with the consent of all co-authors, as well as the original publisher if required.

7.1 Introduction

In this chapter, we analyze the feasibility of performing and mitigating reactive jamming in a satellite communication (SATCOM)-type scenario. Even though jamming techniques such as repeater jamming have been known for decades [16], recently there has been research into other forms of complex jamming with receiving and processing capabilities. This broad category of jamming is known as reactive jamming [51].

A reactive jammer is a type of jammer that has the ability to sense one or more channels, and immediately transmit a jamming signal when it senses a signal it wants to jam [51]. In a frequency-hopping spread spectrum (FHSS) system, this allows the jammer to only transmit when the target is transmitting, as well as only transmit on the target's current frequency, without needing to know the hop sequence. The jammer is then able to conserve power, achieving a gain roughly equal to the FHSS processing gain, and remain harder to detect. Typically a reactive jammer will not simply retransmit the received signal, because this could potentially help the target communications system. Instead, the jammer may noise modulate the received signal, or simply transmit noise on the detected center frequency.

Military satellite communications (MILSATCOM) systems provide worldwide secure communications to military units in the field of operations. In the United States, MILSATCOM systems include Milstar, UFO, MUOS, and AEHF. These communications systems typically include satellites in geosynchronous orbit (35,786 km above the surface), equally spaced along an orbit above the equator. They also tend to use spread spectrum techniques such as direct sequence spread spectrum (DSSS) and FHSS, to provide protection against jamming [15, 107]. The operating frequency of MILSATCOM systems varies between about 2 - 44 GHz [108–110], with the uplink and downlink occupying separate channels. In order to provide secure communications, high-gain steerable antennas (such as phased arrays) are used. In some cases, MILSATCOM systems also include nulling beams that provide out-of-beam interference cancellation.

The analysis in this chapter assumes a 40 GHz uplink carrier (Earth-to-space) and 20 GHz downlink carrier (space-to-Earth). A 2 GHz hopping bandwidth will be used as a case study, because this is roughly the average width of each geostationary SATCOM band [111, 112]. There are some SATCOM systems in Low Earth Orbit (LEO), such as Iridium and Globalstar, but they provide a low data rate and are not considered robust to jamming in the first place. We will remain agnostic of whether the satellite has onboard processing, or whether it uses a transponder, because we will assume the reactive jamming occurs on the uplink or downlink signal between the satellite and ground user/station in a hostile area, as shown in Figure 7.1.

It is well-known that jamming a FHSS system requires either knowing the hopping sequence beforehand and being time-synchronized or dynamically detecting the signal at some frequency and immediately transmitting on that frequency [81]. The latter jammers are called reactive, because they react to the signal they are attacking in realtime. In order for reactive jamming to be feasible, there are three main requirements that must be satisfied:

1. To be able to detect the actual transmission frequency, the jammer must receive the target signal at a high enough signal-to-noise ratio (SNR).
2. To degrade communications, the jamming signal must reach the target receiver with a high enough jammer-to-signal ratio (JSR).
3. The jamming signal must reach the target receiver quickly enough to overlap in time and frequency with the target signal.

The goal of any countermeasure is to prevent or mitigate the threat. Therefore, the countermeasure presented in this chapter will focus on preventing the adversary from meeting one of the above requirements (the one that requires the least amount of effort to foil).

The remainder of this chapter is organized as follows. Related works are discussed in Section 7.2. Before proposing reactive jamming countermeasures, a framework for modeling reactive jamming under a SATCOM-type scenario is developed and organized into the three requirements listed above. These three requirements are discussed in Sections 7.3, 7.4, and 7.5 respectively. As a form of validation, simulation results related to the analysis are provided in Section 7.6. Section 7.7 proposes a reactive jamming mitigation strategy to be used when the hop rate cannot be further increased; a strategy developed using insight from the previous sections. Section 7.8 concludes. It is assumed that the reader has (at a minimum) general knowledge in the field of communications.

7.2 Related Works

Reactive jamming is widely discussed in open literature [53, 56, 113, 114]. The authors of [51] introduce the concept of classifying jammers into the categories of constant, deceptive, random, and reactive. A large portion of reactive jamming literature is based on wireless sensor networks (WSN) [53, 54, 115–117]. While these papers have similar forms of analysis and often focus on jammer detection, the proposed countermeasures are rarely applicable to a SATCOM-type system, due to the fundamental differences between SATCOM and WSN (e.g. most proposed WSN countermeasures involve cooperative sensing).

The author of [81] discusses fundamental limitations on reactive and repeater jamming in frequency-hopping communications, focusing on the geometric constraints. The author comes to the conclusion that this type of jamming can be best mitigated by hopping at a faster rate; of which the specific minimum rate is derived in the paper. This differs from the countermeasure we propose, because many FHSS systems either have a constant hop rate, or are already hopping as quickly as they can.

Several papers investigate the performance of DSSS under reactive/repeater jamming [118–120]. These papers show that reactive jamming is more effective than methods such as

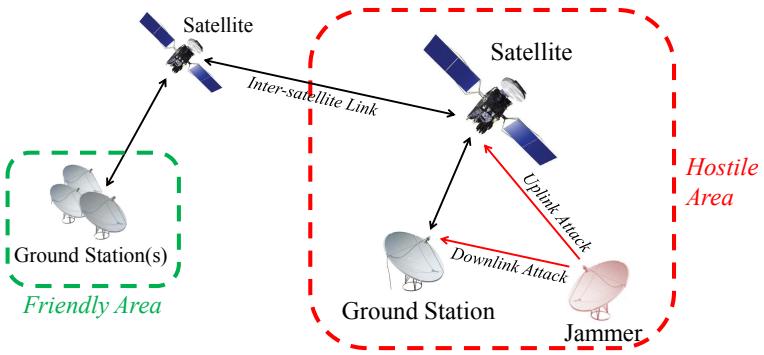

Figure 7.1: A SATCOM jamming scenario involving an inter-satellite link

barrage, partial-band, and multi-tone jamming. The core analysis of these papers have little intersection with that of this chapter, but the conclusions provide a foundation and a reason to perform the analysis in this chapter.

Link budgets related to SATCOM in literature are abundant, such as the ones in [121] and [122]. However, to the best of our knowledge there have been no collections of link budgets and supporting analysis that are specific to **reactive jamming** in a SATCOM-type scenario. This is most likely because reactive jamming is a subset of jamming, and because our analysis focuses on a SATCOM-type scenario.

7.3 Received Signal-to-Noise Ratio at the Jammer

For reactive jamming to be feasible, **the jammer must be able to receive the target signal at a high enough SNR**, so that it can either detect the center frequency, or be able to retransmit it without also transmitting excess out-of-band noise. In order to remain agnostic of the specific waveform (e.g. modulation scheme), we will only be concerned with the received signal power at the jammer P_R, the channel noise power P_{noise}, and their ratio given by $SNR = P_R/P_{\text{noise}}$. We will first determine a rough estimate for the minimum SNR needed at the jammer, and then determine how close a jammer would have to be to the transmitter in order to achieve this SNR requirement. The latter analysis is performed for both an uplink and downlink attack separately. It should be noted that the analysis in this section also applies to the general problem of signal interception.

7.3.1 SNR Threshold for Repeating the Signal

The threshold for an acceptable SNR depends on whether the jammer is digital or analog. If it is a digital reactive jammer that detects the center frequency and bandwidth, and then generates a jamming waveform, the SNR must be high enough for this detection process.

If we assume the jammer has no a priori knowledge of the target waveform, then the only difference between observing a signal and observing noise is the statistical average energy they contain. Therefore, the optimum detector compares the average energy in an observed waveform to a threshold, also known as an energy detector or radiometer [92]. Probability of detection, P_D, of a Neyman-Pearson type energy detector is parameterized by SNR and number of samples, and given by [92]

$$P_D = 1 - \Gamma\left(\frac{n}{2}; \Gamma^{-1}\left(\frac{n}{2}; 1-\alpha\right)(1+SNR)^{-1}\right) \tag{7.1}$$

where $\Gamma(x, y)$ is the incomplete gamma function, α is the false-alarm probability, and n is the number of samples taken from the observed waveform. False-positives are not that critical in this scenario because they only correspond to a slight waste of power consumed by the jammer, so we can set $\alpha = 0.1$. False-negatives have more impact from the perspective of the jammer because they represent communications that won't be jammed, so we will use a stronger value of $P_D = 0.95$, corresponding to a 95% probability of detection when the signal is actually present. The value n determines the delay involved in the detection process, which will come into play when we investigate the geometric component of the problem in Section 7.5. If the jammer must detect the signal reasonably quickly, we can set $n = 10$, which gives us a required SNR of 5 dB. Waiting much longer than around 10 samples would prevent the jammer from being "quick enough", as discussed in Section 7.5, and reducing the amount of samples observed significantly increases the SNR required at the jammer.

For the above analysis to hold, we must assume that the bandwidth after direct-sequence spreading is constant and known to the jammer, which is a realistic assumption because it only needs to be detected once. We also assume that the jammer channelizes the received waveform into bins the size of the bandwidth after direct-sequence spreading and is able to receive the entire band of FHSS operation simultaneously using a wideband front-end (or multiple front-ends).

In the case of an analog jammer, a high SNR would be desired to avoid retransmitting excess noise and thus wasting power. At 5 dB of SNR, the jammer will transmit about 75% of the desired signal and 25% noise ($\frac{0.75}{0.25} = 3 \approx 5$ dB), which is reasonable considering the gain achieved by using reactive jamming. For the remainder of this chapter we will assume a 5 dB SNR is required at the jammer's receiver in order to perform reactive jamming.

Now that we have an estimated SNR threshold, we can assemble an example link budget for a realistic scenario. We will first analyze an uplink jamming attack, then a downlink jamming attack.

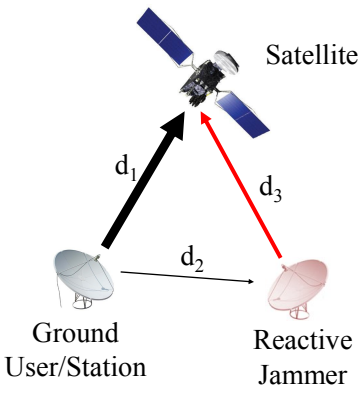

Figure 7.2: Uplink System Diagram

7.3.2 SNR during Uplink Jamming

An uplink attack involves jamming the uplink signal, which is sent from the ground user or ground station and received by the satellite. Thus, the jammer is receiving the ground user's signal and transmitting noise on the detected signal band towards the satellite's receiver, as shown in Figure 7.2. Alternatively, the jammer could retransmit the ground user's signal with a transform applied, as long as the transform does not significantly shift the frequency. In this subsection we are only concerned with how well the jammer can receive the ground user's signal.

The link budget for the target signal is given by [122]

$$P_R = P_T + G_T + G_R - L_{\text{path}} - L_{\text{atm}} - L_{\text{misc}} - L_{\text{process}} \tag{7.2}$$

and noise power is found using the traditional "kTB" method [122]

$$P_{\text{noise}} = k + T + B \tag{7.3}$$

The following list describes each parameter:

- P_R - Received signal power
- P_T - Transmitted signal power
- G_T - Transmitting antenna gain
- G_R - Receiving antenna gain
- L_{path} - Path loss
- L_{atm} - Additional atmospheric attenuation
- L_{misc} - Miscellaneous losses not yet accounted for
- L_{process} - DSSS processing gain

- k - Boltzmann constant (approx. -228.6 dBW/K/Hz)
- T - Noise temperature in dB-Kelvins
- B - Bandwidth of signal in dB-Hz (before spreading)

In order to calculate link budgets for this scenario, we must approximate each parameter listed above. These approximations are made for the sake of obtaining insight into the feasibility of reactive jamming in a SATCOM scenario. Therefore, when possible (and publicly available), parameters will be taken from existing SATCOM systems.

The uplink center frequency f will be centered around 40 GHz as discussed earlier.

P_T is the transmit power of the ground user. To improve traceability of this analysis, we will use information from commercial SATCOM systems in geosynchronous orbit. NASA's Tracking and Data Relay Satellite (TDRS) has a user equivalent isotropically radiated power (EIRP) of 48.5 dBW in the high data rate mode [123], and the ViaSat VMT-1220 ground terminal states a similar figure [124]. Likewise, the ThinSat Talos Integrated Antenna/Terminal Subsystem is specified to have an equivalent EIRP of 47 dBW. Because $EIRP = P_T + G_T$, we must subtract out the estimated main-lobe antenna gain of 35 dB; an approximation made in Section 7.4.1. This leads us to the rough estimation of $P_T = 13.5$ W or roughly 11 dBW, which will be used for analysis.

G_T is the gain of the ground user's transmitting antenna with respect to the jammer in this case, because we are investigating with what power the jammer receives the ground user's signal. This is a tricky value to estimate, because the ground user will be pointed at the satellite, not the jammer. It is highly unlikely that the jammer will be located in the main lobe of the ground user's antenna, unless it is onboard an airborne platform that is flying directly between the ground user(s) and satellite. This modeling dilemma is depicted in Figure 7.3. Clearly there is isolation provided by the directional antenna, but the question is, how much? At any given time, the jammer may be pointed into the peak of a side lobe, or into a null. Regardless, in order to put together a link budget, we need to model the "average gain". For now, a gain of -10 dB will be used for G_T. This can be thought of as isolation provided by the ground user's transmitting antenna. A more elegant solution to this dilemma will be the subject of future research, and may involve a stochastic component.

G_R is the gain of the jammer's receive antenna, which could be anywhere from omni-directional to highly directional. If the jammer knows roughly where the target ground users are, then it can use a directional antenna pointed at them (although it would lack any sort of pointing/beamforming feedback). We will therefore assume a 15 dB antenna is used, to take into account a small amount of directionality.

Path loss L_{path} is a function of the channel between the ground user and jammer. In order to perform uplink reactive jamming, the jammer must be within the same spot beam as the ground user whose signal is being jammed. Therefore, the distance between the ground user and jammer must be no more than the typical beam width, and is likely much less. For the

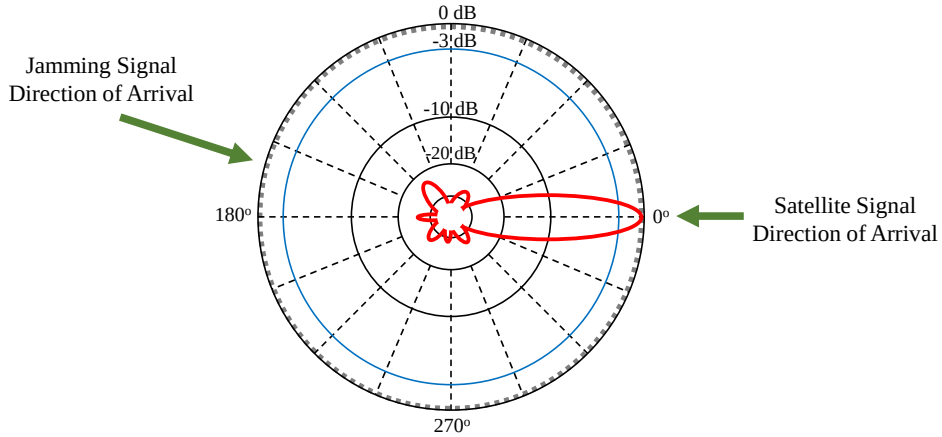

Figure 7.3: Example Radiation Pattern of a Ground User's Directional Receive Antenna

sake of analysis we will vary the distance between 1 - 50 km so that the jammer is likely to be well within the coverage area of the beam. For the channel model we will use a "line-of-sight (LOS) free space plus reflection and multiple diffraction" channel model [121] to capture a rural type environment at 40 GHz, where the jammer places itself at a strategic location (e.g. on the top of a mountain). Recall that the ground station is located in a hostile area, such as a forward operating base. Using this model the attenuation as a function of frequency f in MHz and distance d in km is given in dB using the equation:

$$L_{\text{path}} = 32.45 + 20\log_{10} f + 20\log_{10} d + A_r + A_{fr} \tag{7.4}$$

where A_r and A_{fr} capture reflection attenuation and Fresnel zone obstruction attenuation respectively. We will be using 5 dB for both terms for the sake of approximation, as done in [121]. Rain attenuation is left out of the link budget in order to reduce the amount of variables during analysis. Multipath fading and shadowing are left out so as to not make too many assumptions about the specific geographical features between the jammer and ground user. We note that this channel model is not intended to cover every possible scenario; it is a simple model that is being used an example.

L_{atm} is zero (or at least negligibly small) because the jammer is receiving the ground user's signal.

L_{misc} will be set to 2 dB to take into account miscellaneous losses (e.g. cable loss).

L_{process} is the processing gain (a.k.a. spreading gain) associated with the DSSS. This analysis is focused on detecting the instantaneous center frequency used by the ground user, and therefore the FHSS processing gain (i.e. the ratio of the hopping bandwidth to the bandwidth after DSSS spreading the narrowband signal) does not come into play in the link budget. As a case study we will use the processing gain of the 802.11b/g DSSS implementation, which is about 10 dB [125]. This DSSS spreading gain figure can be thought of as the ratio of the narrowband bandwidth to the instantaneous radio frequency (RF) bandwidth.

The system noise temperature T is the sum of the antenna noise temperature and receiver noise temperature, where the receiver noise temperature is a combination of several noise sources between the antenna terminal and receiver output [107]. A table of typical system noise temperatures in SATCOM links is provided in [107]. The typical value of T for a 40 GHz carrier when the antenna is pointed at Earth is listed as 763 K or 28.8 dBK. Even though a jammer's receiver will likely not be as high quality as that on-board a communications satellite, we will use this figure as a rough approximation.

Any noise temperature can also be expressed as a noise *figure* (denoted as F), where T_0 is a reference temperature.

$$F = 1 + \frac{T}{T_0} \quad \text{[Kelvin]} \tag{7.5}$$

The reference temperature T_0 for a receiver is typically 290 K [107]. The receiver noise figure is largely based on the quality and design of the receiver components. For example, if a receiver is cryogenically cooled it will have a lower noise figure; an example of 29 K is given in [107].

$$T = 290(10^{F_{noise}/10} - 1) \quad \text{[Kelvin]} \tag{7.6}$$

The figure 290 K is based on typical room temperature. In the case of the jammer pointing at the ground user with a 15 dB gain receive antenna, we will approximate $F_{noise} = 3.15$ or $T = 310$ K, due to the fact that the jammer is pointed towards the surface and not the sky.

The bandwidth before direct-sequence spreading B is largely based on the symbol rate, which we will assume to be 1 MSps. The occupied RF bandwidth will be found as if it were a single-carrier signal utilizing raised cosine filters, using the equation below.

$$B = R_S(1 + \alpha) \tag{7.7}$$

where α is the roll-off factor, which is often around the value of 0.25. Using our assumptions, this comes out to $B = 1.25$ MHz, which is a reasonable amount considering 10 dB of DSSS spreading is assumed, and the DSSS signal is then hopped in frequency across 2 GHz of bandwidth. This means the SATCOM link hops across 160 adjacent channels, because $(1.25 \text{ MHz})(10)(160) = 2$ GHz. We note that 160 hopping channels corresponds to a FHSS spreading gain of 22 dB (simply 160 converted to dB).

Using these many approximations, we can now calculate the SNR while varying the distance between the jammer and ground user. The results are shown in Figure 7.4, with the horizontal line representing the threshold for frequency detection (discussed earlier). It can be seen that our SNR threshold is reached at about 8 km. Therefore, the jammer must be within this distance of the target ground user in order to be able to detect the signal and effectively perform reactive jamming.

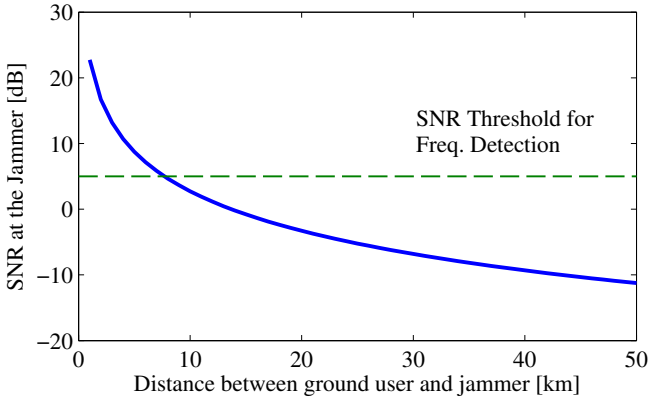

Figure 7.4: SNR at jammer when receiving the ground user's signal

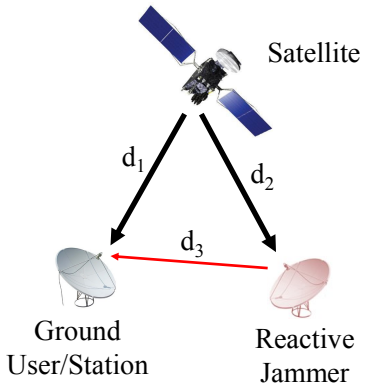

Figure 7.5: Downlink System Diagram

7.3.3 SNR during Downlink Jamming

A downlink attack involves injecting noise into the ground user's receiver, and therefore the jammer must receive the signal transmitted by the satellite, as shown in Figure 7.5. To find how well the jammer can receive this signal, we will re-investigate the parameters described in the uplink attack subsection, but for the downlink jamming scenario.

L_{path} is now the distance between the jammer and satellite, and is considerably larger than that of the uplink attack. Geostationary orbit is located 35,786 km above the Earth's surface, in which the surface is defined by the mean sea level. If we assume that the jammer is in the same meridian as the satellite, then we can use the latitude angle of the jammer to calculate the actual distance d between the two antennas.

$$d = \sqrt{R^2 + r^2 - 2Rr\cos\theta} \tag{7.8}$$

where R is the radius of the Earth (roughly 6.37e6 m), r is the orbital radius of the satellite (42e6 m for geostationary orbit), and θ is the latitude of the jammer. Thus, for our analysis

this equation can be reduced to

$$d = \sqrt{1.818e15 - 5.37e14 \cos\theta} \tag{7.9}$$

To find the distance between the jammer and satellite, we will consider three scenarios, in which the jammer is located at 0°, 45°, and 60° latitude. Table 7.1 shows the distance d between the jammer (or ground user) and satellite, as well as the propagation delay. The downlink carrier is at 20 GHz as discussed earlier, and in this case we calculate L_{path} using the free space path loss equation (with an atmospheric attenuation loss added separately):

$$L_{\text{path}} = (4\pi df/c)^2 \tag{7.10}$$

P_T is the transmit power of the satellite, which we will estimate to be 100 W or 20 dBW, based on real-world examples in [122].

G_T is the satellite's transmit antenna gain with respect to the ground user's location, but because the jammer must be within a beam width of the ground user (as shown later in this chapter), we will assume that the jammer experiences the same transmit antenna gain, estimated to be 40 dB.

G_R is now much higher because we will assume the jammer has a directional antenna pointed at the satellite. It is reasonable to assume that the jammer has accurate knowledge of the satellite's coordinates, because it is in geosynchronous orbit and the orbit locations are in the public domain. We will vary this figure between 10 dB and 40 dB.

L_{atm} is nonzero now that the signal is passing through the entire atmosphere. We will use an approximate figure of 1 dB for the uplink and downlink based on the recommendations in [126].

L_{misc} will again be set to 2 dB to take into account miscellaneous losses.

The noise temperature T is different than the uplink case because we are now dealing with a 20 GHz carrier and a receiver pointed into space (i.e. away from the noise created by Earth). Once again using the reference table in [107], we will approximate T to be 424 K or 26.3 dBK.

We will assume the bandwidth of the downlink signal is the same as the uplink: $B = 1.25$ MHz.

Table 7.1: Distance d between the jammer (or ground user) and satellite, as well as the propagation delay

Latitude	0°	45°	60°
Distance d	35793 km	37928 km	39367 km
Propagation Delay	0.1193 s	0.1264 s	0.1312 s

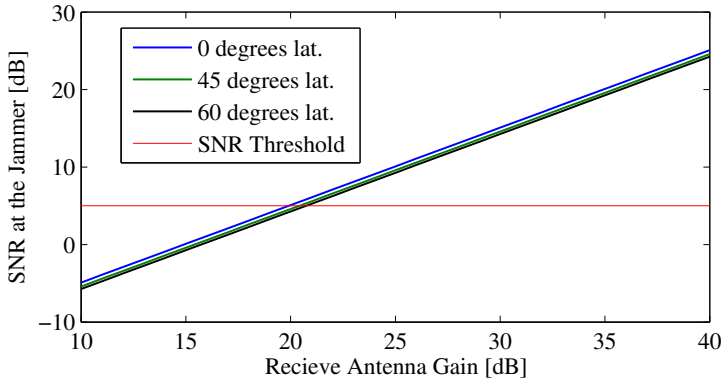

Figure 7.6: SNR at jammer when receiving the satellite's signal

The SNR while varying G_R for the downlink jamming case is shown in Figure 7.6, parameterized by the geographical location of the jammer in latitude. Based on the SNR threshold discussed at the beginning of Section 7.3 (shown as a red line), the jammer must obtain a receive antenna gain of at least 20 dB for reactive jamming to be feasible. For example, the jammer may choose to have a fairly compact 20 dB gain parabolic reflector antenna pointed towards the satellite.

7.4 Jammer-to-Signal Ratio Component

In order to fully corrupt a communications link, **the jammer must cause the bit error rate (BER) after forward error correction to be significantly high**. Either the data must be significantly corrupted, or the control information must be corrupted to the point in which the data can not be successfully decoded. For example, if the control information containing the modulation and coding scheme for the next frame is lost, then data won't be retrieved at the destination. While the data error rate can be discussed in terms of BER, it is also common to consider the block error rate (BLER) or packet delivery ratio (PDR) [51]. In FHSS systems, hop error rate can also be used as a metric. Errors on the control information, on the other hand, are more difficult to quantify because the BER metric has much less meaning. For the rest of this analysis we will only discuss errors on the data, and use the BER metric (this can be thought of as indirectly taking into account erroneous control information).

The worst-case BER is 0.5, although a BER around 0.1 (after forward error correction) will likely cause Denial of Service (DoS), due to the number of retransmissions and dropped packets. Even if the information contains realtime voice or video data, which cannot be retransmitted due to their time sensitivity, the codec would be unable to produce intelligible audio/video. The actual BER/BLER threshold is based off of numerous factors on many different layers, and would be best acquired empirically [11]. Therefore, we will continue by assuming a BER threshold to cause DoS of 0.1 or 10% of the information bits received in error after demodulation and decoding. The minimum JSR to cause a BER of 0.1 is mostly dependent on the modulation scheme and the method of forward error correction (a.k.a. channel coding). Assuming the communication system supports adaptive modulation and coding, in a harsh environment it will likely ratchet down to binary phase-shift keying (BPSK) or quadrature phase-shift keying (QPSK) using a low rate coding scheme. To gain further insight we observe a frame error rate (FER) curve for BPSK in an additive white Gaussian noise (AWGN) channel using low-density parity-check (LDPC) coding at various rates (based on the DVB-S2 protocol), shown in Figure 7.7.

It is clear that around -2 dB of JSR (or 2 dB of SNR), even the most advanced coding reaches a point where the amount of errors would cause link failure.

Given details of a communication link, JSR can be found by calculating the received friendly signal power P_R and received jamming signal power $P_{\mathrm{R-Jam}}$. Channel noise is ignored because the received jamming power will be significantly higher than the received noise power under successful jamming. P_R and $P_{\mathrm{R-Jam}}$ are found using the link budget in Equation 7.2, except $L_{\mathrm{process}} = 0$ for P_R because the communications receiver (either at the ground station or satellite) is able to de-spread the DSSS.

In the following two sections we evaluate these link budgets for the uplink and downlink jamming scenarios respectively.

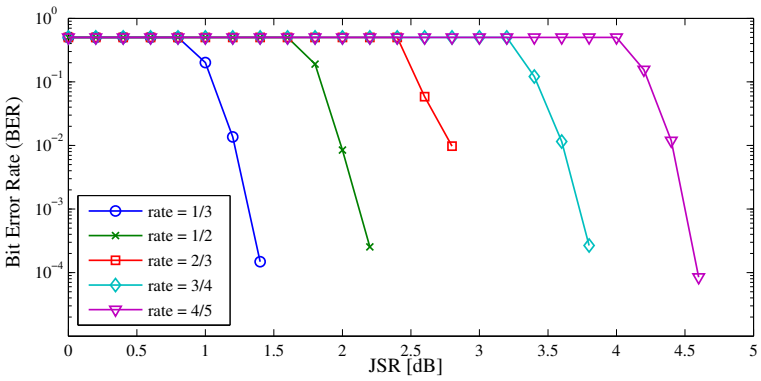

Figure 7.7: BER curve for BPSK in an AWGN channel with LDPC coding at various rates

7.4.1 Uplink Jamming JSR

We will carry over the assumptions made in the previous section, in order to map the analysis to a realistic scenario. In addition, we will assume that the jammer has a directional antenna pointed at the satellite; without this capability the attack is not feasible. We will also assume that the jammer is located close enough to the ground users such that the jammer is located inside the spot beam. Previously we approximated the ground user's transmit antenna gain to have a -10 dB gain in the direction of the jammer, but now we will only care about the main-lobe gain, which we will approximate to 35 dB after cable losses (a figure which is included in [123]). The transmit power of the jammer is set at 10 dB higher than the ground user's, which accounts for the 10 dB loss associated with the DSSS processing gain. The result is a JSR of 0 dB (due to the other gains and losses being equal), which is above our threshold of -2 dB and will likely cause degradation of the link or full denial of service. Table 7.2 provides the full link budgets.

In summary, feasibility is achieved by the jammer transmitting at a much higher power than the user, to compensate for the DSSS processing gain. If the SATCOM system does not use DSSS, then the jammer would simply have to transmit at around the same power as the user.

7.4.2 Downlink Jamming JSR

We will now use the same approach for the downlink jamming scenario. The jammer will still use a 15 dB gain antenna pointed toward the target ground user(s), with a $P_T = 21$ dBW. We will start off by investigating a scenario in which the jammer is 10 km from the ground user. Using the rural path loss model described in Equation 7.4, this comes out to a path loss of 148.5 dB (for the jamming signal). Table 7.3 provides full assembly of the link budgets.

When the jammer uses a transmit antenna gain of 15 dB, as discussed in Section 7.3.2, there is no way to achieve a high enough JSR for successful jamming, mainly due to the high

Table 7.2: Uplink Jamming Attack Link Budget to Calculate JSR

Jammer Power Budget			**Signal Power Budget**		
P_T	21	dbW	P_T	11	dbW
G_T	35	dB	G_T	35	dB
G_R	40	dB	G_R	40	dB
L_{path}	-216	dB	L_{path}	-216	dB
L_{atm}	-1	dB	L_{atm}	-1	dB
L_{misc}	-1	dB	L_{misc}	-1	dB
L_{process}	-10	dB			
			P_R	-132	dbW
$P_{\text{R-Jam}}$	-132	dbW			
	$JSR =$	0.0 dB			

antenna gains associated with the signal power budget. For the jammer to achieve an extra 16.5 dB and even out the budgets, it could either use a 32 dB gain antenna, which is possible if it knows exactly where the target is, or be positioned 1 km from the target instead of 10 km.

7.5 Geometric Component

So far we have investigated the first two reactive jamming constraints. However, even if the SNR and JSR requirements are met, the adversary may still fail to cause denial of service (a desirable outcome for our countermeasure discussed later). In order for a reactive jammer to be effective, **the jamming signal must reach the target receiver before it hops to a new frequency**. Therefore, the geometry of the system is an important factor when analyzing the feasibility of reactive jamming. In addition, the delay between when the jammer receives the target signal and transmits the jamming signal, known as the jammer's processing delay, should be included in this component.

Figure 7.8 shows the time/frequency behavior of a repeater jammer (red) jamming a target signal (green). Despite the time delay between the friendly signal arriving and jamming signal arriving, the signals are vulnerable to jamming.

In order to quantify the impact of jamming, we can abstract the physical layer by assuming a certain fraction of each hop must be jammed at the receiver in order to cause denial of service; a fraction we will denote as η. The value of η largely depends on the channel coding scheme, interleaving, and JSR. Using this abstraction, for a repeater jammer to be successful we must have

Table 7.3: Downlink Jamming Attack Link Budget to Calculate JSR

Jammer Power Budget			**Signal Power Budget**		
P_T	21	dbW	P_T	20	dbW
G_T	15	dB	G_T	40	dB
G_R	-10	dB	G_R	35	dB
L_{path}	-148.5	dB (10 km)	L_{path}	-210	dB
L_{atm}	0	dB	L_{atm}	-1	dB
L_{misc}	-1	dB	L_{misc}	-1	dB
L_{process}	-10	dB			
$P_{\text{R-Jam}}$	-133.5	dbW	P_R	-117	dbW

$$JSR = -16.5 \quad \text{dB}$$

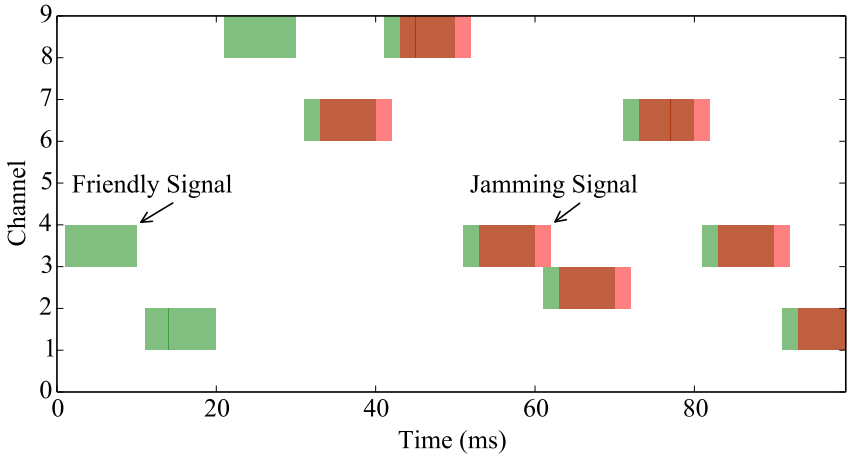

Figure 7.8: Time/Frequency Behavior of a Repeater Jammer

$$\frac{d_2}{c} + T_j + \frac{d_3}{c} \leq \frac{d_1}{c} + (1-\eta)T_h \tag{7.11}$$

where T_j is the processing plus analog delay associated with the jammer, T_h is the duration of each hop, and the distances d are as shown in Figures 7.2 and 7.5 for uplink jamming and downlink jamming respectively. This equation shows that hopping at a faster rate can be used to protect communications from repeater jamming, which is a well-known mitigation strategy [81]. To gain further insight, we must estimate the fraction of each hop that must be jammed, η.

7.5.1 Fraction of Each Hop that must be Jammed (η)

Fundamental limitations on the effectiveness of repeater jamming against FHSS are derived in [81]. For example, the authors show that the average symbol error probability can be formulated by introducing P_j; the probability a symbol is jammed given the jamming signal is present during the reception of that symbol. If we assume the non-jammed symbol error rate is F_{nj} and the symbol error rate under jamming is F_j, then the average symbol error probability is given by [81]

$$P_s = \frac{T_h - T_{nj}}{T_h} P_j F_j + \left(1 - \frac{T_h - T_{nj}}{T_h} P_j\right) F_{nj} \tag{7.12}$$

where T_{nj} is the duration of the hop in which the jamming signal is not present, and is equal to $T_j + (d_2 + d_3 - d_1)/c$. If the hop is not jammed at all, then $P_s = F_{nj}$. It should be noted that F_{nj} is largely determined by the SNR, while F_j is mostly impacted by the JSR.

To estimate T_h, we will assume a symbol rate of 1 MSps as discussed in Section 7.3.2. Using 50 symbols per hop we have a hop duration of about 50 μs, although in the simulation below we vary the number of symbols per hop between 10 and 500 to cover all realistic FHSS configurations (we will not take into account a fractional symbol per hop system). For now we will assume the jammer takes 4 symbols to detect the center frequency, and another 1 symbol worth of time to generate and transmit the jamming waveform (this will include the time it takes to tune the transmitter to the right freq). This results in a jammer processing delay of $T_J = 5$ μs, a figure that is discussed in more detail in Section 7.3.

We will assume that the ground user and jammer are approximately the same distance from the satellite. This is an acceptable assumption because if the jammer is within about 20 km from the ground user, then the most the two distances can deviate is only about 10 - 20 meters, because the distance between the Earth and satellite is three orders of magnitude greater than between the ground user and jammer. In an uplink attack this distance corresponds to d_1 and d_3, while in a downlink attack it is d_1 and d_2. Therefore, the feasibility of an uplink and downlink attack can be analyzed simultaneously because of the system's symmetry. In this case the minimum distance between the ground user and jammer for successful jamming becomes equal to $(T_h - \eta T_h - T_j)c$. Figure 7.9 shows the maximum distance the jammer can be from the ground user in order to successfully jam a hop, as a function of η, parameterized by varying the number of symbols per hop. The saturation on the right hand side of each curve is due to the jammer's processing delay, T_J, of 5 μs.

As seen in these results, for the 50 symbols per hop case, the jammer will have to be within about 10 km from the ground user in order to perform successful reactive jamming, due to the geometric constraints. This assumes the JSR is received at a high enough level and the SNR at the jammer is high enough to repeat the signal. While these results are largely dependent on the assumptions made, they provide insight to the limitations of reactive jamming SATCOM (e.g., it is highly unlikely to perform successful reactive jamming when d_2 is larger than around 10 km).

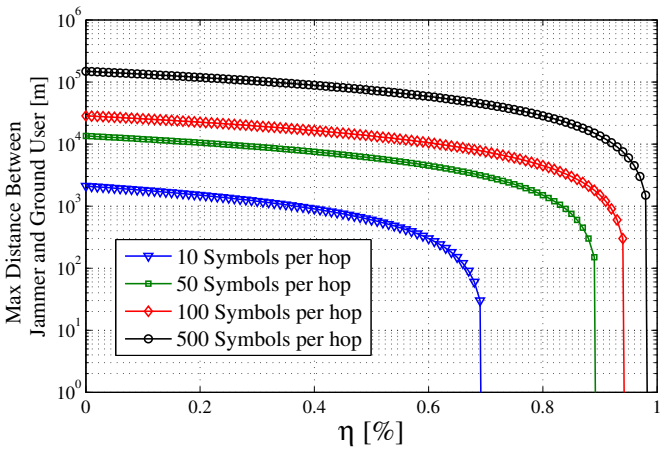

Figure 7.9: Maximum distance the jammer can be from the ground user in order to successfully jam a hop, as a function of η

Table 7.4: STK Simulation Results showing JSR and SNR for the Uplink/Downlink scenarios

Scenario	Notes	Result
Uplink SNR	10 km between user and jammer	3 dB
Downlink SNR	at 0° latitude with $G_R = 30$ dB	15 dB
Uplink JSR	same parameters as in link budget	1 dB
Downlink JSR	same parameters as in link budget	-15 dB

7.6 Simulation using Systems Tool Kit (STK)

As a form of validation, the scenarios discussed in this chapter were simulated using a software package called Systems Tool Kit (STK), a screenshot of which is shown in Figure 7.11. The orbit of a generic geostationary satellite over the Americas was used. The jammer was placed on a ground vehicle, so that the distance between the jammer and ground user could be varied between 0 and 50 km. The distance between the jammer and ground was set to 2 m. All directional antennas used a parabolic antenna model with main-lobe gains as specified earlier in the chapter. The channel and atmospheric attenuation models used were those built into the communications STK package. Figure 7.10 shows the object browser of STK, for the sake of reproducing the results.

These STK scenarios were used to verify the results found in Sections 7.3 and 7.4. The following table highlights some of the numerical results found, with notes describing the specific conditions under which the simulations were performed. In all four cases, these simulated results closely match the analysis in Sections 7.3 and 7.4, which is expected considering the more complex models do not deviate too much from the simple models used during analysis.

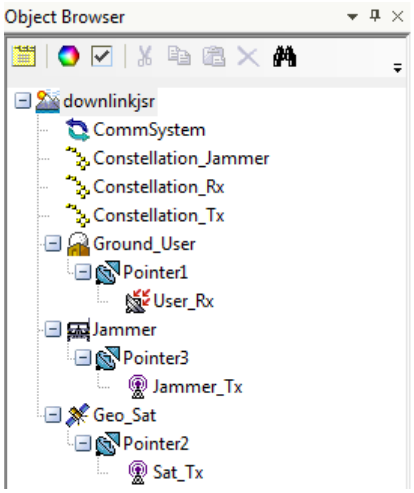

Figure 7.10: Screenshot of the STK Object Browser (downlink scenario)

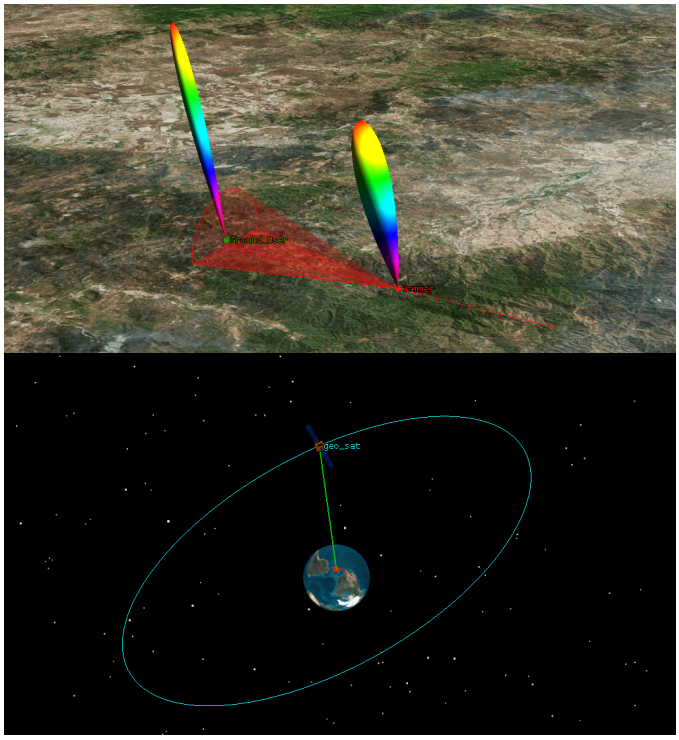

Figure 7.11: Screenshots of the uplink jamming scenario in STK

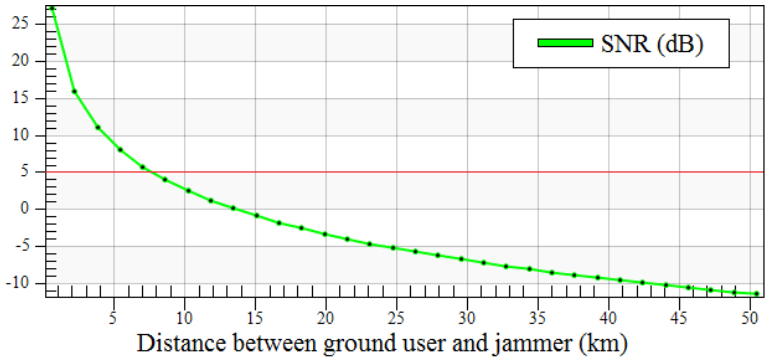

Figure 7.12: STK Simulation showing SNR measured during the uplink attack (at the jammer)

Figure 7.12 shows the results of a simulation that attempts to reproduce the previously presented Figure 7.4. Both methods show that the SNR threshold of 5 dB is reached at about 8 km. It should be noted that the antenna isolation figure of 10 dB, used throughout this analysis, was included in the simulations.

7.7 Reactive Jamming Mitigation

There exist many mitigation strategies for various forms of jamming, and one of the most common mitigation strategies for reactive jamming is to simply hop faster [81]. However, there may be situations in which the radios are hopping as quickly as they can, or do not have an adjustable hop rate. For these situations we propose a mitigation strategy that is based on the insights developed in this chapter. The following is an overview of this technique, to act as a direction for further research and consideration.

The proposed reactive jamming mitigation strategy is to use a coding and interleaving scheme that results in the uncoded transmitted bits appearing at the very beginning of each hop. The remainder of the hop can then be filled with the transmitted bits after channel coding (also known as forward error correction). The proposed mitigation strategy takes advantage of the geometric constraints of reactive jamming. The reasoning behind transmitting uncoded bits, which is unusual in modern wireless communications, is simply to increase the odds of the data being received prior to the jamming waveform. For the link to function only using the uncoded bits (e.g. in the presence of a reactive jammer), there would need to be a fairly strong signal, as well as error detection at the higher layer that checks validity of multiple hops worth of information at once. Most coding literature assumes an AWGN or fading channel in which the erasures do not correspond to the last portion of every hop or frame, but as we found out in the geometric constraints portion, jamming the very beginning of each hop is significantly challenging. This type of mitigation strategy is especially tailored to scenarios where the jammer vastly overpowers the desired

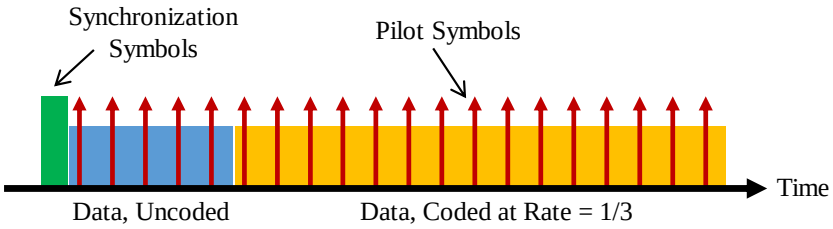

Figure 7.13: An example hop structure using the proposed mitigation strategy

signal at the receiver, leading to no reception of information during jammed portions of the hop. Most modern coding schemes such as turbo coding and LDPC have a very sharp knee in the BER curve, representing a quick drop-off between non-jammed and fully-jammed cases (i.e. there may only be a few dB of SNR difference between the two cases). Note that this strategy does not involve changing where the synchronization symbols and reference symbols (a.k.a. pilots) are placed. Figure 7.13 shows an example of how each hop could be organized according to the proposed mitigation technique. Note that both blocks of data correspond to the same information.

As demonstrated using Figure 7.9 and the related analysis, if the jammer is forced to overlap with over about 90% of the target hop to cause denial of service (i.e. $\eta > 0.9$), then reactive jamming really is not very feasible. While this specific percentage threshold varies based on number of symbols per hop, the jammer's processing delay, and the geometry of the system, the proposed anti-jam strategy can at least attempt to maximize the value of η in order to mitigate most reactive jamming scenarios. **As long as the geographical area within close proximity to the ground user is clear of reactive jamming, using this method of coding and interleaving to raise the value of η will provide strong protection against reactive forms of jamming.**

Figure 7.14 shows simulation results in which the link state is measured while varying η and the symbols per hop. In every single jammed case, it is the uplink that is denied. This is the type of figure that can be used as a guide in determining how high η must be, based on the number of symbols per hop (or how quickly to hop if η is constant).

7.8 Conclusion

In this chapter we analyzed the feasibility of performing reactive jamming (including repeater jamming) in a SATCOM scenario, using a three step approach based on the SNR, JSR, and geometric jamming constraints. Additional clarity was provided by splitting the analysis into uplink and downlink jamming. Even though reactive jamming is a complex form of jamming that requires receiving capabilities and a low processing delay, it allows a jammer to counter the processing gain associated with frequency hopping spread spectrum. It is for this reason

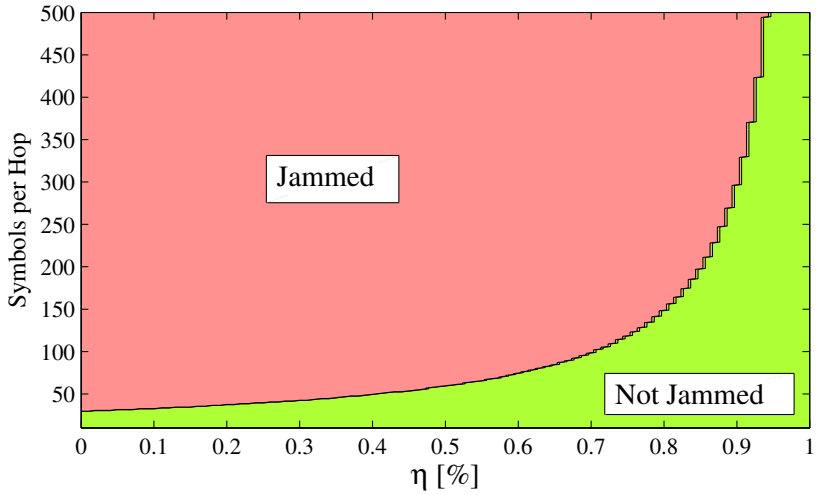

Figure 7.14: Feasible region of countermeasure

Table 7.5: Summary of Feasibility Analysis

Portion	Feasible	Notes
SNR, uplink	Yes	Jammer must be within about 8 km of ground user
SNR, downlink	Yes	20 dB or higher receive antenna gain
JSR, uplink	Yes	Jammer compensates for processing gain with P_T
JSR, downlink	No	Not without described augmentations
Geoemtry, both	Yes	As long as the jammer is $<$ 10 km from user

that we should consider reactive jamming a future threat in the SATCOM domain, especially as software-defined radios become more capable and less costly. Results of each of the three steps are summarized in Table 7.5.

The overall feasibility of uplink and downlink jamming will remain highly dependent on the scenario at hand, although we showed that an example scenario will allow for a jammer to deny communications via uplink jamming but not downlink. Because both the uplink and downlink are vital to the operation of a communications link, only one of the two needs to be jammed. Thus, the overall conclusion of this chapter is that reactive jamming is in fact feasible in a SATCOM scenario, and reactive-specific countermeasures should be considered. This work suggest that as long as the geographical area around user terminals is free of reactive jammers, mitigation can be achieved using the proposed mitigation strategy.

Chapter 8

Reinforcement Learning for Reactive Jamming Mitigation

A portion of the material in this chapter has been previously published in [8], and is being reproduced in this dissertation with the consent of all co-authors, as well as the original publisher if required.

8.1 Introduction

In this chapter, we propose a strategy to mitigate or even avoid reactive forms of jamming using a reinforcement learning (RL) approach, when exploitation is not practical. Through a learning approach, the problem of having to detect and classify which type of jammer is present in real time is avoided. In addition, there is no need to preprogram a radio with specific mitigation strategies; instead the strategy is learned in real time and in the presence of the jammer. The proposed mitigation strategy focuses on finding an effective channel hopping and idling pattern to maximize link throughput. Not only can this approach enable communications, which would otherwise fail in the presence of a sophisticated and reactive jammer, it can also act as an optimization routine that controls the link layer behavior of the radio.

The proposed strategy is well suited for a frequency-hopping spread spectrum (FHSS) system, which are widely used in modern wireless communications. The strategy could also be applied to an orthogonal frequency-division multiple access (OFDMA) system in which users hop between different subcarriers or groups of subcarriers, such as the uplink of LTE. Countless users and systems depend on wireless communications and therefore it is important to secure them against jamming. While there exists many methods to counter barrage jamming (the most basic form of jamming), there are few methods that are designed to address the more intelligent behaviors a jammer can exhibit.

Note that most of this chapter appears in my published work in [8].

8.2 Related Works

Wireless security threats are typically broken up into two categories: cyber-security and electronic warfare (i.e. jamming). Electronic warfare attacks target the physical layer (PHY) and/or media access control layer (MAC) of a communication system, while cyber-security attacks are designed to exploit the higher layers. Here, we are only concerned with jamming, and in particular jamming of an intelligent nature. A series of intelligent jamming attack models are introduced in [51], including the reactive jammer model. The authors propose a basic detection algorithm using statistics related to signal strength and packet delivery ratio. For an overview on electronic warfare and jamming, we refer the reader to [15].

A RL or Markov decision process (MDP) approach has been previously used in the wireless domain for channel assignment [127], general anti-jamming in wireless sensor networks [128], and jammer avoidance in cognitive radio networks [129, 130]. The authors of [127] apply reinforcement learning to the problem of channel assignment in heterogeneous multicast terrestrial communication systems. While [127] does not deal with jamming, it has similar concepts to the techniques proposed in this chapter.

The authors of [128] propose an anti-jamming scheme for wireless sensor networks. To address time-varying jamming conditions, the authors formulate the anti-jamming problem of the sensor network as a MDP. It is assumed that there are three possible anti-jamming techniques: transmission power adjustment, error-correcting code, and channel hopping. These techniques are not explored any further; the set of actions available to the radio is simply which technique is used. While this work is similar to the technique described in this chapter, it greatly generalizes the anti-jamming strategies. In other words, this work does not offer a jamming strategy, it offers a method of choosing the best jamming strategy from a given set.

The authors of [129] use a MDP approach to derive an optimal anti-jam strategy for secondary users in a cognitive radio network. For the jammer model, the authors use reactive jammers seeking to disrupt secondary users and avoid primary users. In terms of actions, in each timeslot the secondary user must decide whether to stay or hop frequencies. The authors propose an online learning strategy for estimating the number of jammers and the access pattern of primary users (this can be thought of as channel availability). Even though the authors use a reactive jammer model similar to the one described throughout this dissertation, they assume the jammer is always successful, and the entire analysis is within the context of dynamic spectrum access.

Additional literature on antijamming using a RL framework that was published after the material in this chapter includes [131–134].

The antijam strategy described in this chapter is the first MDP-based RL approach designed to mitigate a wide range of reactive jamming behaviors. In this chapter we provide initial insights into the feasibility and suitability of the strategy, leaving implementation for future work.

8.3 System Model and Problem Formulation

Consider the typical wireless communications link, with the addition of a jammer that both receives the friendly signal (but not necessarily demodulates it) and transmits a jamming signal, as shown in Figure 8.1. For the sake of simplicity we will only consider a unidirectional link, although this analysis also applies to bidirectional links, that may be unicast or broadcast, as well as a collection of links.

While reactive jamming can take on different forms, we will broadly define the term as any jammer that is capable of sensing the link and reacting to sensed information. We will assume this sensed information is in the form of the presence or absence of energy, because any additional information such as modulation scheme or actual data would be irrelevant for this mitigation strategy. A simple example of a reactive jammer is one that senses the spectrum for activity, and immediately transmits wideband noise when it senses any activity [51]. This strategy allows the jammer to remain idle while the channel is idle and thus save power and avoid being easily detected. Another form of reactive jamming, commonly known

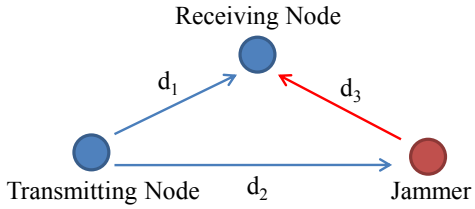

Figure 8.1: System model of a transmitter, receiver, and reactive jammer.

as repeater or follower jamming [81], works by immediately transmitting what it receives with noise added to it. This can be modeled as a jammer that senses a set of channels, and immediately transmits noise on any channel that appears to be active.

Reactive jamming is only feasible when the geometry of the system is such that the jammer's transmitted signal reaches the target receiver before it hops to a new channel or stops transmitting. As such, reactive jamming is only possible when the jammer is physically located near or between the target transmitter and receiver. If η represents the fraction of each hop duration that must remain not jammed for communications to succeed, then we have the following inequality limiting the distances d_2 and d_3 [81]

$$d_2 + d_3 \leq (\eta T_d - T_j)c + d_1 \tag{8.1}$$

where T_d is the hop duration, T_j is the jammer's processing time, c is the speed of light, and d_1, d_2, and d_3 are the distances indicated in Figure 8.1. In addition to this limitation, the jammer-to-signal ratio at the receiving node must be high enough to degrade quality of service. In this chapter we assume the jammer is close enough to the transmitter and receiver, and that the jammer-to-signal ratio is significantly high during periods of jamming.

As part of the analysis and simulation we will investigate two specific reactive jamming models. The first, labeled in this chapter as simply "reactive jamming", will be defined as a jammer that successfully jams any transmission that remains active for too long, regardless of the channel/frequency in use. The second jammer model is based on repeater jamming, and it is described as a jammer which successfully jams any transmission that remains on the same channel/frequency for too long. Note that this convention is not the same as the rest of this dissertation, although it is for a reason that will be clear later. While there are other ways to formulate reactive jamming models, the analysis and simulation in this chapter will focus on these two. More formal definitions of these two jammer models is as follows:

- **Reactive Jammer** - Begins jamming any transmission that remains active for more than N_{REACT} time steps, and will only cease jamming once the target is idle for at least N_{IDLE} time steps.
- **Repeater Jammer** - Begins jamming any transmission that remains on the same channel for more than N_{REP} time steps.

In this analysis we will investigate a transmitter and receiver pair that can hop among a certain number of channels using a FHSS approach, or any other approach that involves radios capable of changing non-overlapping channels. Therefore, at any time step, the transmitter has the option to either remain on the channel, change channel, or go idle. Because the actions of the transmitter must be shared with the receiver beforehand, it is expected that decisions are made in advanced.

It is assumed that channel quality indicators (e.g. whether or not the information was received) are sent back to the transmitter on a per-hop basis. These indicators could be binary (indicating an ACK or NACK), or they could take on a range of values indicating the link quality. Lastly, it is assumed that the receiver is not able to simply detect the presence of a jammer. This is especially true in reactive jamming, where the jamming signal largely overlaps with the target signal, making jammer detection difficult. If the ability to detect the jammer is already on the platform, it does not invalidate the proposed strategy. It would simply alter the assumptions that govern the system model in this particular analysis.

8.4 Strategy for Mitigation of Reactive Jamming

The mitigation (a.k.a. anti-jam) strategy described in this chapter is based on modeling the system as a MDP, where the transmitter is the decision maker, and uses the RL framework to learn a strategy for dealing with a broad array of reactive jamming behaviors. This strategy is in the form of a channel hopping pattern, where going idle is considered as hopping to the "idle channel" for a certain duration. However, we are not concerned with choosing the best channel to transmit on at any given time, nor identifying corrupt channels that have excessive noise. The mitigation strategy described in this chapter is designed to work in tandem with this kind of algorithm, i.e. one that indicates which specific channels are suitable for use and which are not. An example algorithm is described in [130], where the authors propose an approach for choosing the best possible channel in each time slot to avoid the jammer, by formulating the problem as a multi-armed bandit process (which is actually a form of RL, but different than the one used in this work). Lastly, we are not concerned with the PHY-layer waveform characteristics that the transmitter or jammer uses (i.e. bandwidth, modulation, type of noise, etc.). Adaptive modulation and coding can be performed alongside the proposed strategy. All of these assumptions and abstractions were meant to generalize the approach described in this chapter, so it can fit within many different communications systems.

8.4.1 Reinforcement Learning Background

RL is the subset of machine learning (ML) concerned with how an agent should take actions in an environment to maximize some notion of cumulative reward. The agent is the entity

interacting with the environment and making decisions at each time interval, and in this chapter we will consider the transmitter as the agent (although the actions it chooses must be forwarded to the receiver, as in most cognitive radio type applications). An agent must be able to sense some aspect of the environment, and make decisions that affect the agent's state. For example, reinforcement learning can be used to teach a robot how to walk, without explicitly programming the walking action. The robot could be rewarded for achieving movement in a forward direction, and the robot's action at each time step could be defined as a set of angular servo motions. After trying a series of random motions, the robot will eventually learn that a certain pattern of motion leads to moving forward, and thus a high cumulative reward. In this chapter, we apply this concept to a transmitter that learns how to hop/idle in a manner that allows successful communications under a sophisticated reactive jamming attack.

There are four main components of a RL system: a policy, reward, value function, and the model of the environment [135]. A policy (typically denoted as π) defines how the agent will behave at any given time, and the goal of a RL algorithm is to optimize the policy in order to maximize the cumulative reward. A policy should contain a stochastic component, so that the agent tries new actions (known as exploration). A reward, or reward function, maps the current state and action taken by the agent to a value, and is used to indicate when the agent performs desirably. In a communication system, a possible reward function could be a weighted combination of the link throughput, spectral efficiency, and power consumption. While the reward function indicates what is desirable in the immediate sense, the value function determines the long-term reward. A state may provide a low immediate reward, but if it leads to other states that provide a high reward, then it would have a high *value*.

The model of the environment is used to either predict a reward that has not been experienced yet, or simply determine which actions are possible for a given state. For example, it is possible to create a RL agent that learns how to play chess, and the environment would be a model of the chess board, pieces, and set of legal moves.

In RL, the environment is typically formulated as a MDP, which is a way to model decision making in situations where outcomes are partially random and partially under the control of the decision maker. The probability of each possible next sate, s', given the current state s and action a taken, is given by [135]

$$P^a_{ss'} = Pr\{s_{t+1} = s' | s_t = s, a_t = a\} \tag{8.2}$$

Equation 8.2 provides what are known as transition probabilities, and because they are only based on the current state and action taken, it assumes a memoryless system and therefore has the Markov property.

The expected reward (obtained in the next time step) for a certain state-action pair is given by Equation 8.3. The goal of a learning agent is to estimate these transition probabilities and rewards, while performing actions in an environment.

$$R^a_{ss'} = E\{r_{t+1}|s_{t+1} = s', s_t = s, a_t = a\} \tag{8.3}$$

In order for an agent to take into account the long-term reward associated with each action in each state, it must be able to predict the expected long-term reward. For a certain policy π, we calculate the expected return from starting in state s and taking action a as [135]

$$Q^\pi(s,a) = E_\pi \left\{ \sum_{k=0}^{\infty} \gamma^k r_{t+k+1} \middle| s_t = s, a_t = a \right\} \tag{8.4}$$

where γ is known as the discount rate, and represents how strongly future rewards will be taken into account. Equation 8.4 is known as the action-value function, and in a method known as Q-Learning, the action-value function is estimated based on observations. While performing actions in an environment, the learner updates its estimate of $Q(s_t, a_t)$ as follows [135]:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha \left[r_{t+1} + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t) \right] \tag{8.5}$$

where r_{t+1} is the reward received from taking action a, and α is the learning rate, which determines how quickly old information is replaced with new information. Because Q-Learning is an iterative algorithm, it must be programmed with initial conditions ($Q(s_0, a_0)$). Optimistically high values are typically used for initialization, to promote exploration. However, even once some initial exploration is performed, there needs to be a mechanism that prevents the agent from simply sticking to the best policy at any given time. An approach known as Epsilon-greedy forces the agent to take the "best action" with probability $1 - \epsilon$, and take a random action (using a uniform probability) with probability ϵ. Epsilon is usually set at a fairy high value (e.g. 0.95) so that a majority of the time the agent is using what it thinks is the best action. For an in-depth tutorial on MDPs and RL, we refer the reader to [135].

Lastly, as a note to other researchers looking to apply MDP-based RL, it is important to make sure that RL is not applied to systems in which every state can lead to any other state. An example would be the classic cognitive radio application where a radio is tasked with tuning its physical layer parameters to best perform in a given wireless environment. If the state s consists of a tuple containing, for example, {modulation scheme, coding scheme, transmit power}, then all the states are connected to each other, because the radio can choose to use any value of these three parameters at any time (nothing is preventing the radio from switching from QPSK to 16-QAM). As such, modeling the system as a MDP highly overcomplicates things without providing a gain or advantage. While there are published papers within the wireless communications domain that make this very mistake, they will not be specifically called out.

Table 8.1: Four common types of Markov Models

	States are Observable	**States are Partially Observable**
System is Autonomous	Markov Chain	Hidden Markov Model (HMM)
System is Controlled	Markov Decision Process	Partially Observable MDP

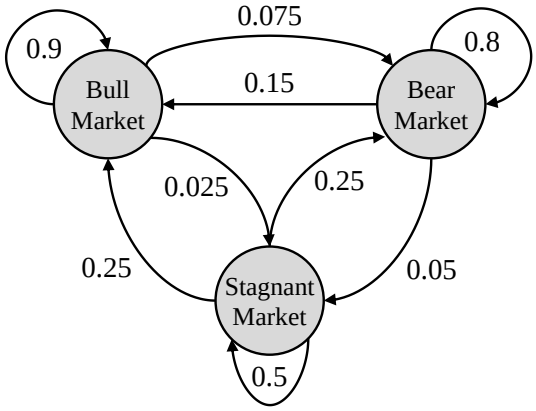

Figure 8.2: Example of a Markov chain, based on the stock market.

8.4.2 Summary of Markov Models

The term Markov occurs often in this field of research, and there is sometimes confusion as to which type of Markov model is being used in a particular situation. This subsection attempts to clear up the major differences between the common Markov models.

A stochastic process has the *Markov property* if the probability distribution of the future states of a process depends only on the present state, not on the sequence of events before the present state. There are four main types of Markov Models, described in the table below (it should be noted that a Markov Random Field or Markov network may be considered to be a generalization of a Markov chain in multiple dimensions).

Markov Chain: A mathematical system that undergoes transitions from one state to another, among a finite or countable number of possible states. The commonly used stock market example is shown in Figure 8.2. In a Markov model diagram, each circle is a state. We assume that the current state is directly observable.

Hidden Markov Model (HMM): Similar to a Markov Chain, but the state is not directly observable. Instead, there are outputs that are dependent on the state, which we observe. Each circle represents a random variable (RV), not a specific state, as shown in Figure 8.3. The top circles are hidden RVs while the bottom ones are observable RVs. HMM are often used to perform inference, which means estimating the hidden states based on observations. The arrows in Figure 8.3 represent conditional dependencies. You can describe a specific

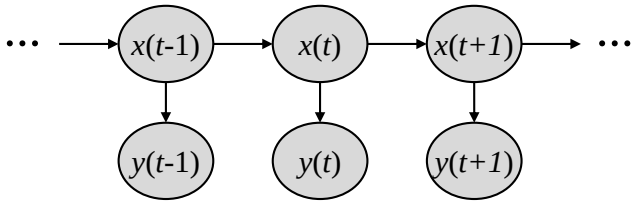

Figure 8.3: Example of a hidden Markov model, in its general form.

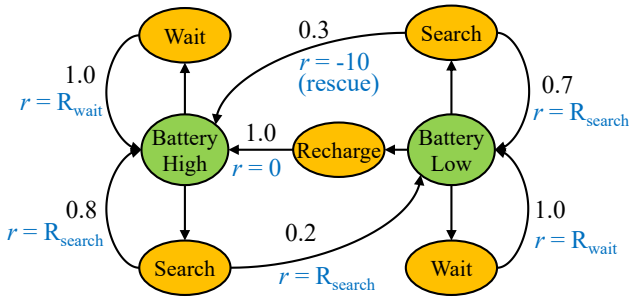

Figure 8.4: Example of a Markov decision process involving a recycling robot that must choose whether it searches for cans or waits for them to be dropped off.

HMM through the transitional probabilities: $y(t)$ based on $x(t)$, $x(t+1)$ based on $x(t)$, and the initial probability, $x(0)$.

Markov Decision Process: A way to model decision making in situations where outcomes are partially random and partially under the control of a decision maker. MDPs play a huge role in RL. At each time step, the process is in some state s, and the decision maker must choose an action a that is available in state s. At the next time step, the process moves randomly into a new state s', giving the decision maker a corresponding reward r which is based solely on s, a, and s' (although it could still be stochastic). An example MDP is shown in Figure 8.4. In this example, a recycling robot decides whether it should search for a can, remain still and wait for someone to bring it a can, or go back to home and recharge. Each green circle represents a state, and there are only two in this example. Each orange circle is an action, and the probabilities show the likelihood of the action leading to each state. R_{search} and R_{wait} are the rewards associated with searching and waiting respectively, and are likely stochastic in nature, with $R_{search} > R_{wait} > 0$. If the robot is in a low battery state, goes to search, and depletes its battery, it must be rescued and recharged, and thus a negative ten reward is assigned to the event.

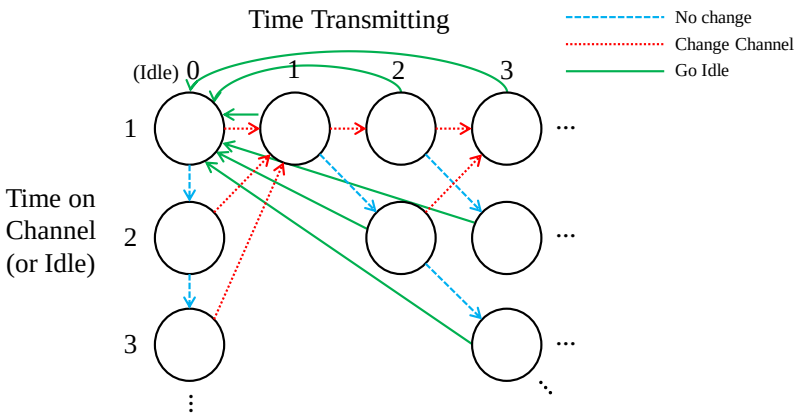

Figure 8.5: Markov Decision Process associated with hopping channels and going idle. These states and actions are independent of the jammer's actions.

8.4.3 Markov Decision Process Formulation

We will now formulate a MDP used to model the transmitter's available states and actions. The states exist on a two dimensional grid, in which one axis corresponds to the time that the transmitter has been on a given channel (including the "idle channel"), and the other axis corresponds to the time the transmitter has been continuously transmitting. Time is discretized into time steps, and we will assume the step size is equal to the smallest period of time in which the transmitter must remain on the same channel. Figure 8.5 shows the state space and actions available in this MDP.

The transmitter will always start in the top-left state, which corresponds to being idle for one time step. It then must choose whether to remain idle, or "change channel" (which can be interpreted as "start transmitting" when coming from an idle state). If it decides to change channel, then in the next time step it must decide whether to remain on that channel or change to a new channel, which we will assume is chosen randomly from a list of candidate channels. It should be noted that the result of each action is deterministic, however the rewards may contain a stochastic component. Due to what the states represent, the MDP is theoretically infinitely long in directions indicated by ellipsis in Figure 8.5. However, in practical systems the width and height of the MDP must be limited, as we discuss later.

The layout of this MDP is different than others that appear in related literature, and the main difference is that the state does not include any specific channels. We believe this is the "special sauce" behind the proposed strategy, because it allows for more rapid learning, and does not "blow up" when the number of channels is increased to 100 or 1000. In approaches that incorporate the specific channel into the state, the radio must build statistics for each channel, one at a time, which can take a while and lead to a very nonadaptive system. On the other hand, generating statistics on each specific channel allows the radio to determine which channels are, in general, more favorable to use. It is for this reason that we recommend

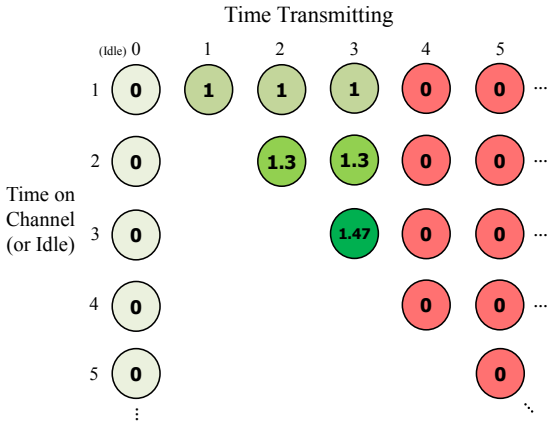

Figure 8.6: Rewards associated with reactive jammer model $N_{REACT} = 3$.

the strategy described in this chapter be used in tandem with a channel selection algorithm, which has the ability to blacklist specific channels (e.g. a channel that has a static barrage jammer jamming it).

The reward associated with each state transition is based on the actual data throughput that occurs during the time step. As such, the rewards are based on the jammer model, and may be stochastic in nature. Figure 8.6 shows the rewards associated with a transmitter and receiver operating in the presence of a reactive jammer with $N_{REACT} = 3$ and $N_{IDLE} = 1$ (model and parameters defined in the previous section). This example shows that when the radio is transmitting for more than three continuous time steps, the link becomes jammed (red states) and the reward becomes zero until the jammer goes idle and then starts transmitting again (the radio is not rewarded while idle). Although the rewards are shown on top of each state, they are actually associated with the previous state and action taken, and won't always be equal for a given resulting state. The numbers 1, 1.3, and 1.47 are examples to demonstrate the fact that remaining on the same channel is more favorable than switching channels, due to the time it takes to switch frequencies. In a real implementation the reward would be based on the achieved throughput or quality of the link; not a model. A summary of how to cast this mitigation approach into a RL framework is given in Table 8.2.

Now that the states, actions, and rewards are established, we can investigate the learning process of the transmitter in the presence of various types of reactive jammers. In RL, the agent (in this case, the transmitter) learns by trying actions and building statistics associated with each state-action pair. At the beginning of the learning process, the agent has no information about the environment, and must try random actions in any state. After a period of learning the agent eventually chooses what it thinks is the best action for each state in terms of the predicted long-term reward. The Epsilon-greedy approach forces the agent to never consider itself "finished" with learning.

Under a reactive jammer with a certain N_{REACT} and when $N_{IDLE} = 1$, the optimal policy is

Table 8.2: Summary of how to cast this mitigation approach into a RL framework

Environment States exist on a two dimensional grid, in which one axis corresponds to the time the transmitter has been on a given channel (including the "idle channel"), and the other axis corresponds to the time the transmitter has been continuously transmitting.
Agent's Actions are to either 1) remain idle or 2) "change channel" which can be interpreted as "start transmitting" when coming from an idle state.
State Transition Rules are deterministic (although a stochastic component due to external factors could be added) and based on the action taken.
Reward Function is a value proportional to the data throughput that occurred during the time step (not known until feedback is received).
Agent's Observations include the state it is currently in, and the reward achieved from each state-action pair.
Exploration vs. Exploitation is achieved using the Epsilon-greedy approach, in which the agent chooses a uniformly random action a small fraction of the time.
Task type is continuing by nature, but could be treated as episodic where each episode is an attempt to transmit for N time steps.

to remain on the same channel for N_{REACT} time steps, and then go idle for one time step. Three optimal policies are shown in Figure 8.7, correspond to $N_{REACT} = 1, 2$, and 3. Each optimal policy resembles a loop that starts at idle for one time step and proceeds to transmit on the same channel for N_{REACT} time steps. In a real-world scenario, it takes the transmitter many time steps to establish that this is the best policy to take, because it must explore each state-action multiple times to build reliable statistics.

The optimal policy for a repeater jammer is shown in Figure 8.8, using $N_{REP} = 1$ and 2. This zigzag pattern indicates constant-rate frequency hopping, which is well-established as the typical method for avoiding repeater jamming [81]. Unfortunately, the optimal policy will always be infinitely long in the horizontal direction. To take this into account, the learning process can involve resetting the current state to the top-left state after a certain period of time continuously transmitting. This will have minimal influence on learning the optimal policy as long as the state space spans enough time steps to take into account the slowest (i.e. the highest value of N_{REP}) perceivable repeater jammer.

Using the approach described in this chapter, there is no need to perform "jammer classification". As such, the mitigation strategy will remain effective over a wide range of jamming behaviors, and may even be able to deal with new jamming behaviors that were not considered during development.

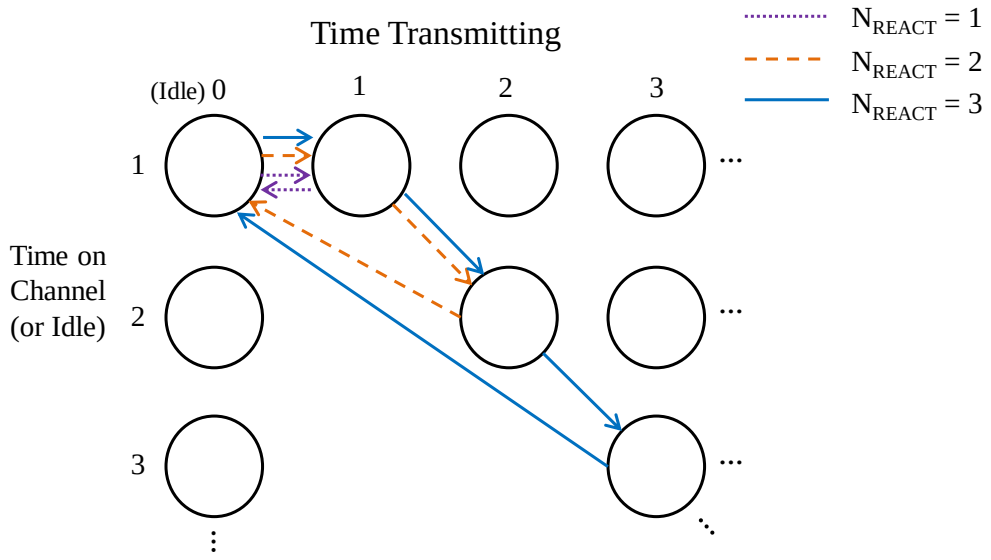

Figure 8.7: Optimal policies in the presence of three different reactive jammers, resembling constant-rate frequency hopping.

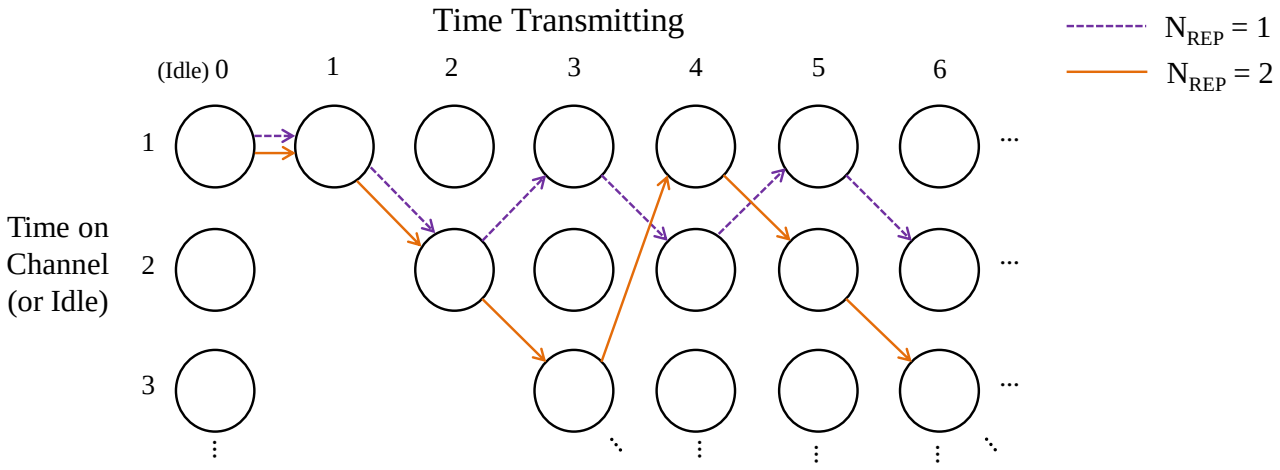

Figure 8.8: Optimal policies in the presence of two different repeater jammers.

8.4.4 Knowledge Decay

The last component of the proposed mitigation strategy is taking into account a changing environment. A given jammer may only be present for a short period of time, and link performance would degrade if the transmitter were sticking to its acquired knowledge. As such, the learning engine must incorporate some form of knowledge decay. Due to the nature of Q-Learning, the learning rate α can be used as a form of knowledge decay, by setting it low enough so that the learner can react to a changing environment. A proper value for α would be based on how quick the transmitter is able to learn optimal policies for a variety of jammers, and how long each jammer remains active of the channel (or rather, how long the radio remains active in the environment). Naturally, these kind of questions are addressed during system design and implementation. A detailed investigation on approaches of knowledge decay/forgetting is beyond the scope of this work, but for more information we refer the reader to [135].

8.4.5 Comparison with Traditional Parameter Optimization

In order to evaluate the usefulness of the proposed approach, we compare it to a much simpler alternative. Finding an effective channel hopping and idling pattern in the presence of a reactive jammer could also be performed by optimizing the hopping rate and transmission duty cycle. This can also be thought of as adjusting T_{ON} and T_{OFF}; the transmission and idle time of a transmitter, assuming it hops frequencies after each transmission. This type of approach is often used in cognitive radio [10]. If $T_{OFF} = 0$, then T_{ON} becomes the hopping rate. Any number of optimization approaches could be used to tune these two parameters. However, even though this simpler approach can take into account the two specific jammer models described in this chapter, it does not have the flexibility inherent to the RL approach. For example, consider the scenario involving a reactive jammer with $N_{REACT} = 4$, $N_{IDLE} = 1$ and a repeater jammer with $N_{REP} = 1$, both targeting the friendly node simultaneously. The optimal transmission strategy would be to hop channels every time step, but also go idle for one time step after the fourth consecutive transmission (a strategy which is likely not possible with traditional parameter optimization). In addition, if the actual jammer behavior experienced by the transmitter does not match any models developed during creation of the mitigation strategy, then added flexibility may mean the difference between communicating and being fully jammed. That being said, we acknowledge that simpler approaches exist for this overall problem, and in many cases they may be better suited to the application at hand.

8.5 Simulation Results

In this section, we present simulation results to show proof of concept of our proposed technique. To simulate this RL based mitigation strategy, a link layer simulation framework

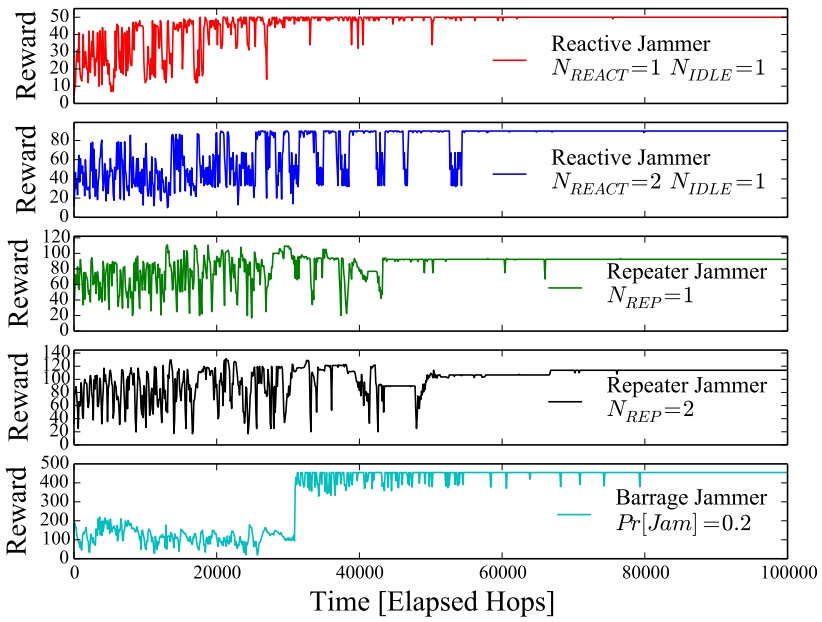

Figure 8.9: Simulation results showing the learning process over time in the presence of different jammers.

was created, which included the jammer models described in this chapter. Q-learning was chosen as the RL technique [135]. In terms of Q-learning parameters, a learning rate, α, of 0.95 (the transmitter will quickly use learned knowledge) and discount factor, γ, of 0.8 was used for the simulations. This relatively low discount factor was used because of the cyclic nature of the optimal policies. Figure 8.9 shows the reward over time for various jamming models, depicting the learning process until saturating to an effective policy with constant reward. Because the reward from each time step is proportional to link throughput, the results can be interpreted as throughput over time. The barrage jammer was modeled by causing jamming with 20% probability at each time step, regardless of how long the transmitter has been transmitting or on a given channel. This can be thought of as a nonreactive jammer that is always transmitting, but at a jammer-to-signal ratio that is not quite high enough to cause complete denial of service. If we were to have modeled the barrage jammer as jamming every transmission successfully, then the results would be of no interest, as the link throughput would simply be near-zero. Note that a partial-band static jammer would have similar results as the barrage jammer, because the simulation framework created here did not include a channel selection algorithm.

The maximum achievable reward under each jamming behavior varies, which is expected (e.g. $N_{REACT} = 2$ will allow using a higher duty cycle than $N_{REACT} = 1$). Although not depicted in Figure 8.9, it should be noted that the transmitter learned the optimal policy (discussed in the previous section) only during the reactive jamming and barrage jamming scenarios. In both repeater jamming scenarios the learned policy did not traverse the entire zigzag pattern

on the MDP, which is the optimal policy for the repeater jamming model as discussed earlier. Rather, the transmitter would go idle on occasion, which would essentially reset the zigzag pattern. Hence, the reward achieved under repeater jamming was not the maximum possible reward. Under barrage jamming the optimal policy for the transmitter would be to remain transmitting on the same channel indefinitely, which occurred after around 50,000 time steps, except for the occasional channel hop (as indicated by the small dips in the plot). This demonstrates how the proposed strategy can work under non-reactive jamming, despite not being designed to do so, and even provide better throughput than a constant-rate FHSS approach by avoiding the overhead associated with changing channels.

It should be noted that the time taken to learn an effective strategy for a given jammer is a function of the learning rate parameter and learning technique (Q-learning in this case). Results in Figure 8.9 show a learning time between 30,000 and 50,000 time steps, which is one or two seconds in a system where the minimum hop duration is on the order of tens of microseconds. While this may seem long compared to a radio that is preprogrammed with specific anti-jam strategies, it is unlikely that the presence of different jammers will change within seconds. In addition, the preprogrammed radio must spend time classifying the type of jammer present in order to know which mitigation scheme to use; a process which is not needed for the proposed strategy. We remind the reader that although wireless channel conditions are known for changing within milliseconds, the proposed strategy is meant to counter the adversary; not traditional channel imperfections such as fading or Doppler shift.

8.6 Conclusions

In this chapter, we have developed a RL based strategy that a communication system can use to deal with reactive jammers of varying behavior by learning an effective channel hopping and idling pattern. Simulation results provide a proof of concept and show that a high-reward strategy can be established within a reasonable period of time (the exact time being dependent on the duration of a time step).

This approach can deal with a wide range of jamming behaviors, not known a priori. Without needing to be preprogrammed with anti-jam strategies for a list of jammers, our approach is able to adapt to the harsh wireless environment. The proposed technique is best used in tandem with an algorithm that finds a favorable subset of channels to use, as well as modern optimization techniques such as adaptive modulation and forward error correction. In future work we will investigate expanding the MDP state space to take into account additional factors, as well as explore more stochastic jammer models. We will also investigate how to deal with reactive jammers that have much longer reaction times relative to the time step interval (e.g., $N_{REACT} = 20$), which would take a long time to learn using the proposed strategy. Lastly, it is likely that the RL procedure can be tuned to provide even greater performance.

Chapter 9

The OFDM Reactive Jammer

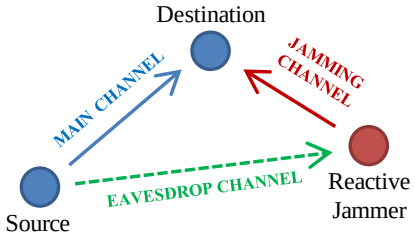

Figure 9.1: The network diagram of a reactive jamming attack, showing the three channels involved.

9.1 Introduction and Problem Formulation

In this chapter we investigate a reactive jammer that senses a large block of spectrum, and transmits noise on any subchannel that appears to have energy on it. By using a system similar to an orthogonal frequency-division multiplexing (OFDM) transceiver, a simple and efficient reactive jammer design is realized. We assume the jammer does not know anything about the target signal(s), other than what block of spectrum they may appear in, and thus uses an energy detector to determine which subchannels have a signal present. Jamming is performed using an OFDM signal that has random symbols assigned to the subchannels with a signal present, and zero assigned to the rest. After describing the design of the jammer, we investigate the impact of energy leakage into adjacent subcarriers that do not overlap with the target signal. Lastly, we compare the OFDM reactive jammer to alternative jamming strategies.

Reactive jamming, also known as time-correlated jamming and follower jamming, is a broad category of jammers that react to sensed energy or signals. Reactive jammers must have receiving capability, and a common form of reactive jamming is repeater jamming (a.k.a. digital RF memory or DRFM jamming) in which the jammer retransmits the signal it receives with a possible transformation applied [3]. A reactive jamming scenario involves three wireless channels, as shown in Figure 9.1. Of the more complex types of jammers, reactive jammers are becoming more prevalent due to low-cost software-defined radios. In this chapter we investigate the design of a reactive jammer that is based on OFDM concepts, and requires no knowledge about the target signals. Figure 9.2 shows the system diagram of the proposed OFDM reactive jammer, split into three parts (receiver, signal detection, transmitter) that correspond to the next three sections of the chapter.

This chapter is organized as follows. Sections 9.2, 9.3, and 9.4 describe the three components of the reactive jammer: the receiver, the detection process, and the transmitter. Section 9.5 provides numerical results showing the performance of the OFDM reactive jammer compared to alternative jamming strategies. Section 9.6 concludes and discusses future work.

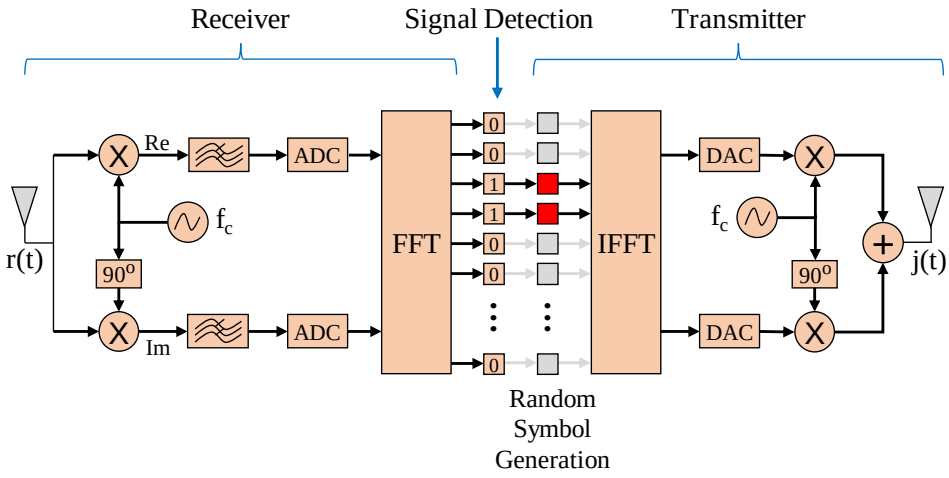

Figure 9.2: The system diagram of the OFDM reactive jammer.

9.2 Jammer's Receiver

Typically, the first step in sense-and-transmit type reactive jamming is to take the received signal and determine if (or how many) signals are present and on what frequencies. We will do this indirectly by using an N-point fast Fourier transform (FFT) to split up the received signal into a series of subchannels, in the same way that power spectral density is typically calculated. Note that this is not the same as using a channelizer, because when using an FFT, energy on one subchannel leaks into the adjacent subchannels. This is not a huge problem because the signals being detected will likely span multiple subchannels (by design). Instead of symbol demodulation, signal detection is performed, which will be discussed in the next section. Using this approach, the jammer's receiver resembles an OFDM receiver, except there is no notion of time and frequency synchronization/offset.

The subcarrier spacing is based on the most narrow (in bandwidth) anticipated signals to be detected (e.g., 10 kHz). This will also determine the total OFDM symbol length, or in this application, the observation interval. Given a certain subcarrier spacing, N will determine how wide a bandwidth is covered by the jammer, as well as how many samples are needed for each observation (limited by the speed of the analog-to-digital converter).

In terms of the jammer covering a large portion of spectrum, the jammer can use multiple receive RF chains, with a single N-point FFT per RF chain. Using this approach, a single reactive jammer can cover nearly all spectrum used for modern communications, and conserve energy by only jamming signals when they are present.

In terms of notation, let $\boldsymbol{r}[n]$ (dimensions Nx1) be the received signal at the jammer after sampling, and $\boldsymbol{R}[n]$ (dimensions Nx1) be the N-point FFT of $\boldsymbol{r}[n]$.

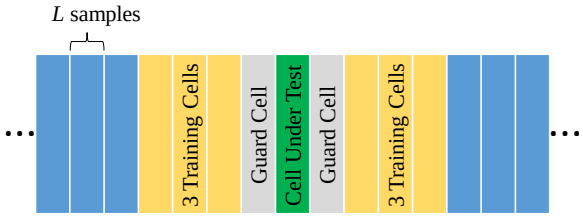

Figure 9.3: Cell relation using the Constant False Alarm Rate (CFAR) technique

9.3 Jammer's Detection Process

9.3.1 Detection Process

In each subchannel k we have a binary signal detection problem, where we would like to determine whether or not a signal is present. The two hypotheses are therefore:

$$H_0 : r_k = n \qquad H_1 : r_k = s + n \tag{9.1}$$

Because we are assuming the jammer does not know anything about the signal (i.e., blind sensing), it has to perform non-coherent detection and the optimal detection method is an energy detector [136]. The detection occurs in the frequency domain, after the signal has gone through a FFT. We will model this process as the detection of an unknown signal in Gaussian noise, and the detection will be performed across N subchannels/subcarriers simultaneously and independently. In this application, false alarms do not have serious consequences, they simply lead to a small amount of jamming energy being wasted from the perspective of the jammer. We will therefore set the false alarm rates fairly high (e.g., 1%). In terms of the noise, we will assume zero-mean Gaussian noise that is white over the bandwidth of each subchannel, but may be colored over the entire FFT bandwidth.

A basic energy detector works as follows. For each subcarrier k, samples are split into blocks of L samples (i.e., L is the number of OFDM symbols needed). For each block, a test statistic, denoted as Λ_k, is calculated:

$$\Lambda_k = \frac{1}{L} \sum_{n=1}^{L} |R_k[n]|^2 \tag{9.2}$$

where $R_k[n]$ is the n'th sample of the k'th subcarrier in the current block. Parseval's theorem tells us that this test statistic is equivalent to the time-domain version, $\Lambda = \sum |r_k[n]|^2$ [137]. If L is large, than the Central Limit Theorem can be applied to approximate the test statistic's distribution as a Gaussian random variable with mean $P_s + \sigma_n^2$ and variance $(P_s + \sigma_n^2)^2/L$ where P_s is the signal power.

The test statistic is compared to a threshold to determine whether to select H_0 or H_1. The value of the threshold is a function of the noise power and the desired false-alarm rate. The

detection problem revolves around maximizing the probability of detection while keeping the probability of false alarm (a.k.a. false alarm rate) below a reasonable level. The problem with this basic method is that it assumes you know the noise power. To get around the problem of not knowing the noise variance and the possibility of colored noise, a class of detection techniques called Constant False Alarm Rate (CFAR) were created. CFAR techniques are heavily used in radar. The basic principle is that the noise variance is estimated using samples that are neighboring to those under test. For example, when calculating Λ_k for a block of L samples, the noise variance can be estimated from the three blocks before and three blocks after the one being tested (with one block as a guard cell on both sides), as depicted in Figure 9.3. Often these blocks of samples are referred to as cells. Using this technique, the decision threshold, T, is given by $T = \alpha\sigma_n^2$ [138]. Parameter α is a scaling factor and σ_n^2 is estimated from samples taken from neighboring blocks:

$$\sigma_n^2 = \frac{1}{L}\sum_{n=1}^{L} |R_k[n]|^2 \tag{9.3}$$

The scaling factor α is tuned based on the desired false alarm rate, denoted as P_{fa}, as follows [138]:

$$\alpha = L(P_{fa}^{-1/L} - 1) \tag{9.4}$$

Using this new detection statistic and decision threshold, the jammer can determine whether or not a signal is present on each subcarrier.

In [137] it is shown that the number of samples needed to achieve a certain probability of detection, P_d, and false alarm rate, P_{fa}, is independent of the threshold, and only a function of the signal-to-noise ratio (SNR), denoted as γ. This can be approximated as:

$$L \approx \left[Q^{-1}(P_{fa}) - Q^{-1}(P_d)\sqrt{2\gamma + 1}\right]^2 \gamma^{-2} \tag{9.5}$$

Equation 9.5 can be used to determine how large L should be during the design process.

Note that this frequency-domain detection approach is not the same as a "channelized-radiometer". A channelized-radiometer uses a channelizer, which has dedicated band-pass filters for each channel, instead of a simple FFT, so that there is no (or very little) overlap between channels. Meanwhile, the FFT approach causes significant overlap, due to the shape of the sinc function.

9.3.2 Numerical Results

We will now investigate how well the OFDM reactive jammer can receive, detect, and jam signals. As mentioned before, the jammer knows nothing about the target signals. However, it cannot detect and jam signals below the noise floor, such as a direct-sequence spread spectrum signal of low spectral density. We will assume the bandwidth of the target signal

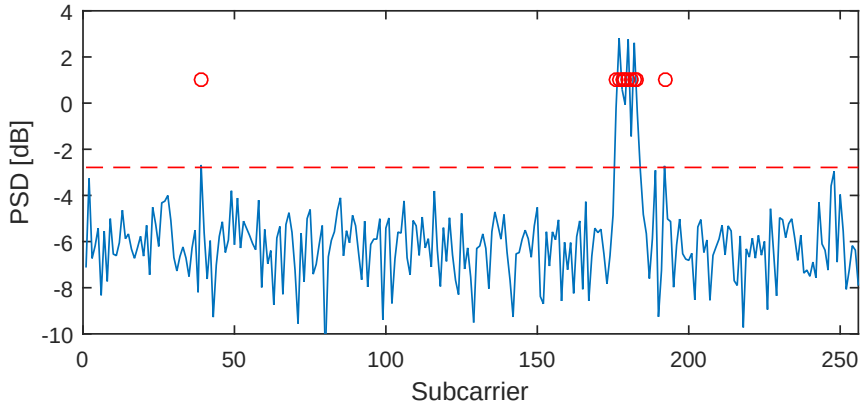

Figure 9.4: Example of the detection process, showing PSD of the single-carrier signal, as well as which subcarriers registered positive detection under an SNR of 6 dB.

is greater than or equal to the subcarrier spacing. We will also assume that the signals are within the jammer's receive frequency range. To test how well the OFDM reactive jammer can receive, detect, and jam signals, the full jammer has been implemented in MATLAB.

The specific signal that has been simulated for the purpose of experimentation is a single-carrier binary phase-shift keying (BPSK) signal that uses a root-raised cosine filter of roll-off factor 0.3. This represents a modern digital signal that can have varying bandwidth. The signal bandwidth used is varied between what is equivalent to one subcarrier and 16 subcarriers, in terms of the jammer's receiver. This provides insight regarding how easy it is to detect signals that are only covered by a small number of subcarriers. It is important to note that the signals being detected are not synchronized in any way to the jammer's receiver in time or frequency.

Figure 9.4 shows the OFDM reactive jammer detecting and jamming a single-carrier signal that occupies roughly 8 subcarriers, under a SNR of 6 dB. The red dashed line represents the decision threshold associated with a false alarm rate of 1%, and the red dots show which subcarriers had a positive detection, and were subsequently jammed. All subcarriers overlapping with the single-carrier signal triggered detection, although two extra subcarriers had a false detection, which makes sense because there are 248 subcarriers that have no signal present which equates to about a 1% false alarm rate. In this example, a 256-point FFT is used, and 2560 samples are taken, representing 10 power spectral density (PSD) observations being averaged together before making a detection decision (i.e., $N = 256$ and $L = 10$).

To investigate the rate of correct detection, we have varied SNR between zero and 10 dB, and used a Monte Carlo type simulation to find the probability of detection. Figure 9.5 shows the results, where each curve corresponds to a signal that covers between one and 16 subcarriers. As before, the false alarm rate for each curve is a constant 1%.

When the signal spans at least two subcarriers, it only takes about 6 dB of SNR to consistently

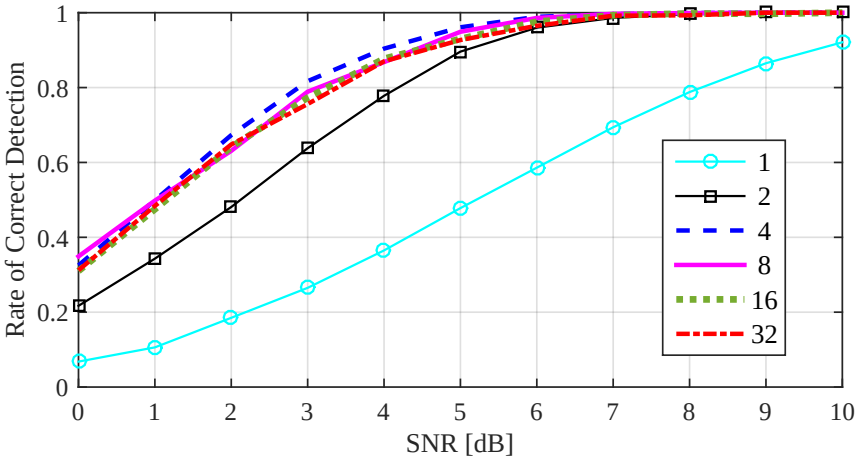

Figure 9.5: Rate of Detection of a single-carrier signal of varying bandwidth (the figure is parameterized by the number of subcarriers the single-carrier signal overlaps with). The signal is modulated with BPSK and passed through a root-raised cosine filter. A constant false alarm rate of 1% is used.

detect and jam it. However, the signal that is only one subcarrier wide in bandwidth is fairly hard to detect every time. However, even detecting a signal half the time corresponds to jamming it with a duty cycle of 50%, where the jamming signal has a period roughly equal to L (2560 samples in this case). In most situations this would likely cause Denial of Service (DoS).

9.4 Jammer's OFDM Transmitter

The third piece of the reactive jammer, the transmitter, is the piece most similar to an OFDM radio. For each subchannel that a signal was detected on, the jammer assigns a random complex variable as the symbol associated with that subcarrier. The jammer assigns a zero to the remaining subcarriers, and then performs an IFFT, generating the OFDM signal used for jamming. Under this strategy, the jammer places energy on top of the detected signals with minimal wasted energy.

9.4.1 Signal Model

The jammer's transmitted signal, $j(t)$, can be represented as a typical OFDM signal,

$$j(t) = \sum_{k=-N/2}^{N/2-1} S_k e^{j2\pi k \Delta f t} \tag{9.6}$$

where N is the number of subcarriers, S_k is the complex symbol carrying data on the k'th subcarrier, and Δf is the subcarrier spacing.

9.4.2 Roll-off Factor of OFDM

Because the jammer's goal is to inject noise on top of the target signals, it is worthwhile to investigate the roll-off factor of an OFDM signal, in terms of how much energy is leaked into subcarriers/subchannels that are not intended to be jammed (i.e., how big are the spectral sidelobes). The average sidelobe power, denoted as A_p, for an OFDM system without windowing is given by a sum of sinc functions [139]:

$$A_p = \frac{1}{K} \sum_{k \in \boldsymbol{N_s}} \left| \sum_{n \in \boldsymbol{N_a}} S_n sinc(\pi(y_k - x_n)) \right|^2 \tag{9.7}$$

where $\boldsymbol{N_a}$ are the active subcarrier indexes and $\boldsymbol{N_s}$ are the K subcarriers associated with the sidelobe being investigated. x_n and y_k corresponds to the normalized subcarrier frequency associated with subcarrier n and k respectively. Unfortunately this expression is dependent on the specific values of the symbols, and there are several possible constellations that can be used for jamming, as discussed in Section 9.4.3. We therefore use simulation to find the average roll-off when using complex Gaussian random symbols.

Figure 9.6 shows the power spectral density of a 256 subcarrier OFDM signal that has groups of 1, 5, 10, and 15 active adjacent subcarriers. The groups are spaced such that there is about 30 dB of attenuation between each one. This represents jamming several spaced out signals with varying bandwidths, using a single OFDM signal. It is clear that there is energy located in the several subcarriers around the ones that are active, effectively causing energy leakage into subchannels that don't have to be jammed. This roll-off is not necessarily a big problem for the jammer, it just represents a small amount of energy that is being wasted. However, in some circumstances, such as when there are a lot of friendly communications signals within the band being jammed, energy leakage into adjacent subcarriers could cause unintended self-jamming. Sidelobe suppression is also an important factor for the same reason, as there could be friendly communications signals adjacent to the band being jammed.

Figure 9.7 shows the fraction of energy that is within the active subcarriers (representing an efficiency-type metric of the jammer), for a varying number of adjacent active subcarriers within a 256 subcarrier OFDM signal. This equates to $1 - A_p/P_s$ where P_s is the OFDM signal's power. For example, if only five adjacent subcarriers are active, in order to jam a signal that spans five subcarriers, then roughly 83% of the energy transmitted by the jammer actually jams the signal, while the rest is wasted.

Note that the roll-off can be improved through raised cosine windowing, but instead of losing out on throughput like in an OFDM radio, the jammer effectively has a duty-cycle where it

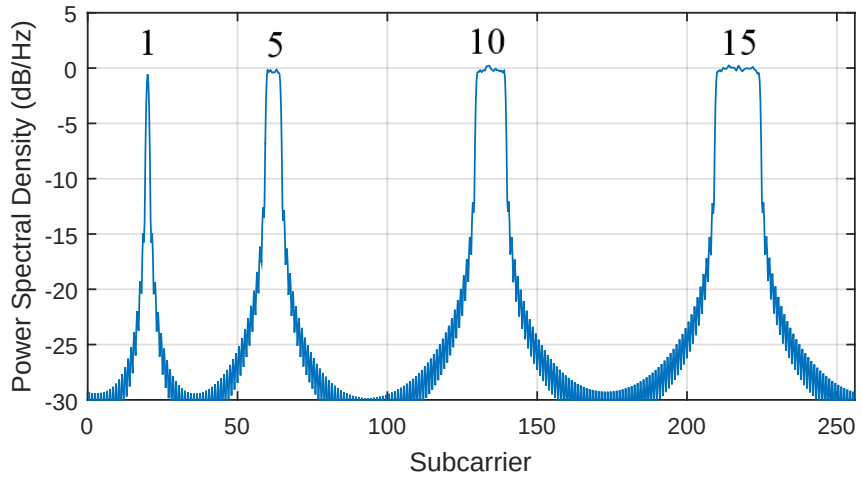

Figure 9.6: PSD of an OFDM signal showing groups of 1, 5, 10, and 15 active adjacent subcarriers.

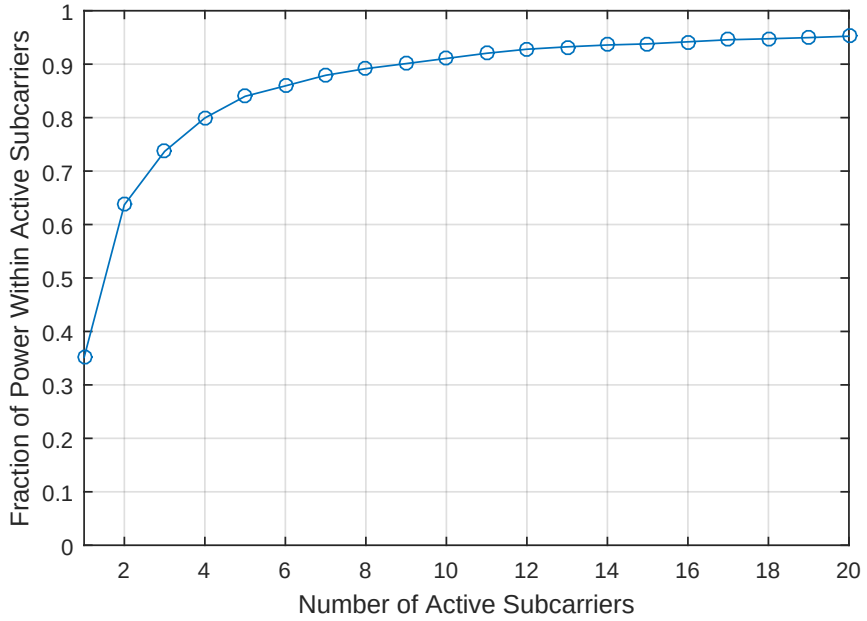

Figure 9.7: Fraction of power within the active subcarriers, for a given number of active subcarriers.

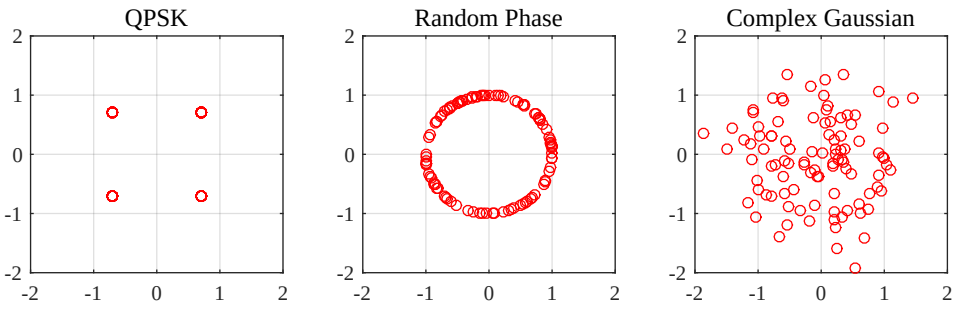

Figure 9.8: Three possible modulation schemes that can be used to create noisy symbols, as part of a jamming waveform, all with unit variance.

is off for short periods of time, which may or may not be desirable. Additional techniques for reducing the sidelobe size were developed within the context of cognitive radio, including the use of "cancellation carriers" [140], subcarrier weighting [141], and a strategy called "adaptive symbol transition" [142]. However, these require high computational complexity.

9.4.3 Transmitted Symbols

The jammer must change symbols periodically, else the jammer effectively transmits a continuous sinusoid (or multiple sinusoids) that the target receivers can more easily deal with. Thus, the jammer must generate a complex symbol, similar to a data symbol, for each subcarrier assigned to jam. Given that the jammer knows nothing about the target signals, there are three possible strategies that make sense when it comes to generating the noise-like symbol on each subcarrier. These three schemes are:

1. A random quadrature phase-shift keying (QPSK) symbol
2. A symbol with amplitude of 1 and phase chosen uniformly random between 0 and 2π
3. A complex Gaussian random variable

QPSK symbols and random phase symbols cause every symbol to have the same power, while complex Gaussian symbols do not. In terms of which scheme provides better jamming, it has been shown that there is very little differences between these strategies [61]. Figure 9.8 shows the constellation of 100 symbols under each scheme.

9.4.4 Peak-to-Average Power Ratio

The peak-to-average power ratio (PAPR) of OFDM is notoriously bad, requiring expensive power amplifiers that have a wide linear region relative to the average transmit power, among other complications. Because we are not transmitting data, we don't care about maintaining

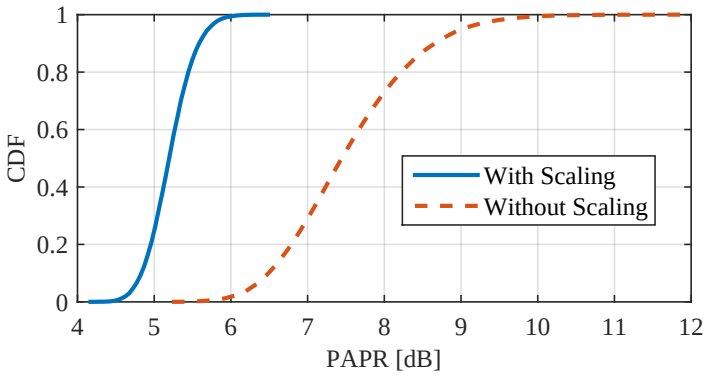

Figure 9.9: Cumulative distribution function of PAPR, before and after the envelope scaling is applied.

the orthogonality between subcarriers. We do, however, care about reducing emissions that are out of band, so as to avoid jamming our own communications. It therefore makes sense to employ simple PAPR reduction, so that it is easier to stay within the linear region of the jammer's power amplifier. While there are possibly hundreds of PAPR reduction techniques in literature, we chose to use a simple yet effective one, envelope scaling, because this OFDM-based reactive jamming approach emphasizes simplicity. To perform envelope scaling we take any points that are above a certain threshold (after the IFFT), and scale them down using a preset factor, usually based on the peak value for each OFDM symbol.

To investigate the benefit of PAPR reduction, we have simulated an example OFDM reactive jammer. It is configured to have 256 subcarriers (the bandwidth is not a factor), and for an example scenario we have placed two signals within the band being jammed. One signal spans ten subcarriers of the jammer, and the other spans 50 subcarriers. We assume that the jammer's detection process is without error, to reduce the simulation framework.

Figure 9.9 shows the cumulative distribution function (CDF) of PAPR, before and after the envelope scaling is applied. We can see that the average PAPR has been reduced by about 2.5 dB, and the variance is significantly reduced, which is just as important as the average. Figure 9.10 shows the PSD of the transmitted jamming signal, with and without envelope scaling. In this figure it is clear which subcarriers are active, corresponding to the location of the two signals being jammed. The "jamming efficiency", which we define as the ratio of jamming energy within the subcarriers being jammed over the total energy, dropped from 100% to 75% when envelope scaling was added. However, this energy is spread fairly evenly across the other subcarriers. If there were more signals to jam within these 256 subcarriers, then this drop in efficiency would be less severe.

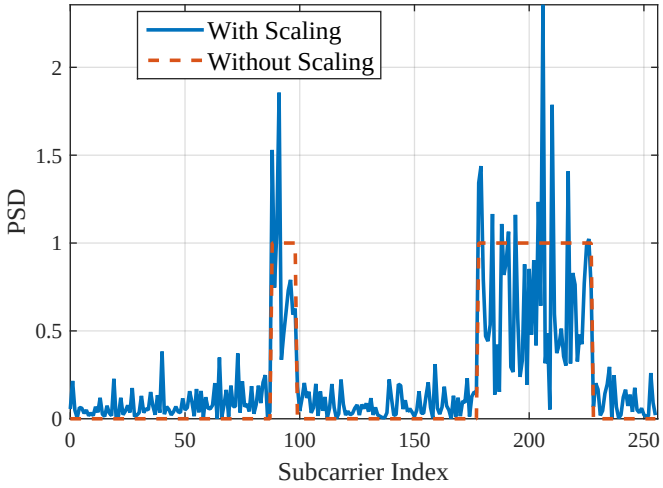

Figure 9.10: Power spectral density of the jamming signal, before and after the envelope scaling is applied.

9.4.5 Alternative to an OFDM-based Transmitter

If PAPR is a huge issue, then an OFDM-based transmitter may not be the best choice for a reactive jammer. One alternative is to simply create a single-carrier signal for each signal to be jammed. Unfortunately this complicates the detection engine, because it must distinguish boundaries of each signal present, instead of just deciding whether each subcarrier should be active or not. In the scenario created as part of the previous subsection, there were two signals, so the reactive jammer would simply generate two single-carrier signals with one of the modulation schemes discussed in Section 9.4.3. It would then add the two signals together and transmit them as the jamming signal. The PAPR of this alternative approach is heavily based on the number of signals present, the width of each signal, and how often the random symbol changes (i.e. the "symbol rate" of the jammer). If there are a large number of narrow-band signals, then the resulting jamming signal is a sum of many sinusoids, resembling OFDM and carrying the same high PAPR issue has before.

9.5 Performance Comparison

We will now compare the OFDM reactive jammer to similar jammers that may be used, including reactive jamming using a single-carrier (SC) signal, which assumes knowledge of the carrier frequency and bandwidth of the target signal, as well as barrage jamming that transmits white Gaussian noise across the entire band. The SC reactive jammer is able to almost perfectly overlap with the target signal in both time and frequency, using its knowledge of the target signal. Bit error rate (BER) over jammer-to-signal ratio (JSR) is the primary metric used for comparison, as it takes into account the inefficiencies of the jammers. Both a

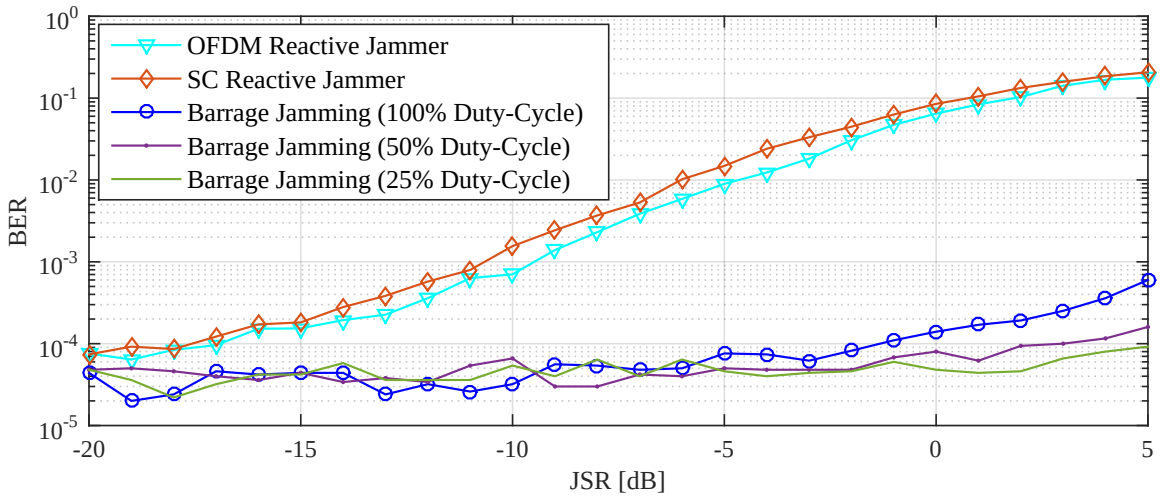

Figure 9.11: Simulation of the OFDM Reactive Jammer, a single-carrier reactive jammer that has knowledge of the target signal's carrier frequency and bandwidth, and barrage jamming. Barrage jamming involves transmitting noise on top of all 256 subcarriers, using a 100% duty cycle, leading to an enormous performance loss. This simulation shows a SNR at the jammer of 6 dB and a SNR at the destination node of 10 dB.

continuous signal and an intermittent signal will be used as the target signal, to show off the gains associated with reactive jamming.

Figure 9.11 shows the simulation when the SNR at the jammer is 6 dB and the SNR at the destination node (not taking into account the jamming signal) is 10 dB. The signal being jammed is roughly four subcarriers wide, which means only about 2% of the barrage jammer's energy is overlapping with the signal in frequency. Three experiments were done, using a target signal with a duty-cycle of 100%, 50%, and 25%. However, because the reactive jammers only jam when they detect the signal, the barrage jammer is the only one that has a significant performance loss due to the non-100% duty-cycle. Thus, five curves are shown.

The two reactive jammer curves resemble a BER curve flipped horizontally, except the curves "bottoms out" at the theoretical BER for BPSK with 10 dB of SNR (i.e., when the jamming signal is negligible). The barrage jammer is expected to perform poorly, considering it is wasting most of its energy, but it was included for the sake of demonstration.

Although we show that reactive jamming can provide enormous gains compared to non-reactive jamming, it is worth noting that reactive jamming will always involve a slight delay before the target signal is jammed (from the perspective of the destination node receiving the signal). If the jammer is located too far away from the target radios, or it is too slow to start jamming (relative to the packet length of the target signal), then it may fail to successfully jam. A solution that combines the best of both worlds is a reactive jammer that places non-reactive (constant) interference on subcarriers that are found to have high activity over a long period of time.

Lastly, we note that this OFDM reactive jammer is especially well-suited for jamming the uplink of LTE, due to the high number of users that each get a block of resources in time and frequency. Using this approach, only the nearby active users are jammed.

9.6 Conclusion and Future Work

In this chapter, we have described a method for reactive jamming inspired by an OFDM transceiver. While the jammer design is fairly straightforward, we have focused on specific issues that arise when using such an approach, such as the poor roll-off of OFDM, and the observation interval needed for accurate signal detection. Simulation results were used to gain further insight into these issues. Lastly, performance of the OFDM jammer was compared to alternative approaches, and numerical results showing BER over JSR were provided.

Further work includes investigating the use of the all-phase FFT (ap-FFT) for the receiver instead of a regular FFT. In addition, Eigenvalue-based blind sensing may provide better detection than using an energy detector. There are plenty of different aspects that could be investigated, regarding the jammer itself.

Chapter 10

Conclusion

As the sophistication of communications systems increases, sophisticated jamming will likely become a bigger threat in public safety, military, and other mission-critical domains. The work presented in Chapter 3 introduced a taxonomy that frames the organization of jammer classes by what information they possess and their capacity to act on that information. This new view of jammers emerges naturally from the way present day wireless technology relies so extensively on software-driven behavior. In addition, understanding the key capabilities that distinguish major classes of jamming, as well as the multidimensional parameter space, can aid in the correct application of anti-jam and detection strategies. Further research regarding the jammer taxonomy includes the design of a radar jamming taxonomy and radio navigation jamming taxonomy. It may be possible to formulate a taxonomy that applies to all forms of jamming.

In Chapters 4, 5, and 6, we introduced a new research area within wireless communications, based on the antifragile paradigm. It has provided a foundation for this research, specific to the electronic warfare domain. By performing jammer exploitation, instead of jammer mitigation, we are able to introduce new tools to the electronic warfare toolbox, and show that it is theoretically possible to achieve a communications gain as a result of a jamming attack. On the one hand, it seems silly to develop antifragile electronic warfare techniques when state-of-the-art anti-jamming is not perfected, but on the other hand, antifragile electronic warfare provides new techniques that could be used alongside anti-jam techniques. Through the antifragile lens, we are able to see the problem under a new perspective, leading to new approaches to an old (yet evolving) threat. An antifragile communication system is something has not yet been realized. However, this dissertation quantified potential antifragile gains in the face of an enemy jammer, in order to motivate its implementation at some future date. Future research within this area should focus on one thing: implementation of antifragile radios, tested with a real reactive jammer.

To tie the analysis back to a real-world electronic warfare application, we analyzed the feasibility of performing and mitigating reactive jamming in a satellite communication (SATCOM) type scenario. Even though reactive jamming is a complex form of jamming that requires receiving capabilities and a low processing delay, it allows a jammer to counter the processing gain associated with frequency hopping spread spectrum. It is for this reason that we should consider reactive jamming a future threat in the SATCOM domain, especially as software-defined radios become more capable and less costly. In our feasibility analysis, we showed that uplink and downlink jamming will remain highly dependent on the scenario at hand, but in many cases are vulnerable. Because both the uplink and downlink are vital to the operation of a communications link, only one of the two needs to be jammed. Thus, the overall conclusion of this chapter was that reactive jamming is in fact feasible in a SATCOM scenario, and reactive-specific countermeasures should be considered.

In Chapter 8, we introduced a strategy to mitigate or even avoid reactive forms of jamming using a reinforcement learning (RL) approach, when jammer exploitation is not practical. Simulation results provided a proof of concept and showed that an effective strategy can be established within a reasonable period of time. This approach can deal with a wide range

of jamming behaviors, not known a priori. Without needing to be preprogrammed with anti-jam strategies for a list of jammers, our approach is able to adapt to the harsh wireless environment. Future work involves implementation of this jamming mitigation approach.

Lastly, to get a feel for what a real implementation of a reactive jammer would look like, we investigated a reactive jammer implementation based on OFDM concepts. This jammer design is based on sensing a large block of spectrum, and transmitting noise on any subchannel that appears to have energy on it. By using a system similar to an OFDM transceiver, a simple and efficient reactive jammer design is realized. Simulation results were used to gain further insight into these issues.

As mentioned in Section 2.2, we recognize that the material in this dissertation is certainly not a perfect embodiment of antifragility. Antifragility within engineering systems is a great challenge, as engineering systems by definition will not benefit from disorder in general. For this research, we have simply used the concept of antifragility as inspiration for the techniques developed during the course of this research, and as a way to view the problem of anti-jamming from a different perspective. We hope that this work can pave the way for further research into antifragile communications.

Bibliography

[1] M. Lichtman, M. T. Vondal, T. C. Clancy, and J. H. Reed, "Antifragile Communications," *IEEE Systems Journal*, vol. PP, no. 99, pp. 1–12, 2016.

[2] M. Lichtman, R. P. Jover, M. Labib, R. Rao, V. Marojevic, and J. H. Reed, "LTE/LTE-A jamming, spoofing, and sniffing: threat assessment and mitigation," *IEEE Communications Magazine*, vol. 54, no. 4, pp. 54–61, April 2016.

[3] M. Lichtman, J. D. Poston, S. Amuru, C. Shahriar, T. C. Clancy, R. M. Buehrer, and J. H. Reed, "A Communications Jamming Taxonomy," *IEEE Security and Privacy*, February 2016.

[4] M. Lichtman and J. H. Reed, "Anomaly-based intrusion detection of protocol-aware jamming," in *IEEE Military Communications Conference (MILCOM)*, Oct 2015, pp. 269–274.

[5] ——, "Analysis of reactive jamming against satellite communications," *International Journal of Satellite Communications and Networking*, vol. 34, no. 2, pp. 195–210, 2016.

[6] M. Lichtman, T. Czauski, S.-K. Ha, P. David, and J. H. Reed, "Detection and Mitigation of Uplink Control Channel Jamming in LTE," in *IEEE Military Communications Conference (MILCOM)*, 2014, pp. 1187–1194.

[7] J. Kakar, K. McDermott, V. Garg, M. Lichtman, V. Marojevic, and J. H. Reed, "Analysis and Mitigation of Interference to the LTE Physical Control Format Indicator Channel," in *IEEE Military Communications Conference (MILCOM)*, 2014, pp. 228–234.

[8] M. Lichtman and J. H. Reed, "Reinforcement Learning for Reactive Jamming Mitigation," *Journal of Cyber Security and Mobility*, vol. 3, no. 2, pp. 213–229, 2014.

[9] C. Shahriar, M. L. Pan, M. Lichtman, T. C. Clancy, R. McGwier, R. Tandon, S. Sodagari, and J. H. Reed, "PHY-Layer Resiliency in OFDM Communications: A Tutorial," *IEEE Communications Surveys Tutorials*, vol. 17, no. 1, pp. 292–314, January 2015.

[10] S. Dudley, W. Headley, M. Lichtman, E. Imana, X. Ma, M. Abdelbar, A. Padaki, A. Ullah, M. Sohul, T. Yang, and J. Reed, "Practical Issues for Spectrum Management With Cognitive Radios," *Proceedings of the IEEE*, vol. 102, no. 3, pp. 242–264, March 2014.

[11] M. Lichtman, J. Reed, T. Clancy, and M. Norton, "Vulnerability of LTE to hostile interference," in *IEEE GlobalSIP*, Dec 2013, pp. 285–288.

[12] M. Lichtman, W. C. Headley, and J. H. Reed, "Automatic Modulation Classification under IQ Imbalance Using Supervised Learning," in *IEEE Military Communications Conference (MILCOM)*, Nov 2013, pp. 1622–1627.

[13] T. C. Clancy, M. Norton, and M. Lichtman, "Security Challenges with LTE-Advanced Systems and Military Spectrum," in *IEEE Military Communications Conference (MILCOM)*, Nov 2013, pp. 375–381.

[14] M. J. L. Pan, M. Lichtman, T. C. Clancy, and R. W. McGwier, "Protecting physical layer synchronization: mitigating attacks against OFDM acquisition," in *Wireless Personal Multimedia Communications (WPMC)*, June 2013, pp. 1–6.

[15] D. L. Adamy, *EW 101*. London: Artech House, 2001.

[16] J. A. Boyd, D. B. Harris, and D. King, *Electronic Countermeasures*. Peninsula Pub, 1978, vol. 1.

[17] N. N. Taleb, *Antifragile: Things that gain from disorder*. Random House LLC, 2012.

[18] R. W. Lucky, "Antifragile Systems [Reflections]," *IEEE Spectrum*, vol. 50, no. 3, pp. 28–28, 2013.

[19] V. De Florio, "Antifragility = Elasticity + Resilience + Machine Learning Models and Algorithms for Open System Fidelity," *Procedia Computer Science*, vol. 32, pp. 834–841, 2014.

[20] A. Danchin, P. M. Binder, and S. Noria, "Antifragility and tinkering in biology (and in business) flexibility provides an efficient epigenetic way to manage risk," *Genes*, vol. 2, no. 4, pp. 998–1016, 2011.

[21] E. Verhulsta, "Applying systems and safety engineering principles for antifragility," *Procedia Computer Science*, vol. 32, pp. 842–849, 2014.

[22] K. H. Jones, "Engineering Antifragile Systems: A Change In Design Philosophy," *Procedia Computer Science*, vol. 32, pp. 870–875, 2014.

[23] A. Tseitlin, "The antifragile organization," *Communications of the ACM*, vol. 56, no. 8, pp. 40–44, 2013.

[24] A. Abid, M. T. Khemakhem, S. Marzouk, M. B. Jemaa, T. Monteil, and K. Drira, "Toward antifragile cloud computing infrastructures," *Procedia Computer Science*, vol. 32, pp. 850–855, 2014.

[25] T. E. Marler, "Promoting the confluence of tropical cyclone research," *Communicative & integrative biology*, vol. 8, no. 2, p. e1017165, 2015.

[26] K. J. Hole, "Toward Anti-fragility: A Malware-Halting Technique," *IEEE Security & Privacy*, no. 4, pp. 40–46, 2015.

[27] C. A. Ramirez and M. Itoh, "An initial approach towards the implementation of human error identification services for antifragile systems," in *SICE Annual Conference*. IEEE, 2014, pp. 2031–2036.

[28] M. S. Asif, M. Shafiq, J.-G. Choi, M. Iqbal, and A. Irshad, "Flexible and efficient aggregation framework for antifragile wireless mesh networks," *Journal of Reliable Intelligent Environments*, vol. 1, no. 2-4, pp. 159–171, 2015.

[29] K. J. McKernan, "The chloroplast genome hidden in plain sight, open access publishing and anti-fragile distributed data sources," *Mitochondrial DNA*, pp. 1–2, 2015.

[30] A. Gorgeon, "Anti-Fragile Information Systems," *International Conference on Information Systems*, 2015.

[31] T. Bendell, "Stress to impress," *Significance*, vol. 11, no. 5, pp. 81–81, 2014.

[32] J. S. Levin, S. P. Brodfuehrer, and W. M. Kroshl, "Detecting antifragile decisions and models lessons from a conceptual analysis model of Service Life Extension of aging vehicles," in *IEEE Systems Conference (SysCon)*. IEEE, 2014, pp. 285–292.

[33] N. M. Weber, K. S. Baker, A. K. Thomer, T. C. Chao, and C. L. Palmer, "Value and context in data use: Domain analysis revisited," *Proceedings of the American Society for Information Science and Technology*, vol. 49, no. 1, pp. 1–10, 2012.

[34] S. Mittal, "Model engineering for cyber complex adaptive systems," in *Workshop on Model Engineering, European Modeling and Simulation Symposium*, 2014.

[35] L. Guang, E. Nigussie, J. Plosila, and H. Tenhunen, "Positioning antifragility for clouds on public infrastructures," *Procedia Computer Science*, vol. 32, pp. 856–861, 2014.

[36] V. De Florio, "On environments as systemic exoskeletons: crosscutting optimizers and antifragility enablers," *Journal of Reliable Intelligent Environments*, vol. 1, no. 2-4, pp. 61–73, 2015.

[37] K. D. Swenson, "Designing for an Innovative Learning Organization," in *Enterprise Distributed Object Computing Conference (EDOC)*, Sept 2013, pp. 209–213.

[38] F. Bonsignorio, A. P. del Pobil, and E. Messina, "Fostering Progress in Performance Evaluation and Benchmarking of Robotic and Automation Systems [TC Spotlight]," *IEEE Robotics Automation Magazine*, vol. 21, no. 1, pp. 22–25, March 2014.

[39] A. P. Tth and D. Svetinovic, "Identifying signs of systems fragility: A crowdsourcing requirements case study," in *Industrial Engineering and Engineering Management (IEEM)*, Dec 2013, pp. 356–360.

[40] G. Primiero and F. Raimondi, "Software theory change for resilient near-complete specifications," *The 5th International Symposium on Frontiers in Ambient and Mobile Systems*, vol. 52, pp. 988–995, 2015.

[41] V. De Florio and G. Primiero, "A framework for trustworthiness assessment based on fidelity in cyber and physical domains," *The 6th International Conference on Ambient Systems, Networks and Technologies*, vol. 52, pp. 996–1003, 2015.

[42] S. Marrone and R. Nardone, "Automatic resource allocation for high availability cloud services," *Procedia Computer Science*, vol. 52, pp. 980–987, 2015.

[43] R. Melo, A. Santos, M. Nogueira, and D. Medhi, "Resilience and Knowledge in a Metric for Heterogeneous Wireless Connectivity," Technical Report RT-DINF 003/2013, 2013, University of Missouri-Kansas City.

[44] V. De Florio, "On resilient behaviors in computational systems and environments," *Journal of Reliable Intelligent Environments*, vol. 1, no. 1, pp. 33–46, 2015.

[45] D. Russo and P. Ciancarini, "A Proposal for an Antifragile Software Manifesto," *Procedia Computer Science*, vol. 83, pp. 982–987, 2016.

[46] S. F. Haider, L. Abbas, A. Ali, M. Iqbal, I. Raza, S. A. Hussain, and D. Y. Suh, "Taxonomy and issues for antifragile-based multimedia cloud computing," *Journal of Reliable Intelligent Environments*, vol. 2, no. 1, pp. 37–49, 2016.

[47] J. Mitola, "Cognitive radio—an integrated agent architecture for software defined radio," 2000.

[48] W.-B. Yang and M. Souryal, "LTE Physical Layer Performance Analysis," *National Institute of Standards and Technology US Department of Commerce*, 2014.

[49] "Common Attack Pattern Enumeration and Classification (CAPEC)," MITRE. [Online]. Available: https://capec.mitre.org

[50] R. Poisel, *Modern Communications Jamming: Principles and Techniques*. London: Artech House, 2011.

[51] W. Xu, W. Trappe, Y. Zhang, and T. Wood, "The feasibility of launching and detecting jamming attacks in wireless networks," in *ACM international symposium on Mobile ad hoc networking and computing*. ACM, 2005, pp. 46–57.

[52] K. Pelechrinis, M. Iliofotou, and S. V. Krishnamurthy, "Denial of service attacks in wireless networks: The case of jammers," *IEEE Communications Surveys & Tutorials*, vol. 13, no. 2, pp. 245–257, 2011.

[53] M. Wilhelm, I. Martinovic, J. B. Schmitt, and V. Lenders, "Short paper: Reactive jamming in wireless networks: How realistic is the threat?" in *Proceedings of the ACM WiSec*. ACM, 2011, pp. 47–52.

[54] W. Xu, K. Ma, W. Trappe, and Y. Zhang, "Jamming sensor networks: Attack and defense strategies," *IEEE Network*, vol. 20, no. 3, pp. 41–47, 2006.

[55] M. Cakiroglu and A. T. Ozcerit, "Jamming detection mechanisms for wireless sensor networks," in *Proceedings of the 3rd international conference on Scalable information systems*, 2008, p. 4.

[56] Y. Liu and P. Ning, "BitTrickle: Defending against broadband and high-power reactive jamming attacks," in *INFOCOM, 2012 Proceedings IEEE*. IEEE, 2012, pp. 909–917.

[57] L. Wang and A. M. Wyglinski, "A combined approach for distinguishing different types of jamming attacks against wireless networks," in *IEEE Pacific Rim Conference on Communications, Computers and Signal Processing*. IEEE, 2011, pp. 809–814.

[58] D. R. Raymond and S. F. Midkiff, "Denial-of-service in wireless sensor networks: Attacks and defenses," *Pervasive Computing, IEEE*, vol. 7, no. 1, pp. 74–81, 2008.

[59] D. Thuente and M. Acharya, "Intelligent jamming in wireless networks with applications to 802.11 b and other networks," in *IEEE MILCOM*, vol. 6, 2006.

[60] R. Kohavi and F. Provost, "Glossary of terms," *Machine Learning*, vol. 30, no. 2-3, pp. 271–274, 1998.

[61] S. Amuru and R. M. Buehrer, "Optimal jamming strategies in digital communications-Impact of modulation," in *IEEE Global Communications Conference (GLOBECOM)*, 2014, pp. 1619–1624.

[62] R. Chen, J.-M. Park, and J. H. Reed, "Defense against primary user emulation attacks in cognitive radio networks," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 1, pp. 25–37, 2008.

[63] J. Mitola and G. Q. Maguire Jr, "Cognitive radio: making software radios more personal," *IEEE Personal Communications*, vol. 6, no. 4, pp. 13–18, 1999.

[64] T. C. Clancy and N. Goergen, "Security in cognitive radio networks: Threats and mitigation," in *Cognitive Radio Oriented Wireless Networks and Communications (CrownCom)*, 2008.

[65] T. Basar, "The Gaussian test channel with an intelligent jammer," *IEEE Transactions on Information Theory*, vol. 29, no. 1, pp. 152–157, Jan 1983.

[66] T. C. Clancy, "Efficient OFDM Denial: Pilot Jamming and Pilot Nulling," in *IEEE International Conference on Communications (ICC)*, June 2011.

[67] S. Amuru and R. M. Buehrer, "Optimal Jamming using Delayed Learning," in *IEEE MILCOM*, Oct. 2014.

[68] M. Labib, V. Marojevic, and J. Reed, "Analyzing and Enhancing the Resilience of LTE/LTE-A Systems to RF Spoofing," in *IEEE Conference on Standards for Communications and Networking Proces. (CSCN)*, Oct 2015.

[69] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) procedures in idle mode (Release 12)," 3rd Generation Partnership Project (3GPP), TS 36.304, Mar. 2015. [Online]. Available: http://www.3gpp.org/dynareport/36304.htm

[70] ——, "Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC) (Release 12)," 3rd Generation Partnership Project (3GPP), TS 36.331, Mar. 2015. [Online]. Available: http://www.3gpp.org/dynareport/36331.htm

[71] C. S. Patel, G. L. Stuber, G.ber, and T. G. Pratt, "Analysis of OFDM/MC-CDMA under channel estimation and jamming," in *IEEE Wireless Communications and Networking Conference*, vol. 2, 2004, pp. 954–958.

[72] M. Baker and T. Moulsley, "Downlink Physical Data and Control Channels," in *LTE, The UMTS Long Term Evolution: From Theory to Practice*, 2nd ed., S. Sesia, I. Toufik, and M. Baker, Eds. Chichester, West Sussex, United Kingdom: John Wiley & Sons Ltd, 2011, ch. 9.

[73] Sanjole, "WaveJudge 5000 LTE Analyzer," http://www.sanjole.com/our-products/lte-analyzer/.

[74] G. Gorbil, O. H. Abdelrahman, and E. Gelenbe, "Storms in mobile networks," in *Proceedings of the 10th ACM symposium on QoS and security for wireless and mobile networks*. ACM, 2014, pp. 119–126.

[75] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA); Physical channels and modulation (Release 8)," 3rd Generation Partnership Project (3GPP), TS 36.211, Dec. 2009. [Online]. Available: http://www.3gpp.org/dynareport/36211.htm

[76] R. P. Jover, J. Lackey, and A. Raghavan, "Enhancing the security of LTE networks against jamming attacks," *EURASIP J. on Inform. Security*, 2014.

[77] W. Xu, W. Trappe, and Y. Zhang, "Anti-jamming timing channels for wireless networks," in *Proceedings of the first ACM conference on Wireless network security*, 2008, pp. 203–213.

[78] S. D'Oro, L. Galluccio, G. Morabito, S. Palazzo, L. Chen, and F. Martignon, "Defeating Jamming With the Power of Silence: A Game-Theoretic Analysis," *IEEE Transactions on Wireless Communications*, vol. 14, no. 5, pp. 2337–2352, May 2015.

[79] Y. Zhao, "Enabling cognitive radios through radio environment maps," *Virginia Tech Electronic Theses and Dissertations*, 2007.

[80] S. Verdu and T. Han, "A general formula for channel capacity," *IEEE Transactions on Information Theory*, vol. 40, no. 4, pp. 1147–1157, 1994.

[81] D. J. Torrieri, "Fundamental limitations on repeater jamming of frequency-hopping communications," *IEEE Journal on Selected Areas in Communications*, vol. 7, no. 4, pp. 569–575, 1989.

[82] D. L. Adamy, *EW 102: a second course in electronic warfare.* Artech House, 2004.

[83] C. Ko, H. Nguyen-Le, and L. Huang, "ML-based follower jamming rejection in slow FH/MFSK systems with an antenna array," *IEEE Transactions on Communications*, vol. 56, no. 9, pp. 1536–1544, 2008.

[84] C. J. Burges, "A tutorial on support vector machines for pattern recognition," *Data mining and knowledge discovery*, vol. 2, no. 2, pp. 121–167, 1998.

[85] R. Caruana and A. Niculescu-Mizil, "An empirical comparison of supervised learning algorithms," in *Proceedings of the 23rd international conference on Machine learning*. ACM, 2006, pp. 161–168.

[86] M. Z. Win and J. H. Winters, "Analysis of hybrid selection/maximal-ratio combining of diversity branches with unequal SNR in Rayleigh fading," in *IEEE Vehicular Technology Conference*, 1999.

[87] C. E. Shannon and W. Weaver, "A Mathematical Theory of Communication," *University of Illinois Press*, 1949.

[88] P. Anghel, M. Kaveh *et al.*, "Exact symbol error probability of a cooperative network in a Rayleigh-fading environment," *IEEE Transactions on Wireless Communications*, vol. 3, no. 5, pp. 1416–1421, 2004.

[89] J. N. Laneman and G. W. Wornell, "Energy-efficient antenna sharing and relaying for wireless networks," in *IEEE Wireless Communications and Networking Conference*, 2000.

[90] J. M. Wozencraft and I. M. Jacobs, *Principles of Communication Engineering*. Wiley New York, 1965, vol. 32.

[91] P. Lee, "Computation of the bit error rate of coherent M-ary PSK with Gray code bit mapping," *IEEE Transactions on Communications*, vol. 34, pp. 488–491, 1986.

[92] H. V. Poor, *An Introduction to Signal Detection and Estimation*. Berlin: Springer, 1994.

[93] J. Proakis and M. Salehi, *Digital Communications*, ser. McGraw-Hill higher education. McGraw-Hill Education, 2007.

[94] K. Cho and D. Yoon, "On the general BER expression of one-and two-dimensional amplitude modulations," *IEEE Transactions on Communications*, vol. 50, no. 7, pp. 1074–1080, 2002.

[95] W. E. Stark, "Capacity and cutoff rate of noncoherent FSK with nonselective Rician fading," *IEEE Transactions on Communications*, vol. 33, no. 11, pp. 1153–1159, 1985.

[96] I. Jacobs and J. Wozencraft, *Principles of communication engineering*. Waveland Pr Inc, 1965.

[97] K. Jordan Jr, "The performance of sequential decoding in conjunction with efficient modulation," *IEEE Transactions on Communication Technology*, vol. 14, no. 3, pp. 283–297, 1966.

[98] Y. Xiao, "IEEE 802.11 n: enhancements for higher throughput in wireless LANs," *Wireless Communications, IEEE*, vol. 12, no. 6, pp. 82–91, 2005.

[99] D. Nguyen, C. Sahin, B. Shishkin, N. Kandasamy, and K. R. Dandekar, "A real-time and protocol-aware reactive jamming framework built on software-defined radios," in *Proceedings of the 2014 ACM workshop on Software radio implementation forum*. ACM, 2014, pp. 15–22.

[100] T. S. Rappaport *et al.*, *Wireless communications: principles and practice*. Prentice Hall PTR New Jersey, 1996, vol. 2.

[101] N. Vaidya, A. Dugar, S. Gupta, and P. Bahl, "Distributed fair scheduling in a wireless LAN," *Mobile Computing, IEEE Transactions on*, vol. 4, no. 6, pp. 616–629, 2005.

[102] I. Rhee, A. Warrier, J. Min, and L. Xu, "DRAND: distributed randomized TDMA scheduling for wireless ad-hoc networks," in *Proceedings of the 7th ACM international symposium on Mobile ad hoc networking and computing*. ACM, 2006, pp. 190–201.

[103] A. B. Forouzan, *Data Communications & Networking*. Tata McGraw-Hill Education, 2006.

[104] S. Amuru and R. M. Buehrer, "Optimal Jamming Against Digital Modulation," *IEEE Transactions on Information Forensics and Security*, vol. 10, no. 10, pp. 2212–2224, Oct 2015.

[105] L. Kleinrock and F. A. Tobagi, "Packet switching in radio channels: Part I–carrier sense multiple-access modes and their throughput-delay characteristics," *Communications, IEEE Transactions on*, vol. 23, no. 12, pp. 1400–1416, 1975.

[106] N. Abramson, "THE ALOHA SYSTEM: another alternative for computer communications," in *Proceedings of the November 17-19, 1970, fall joint computer conference*. ACM, 1970, pp. 281–285.

[107] J. R. Wertz and W. J. Larson, *Space mission analysis and design*. Berlin: Microcosm, 1999.

[108] Northrop Grumman, "AEHF Payload: Assured, protected, survivable communications," 2014.

[109] A. Einhorn and J. Miller, "Spectrum management issues related to the AEHF system," in *IEEE MILCOM*, 2007, pp. 1–7.

[110] P. Muri and J. McNair, "A survey of communication sub-systems for intersatellite linked systems and cubeSat missions," *Journal of Communications*, vol. 7, no. 4, pp. 290–308, 2012.

[111] National Telecommunications and Information Administration, "United States Frequency Allocations," http://www.ntia.doc.gov/files/ntia/publications/spectrum_wall_chart_aug2011.pdf, 2011, accessed: 2014-12-1.

[112] A. N. Ince, *Digital Speech Processing: Speech Coding, Synthesis and Recognition*. Berlin: Springer, 1992, vol. 155.

[113] M. Strasser, B. Danev, and S. Capkun, "Detection of Reactive Jamming in Sensor Networks," *ACM Transactions on Sensor Networks*, vol. 7, no. 2, pp. 16:1–16:29, Sep. 2010.

[114] I. Shin, Y. Shen, Y. Xuan, M. T. Thai, and T. Znati, "A Novel Approach Against Reactive Jamming Attacks," *Ad Hoc & Sensor Wireless Networks*, vol. 12, no. 1-2, pp. 125–149, 2011.

[115] A. Perrig, J. Stankovic, and D. Wagner, "Security in wireless sensor networks," *Communications of the ACM*, vol. 47, no. 6, pp. 53–57, 2004.

[116] M. Li, I. Koutsopoulos, and R. Poovendran, "Optimal jamming attacks and network defense policies in wireless sensor networks," in *IEEE INFOCOM*, 2007, pp. 1307–1315.

[117] A. Wood and J. A. Stankovic, "Denial of service in sensor networks," *IEEE Computer*, vol. 35, no. 10, pp. 54–62, 2002.

[118] H. Wang, J. Guo, and Z. Wang, "Feasibility assessment of repeater jamming technique for DSSS," in *IEEE Wireless Communications and Networking Conference*. IEEE, 2007, pp. 2322–2327.

[119] S. Chuprun and C. S. Bergstrom, "Comparison of FH/CDMA and DS/CDMA for wireless survivable networks," in *IEEE Conference on Universal Personal Communications (Florence, Italy)*, vol. 2. IEEE, 1998, pp. 1305–1309.

[120] H. Wang, J. Guo, and Z. Wang, "Evaluation of security for DSSS under repeater jamming," in *IEEE International Conference on Communications*. IEEE, 2007, pp. 5525–5530.

[121] H. R. Anderson, *Fixed Broadband Wireless System Design*. New York: John Wiley & Sons, 2003.

[122] T. Pratt, C. Bostian, and J. Allnutt, *Satellite Communications*. Electronic Industry Press, Beijing, 2005.

[123] T. Lewis, "TDRSS 2nd Workshop," http://msp.gsfc.nasa.gov/TUBE/pdf/infopack.pdf, Omitron, Inc., 1996, accessed: 2014-12-1.

[124] ViaSat, "VMT-1220 Ground Mobile Terminal," 2013.

[125] "IEEE 802.11b-1999," http://standards.ieee.org/getieee802/download/802.11b-1999.pdf, accessed: 2014-12-1.

[126] R. ITU-R P.676-2, "Attenuation by atmospheric gases," http://www.itu.ch/itudoc/itu-r/rec/p/676-2_29169.html, 1995, accessed: 2014-12-1.

[127] M. Yang and D. Grace, "Cognitive radio with reinforcement learning applied to heterogeneous multicast terrestrial communication systems," in *Cognitive Radio Oriented Wireless Networks and Communications, 2009*. IEEE, 2009, pp. 1–6.

[128] Y. Zhu, X. Li, and B. Li, "Optimal Adaptive Antijamming in Wireless Sensor Networks," *International Journal of Distributed Sensor Networks*, vol. 2012, 2012.

[129] Y. Wu, B. Wang, and K. R. Liu, "Optimal defense against jamming attacks in cognitive radio networks using the markov decision process approach," in *IEEE Global Telecommunications Conference 2010*. IEEE, 2010, pp. 1–5.

[130] S. Sodagari and T. C. Clancy, "An anti-jamming strategy for channel access in cognitive radio networks," in *Decision and Game Theory for Security*. Springer, 2011, pp. 34–43.

[131] S. Amuru, C. Tekin, M. van der Schaar, and R. Buehrer, "Jamming Bandits - A Novel Learning Method for Optimal Jamming," *IEEE Transactions on Wireless Communications*, vol. PP, no. 99, pp. 1–1, 2015.

[132] S. Amuru, C. Tekin, M. van der Schaar, and R. M. Buehrer, "A systematic learning method for optimal jamming," in *IEEE International Conference on Communications (ICC)*, June 2015, pp. 2822–2827.

[133] S. Amuru, Y. Xiao, M. van der Schaar, and R. M. Buehrer, "To Send or Not to Send - Learning MAC Contention," in *2015 IEEE Global Communications Conference (GLOBECOM)*, Dec 2014, pp. 1–6.

[134] S. Amuru, "Intelligent Approaches for Communication Denial," 2015.

[135] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. Cambridge Univ Press, 1998, vol. 1, no. 1.

[136] R. Tandra and A. Sahai, "SNR walls for signal detection," *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 1, pp. 4–17, 2008.

[137] S. Atapattu, C. Tellambura, and H. Jiang, *Energy Detection for Spectrum Sensing in Cognitive Radio*. Springer, 2014.

[138] M. A. Richards, *Fundamentals of radar signal processing*. Tata McGraw-Hill Education, 2005.

[139] I. Cosovic and V. Janardhanam, "Sidelobe suppression in OFDM systems," in *Multi-Carrier Spread-Spectrum*. Springer, 2006, pp. 473–482.

[140] S. Brandes, I. Cosovic, and M. Schnell, "Reduction of out-of-band radiation in OFDM systems by insertion of cancellation carriers," *IEEE Communications Letters*, vol. 10, no. 6, pp. 420–422, 2006.

[141] I. Cosovic, S. Brandes, and M. Schnell, "Subcarrier weighting: a method for sidelobe suppression in OFDM systems," *IEEE Communications Letters*, vol. 10, no. 6, pp. 444–446, 2006.

[142] H. Mahmoud, H. Arslan *et al.*, "Sidelobe Suppression in OFDM-Based Spectrum Sharing Systems Using Adaptive Symbol Transition," *IEEE Communications Letters*, vol. 12, no. 2, pp. 133–135, 2008.