Institut für
Technische Informatik und
Kommunikationsnetze

Jamming Detection in Wireless Ad Hoc Networks

Markus Schafroth

Masterthesis MA-2008-21

September 2008 - March 2009

Computer Engineering and Networks Laboratory

Advisors: Dr. V. Lenders and Dr. F. Legendre
Professor: Prof. Dr. B. Plattner

ETH Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich
Ecole polytechnique fédérale de Zurich
Politecnico federale di Zurigo

Abstract

As wireless networks are gaining higher popularity, their security is becoming a critical aspect. If adequate countermeasures should be taken against denial of service (DoS) attacks, known as jamming, the first line of defence is the differentiation between jamming and poor network conditions due to environmental influences. Therefore, in this practical work we deploy and evaluate several algorithms for the detection of jamming. Based on metrics like the received signal strength (RSSI) and packet delivery ratio (PDR), we formulate measurement-based models of normal and jammed network behaviour. Along with the self-contained detection at each node, we also investigate the use of a cooperative detection algorithm and the detection at dedicated nodes. Results gained from our indoor WLAN testbed show a high detection performance of our algorithms in static networks. In particular, results show that the dedicated jamming detection achieves a high precision. This allows for saving energy and processing power resources, since the jamming detection is performed by a few dedicated nodes only. Although the detection is affected by node mobility and high background traffic load, the detection performance is still remarkable under such circumstances.

Mit der zunehmenden Verbreitung von drahtlosen Kommunikationsnetzwerken gewinnen Denial of Service (DoS) Attacken, auch Jamming genannt, an Bedeutung. Um geeignete Gegenmassnahmen zu treffen, ist eine Unterscheidung zwischen Jamming und natürlichen Einflüssen notwendig. In dieser praktischen Arbeit werden mehrere Algorithmen zur Erkennung von Jamming entwickelt und getestet. Basierend auf Metriken wie der Signalstärke (RSSI) und der Paketzustellrate (PDR) werden experimentell Modelle von normalem und gejammtem Netzwerkverkehr erstellt. Neben der eigenständigen Detektion durch alle beteiligten Netzwerkknoten wird auch untersucht, wie sich ein kooperativer Algorithmus und insbesondere die dedizierte Detektion bei ausgewählten Knoten bewähren. Die Resultate der auf WLAN basierenden Experimente zeigen eine sehr hohe Detektionsgenauigkeit unserer Algorithmen in statischen Netzwerken. Insbesondere die dedizierte Detektion erwies sich als vielversprechend. Somit können für die Jamming Detektion einzelne Knoten mit mehr Energie- und Rechenkapazität eingesetzt werden, während die übrigen Knoten keine Ressourcen dafür benötigen. Auch wenn Mobilität und starke Netzwerkauslastung die Detektion erschweren, werden trotzdem noch gute Resultate erzielt unter solch erschwerten Bedingungen.

Contents

List of Figures

List of Tables

1 Introduction

Wireless communication systems are encountered in many situations of our daily life. Along with radio and television broadcasting, mobile phones and global positioning devices, there are many other applications where wireless communication technology is used today. Further examples are radio communication in the aviation or in the police and rescue sector. Many people also run wireless access points at home to share their internet connection all over the house. Most universities cover their campuses with wireless local area networks (WLAN) and an increasing number of restaurants and hotels offer wireless internet access to their customers. A newer trend is the equipment of trains and aeroplanes with wireless networks.

Other applications of wireless networks are sensor networks as they are increasingly used in the building automation sector. In this latter scenario, sensors placed all over the building collect data like temperature, humidity, brightness etc. and transmit it to the control room, where the appropriate actions are triggered. This might include the regulation of heating or of net curtains for example. There exist also wireless fire alarm systems for the usage in larger buildings. Such installations based on wireless communication systems provide more flexibility and low cost installation over the usage of wired systems.

Depending on the intended use of the wireless communication system, the confidentiality, integrity and availability of the data transfer is of particular importance. Therefore it is crucial that all possible countermeasures are taken to prevent attacks against those three basic requirements. If wireless communication systems are employed in a business or even military environment, their security is becoming even more critical.

1.1 Motivation

In this thesis, we focus on the availability of wireless communication networks. The first step towards a higher level of availability is the analysis of reasons which cause decreased performance or even the complete failure of the communication system. Since there are many effects which affect the network performance, it is worth trying to distinguish between poor network conditions due to natural environmental influences and intentional attacks driven by adversaries. Reasons for natural degradation of the network performance are

- congestion due to a high traffic load in the network
- large distance between transmitter and receiver (high path loss)
- obstacles or walls in the line of sight between transmitter and receiver (shadowing)
- mobility of network nodes (multipath fading)
- other wireless communication systems using the same frequency band (interference)
- other electronic devices which may cause interference (e.g. microwave oven)

Similar effects as they occur due to environmental conditions may also be forced intentionally by attackers. Such denial of service (DoS) attacks, which are targeted against the medium access and network availability, are called "jamming" and the executors are called "jammers". There is virtually no limit in the number of ways a wireless communication system may be jammed. Chapter 3 gives an overview of several strategies used in different jamming models. The main goal remains the same among most of those techniques: either artificially lower the signal-to-noise ratio (SNR) at the receiver, occupy the channel all the time so as no other transmitting station ever detects a free time slot or lower the throughput by causing collisions with specific frames of the attacked data transfer. For the sake of completeness, it has to be stated that there are even more ways of jamming depending on the actual communication standard used in the targeted wireless network. Some well known attacks against WLANs for example use a vulnerability

in the way mobile stations are logically connected to the network. In short, it is possible to disconnect all nodes by sending only a few spoofed frames. If this is done repeatedly, the nodes are not able to create a network again and therefore the data transfer collapses. Even if such attacks are very efficient at this time, they might become harder if not impossible with upcoming WLAN standards in the future.

1.1.1 On the Significance of Jamming Detection

Due to the diversity of impacts affecting the performance of a wireless communication system, it is not possible to give general advices which guarantee its availability in every situation. Especially, there are different measures to be taken in case of natural environmental influences or intended denial of service attacks respectively. Therefore, the first line of defence is the detection of jamming and its discrimination from other poor network conditions.

In the past, there have been investigations to counter denial of service attacks on the physical layer. Proposed mechanisms are the use of spread spectrum techniques either by using direct-sequence spread spectrum (DSSS) or frequency hopping spread spectrum (FHSS) mechanisms. Even though this reduces the impact of jamming, a jammer operating at higher power might be able to limit or even prevent the operability of a wireless communication system. Therefore, jamming is still a real threat and further research in the field of jamming detection is necessary.

1.1.2 Investigated Wireless LAN Scenario

We consider a military scenario in this thesis as illustrated in figure 1.1. There is a wireless network for the communication among a group of soldiers by foot and on vehicles. The communication within close proximity is realized by the use of a self-organizing ad hoc network, while data targeted at remote locations is routed via satellite links installed on vehicles. We further assume that all communication links are encrypted at the link layer using strong encryption. This means that only authorized entities are able to decrypt data traffic and generate valid data frames with arbitrary contents.

Since we consider an operation, we expect a certain mobility among the in-

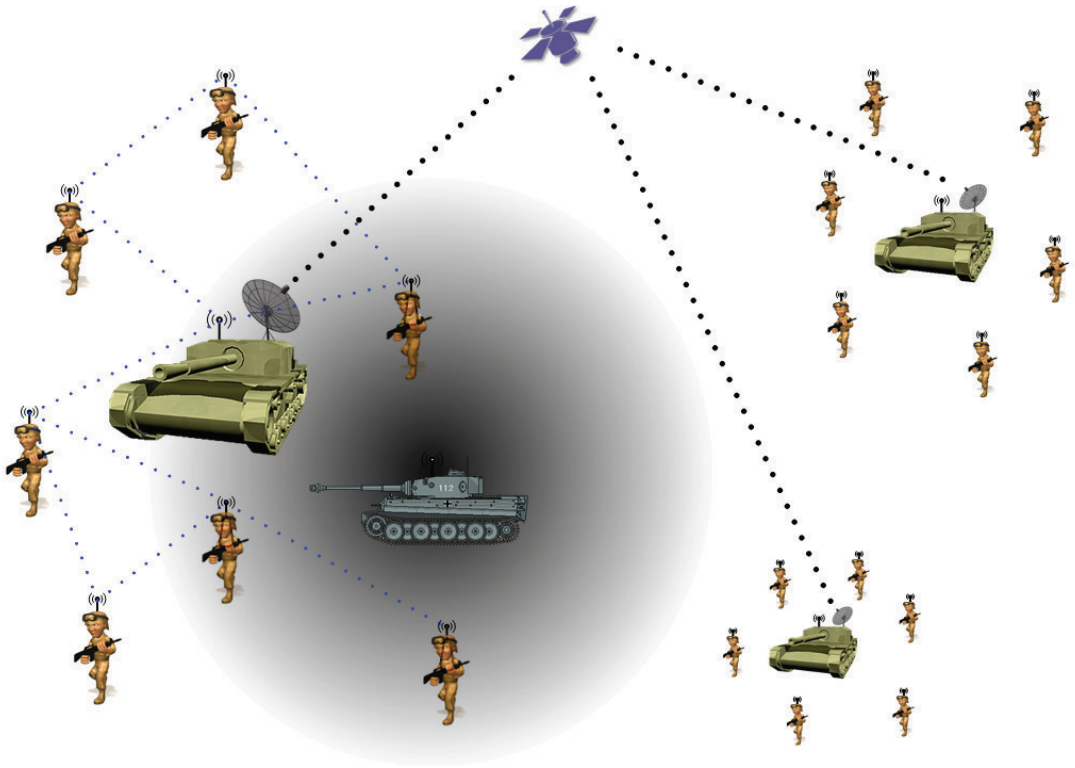

Figure 1.1: The investigated wireless LAN scenario: A military setting where several mobile entities communicate using a wireless ad hoc network. Long-distance communication is routed via satellite links installed on vehicles. Jamming attacks (depicted as the grey circle) should be detected inside the ad hoc network. Source: by author

volved individuals and vehicles. This may lead to poor network conditions due to increasing distances between the nodes or caused by multipath effects for instance. Furthermore, there may be obstacles in the line of sight which influence the communication. Even though the strong encryption protects our network from eavesdropping, we are not immune to attacks targeting the medium access. Therefore, we need an algorithm which allows to detect such physical layer jamming attacks. In figure 1.1, the jammer is illustrated by the single adversarial tank and the grey circle denoting the range of its transmission area.

There are several possibilities on how and where exactly the jamming detection is performed in the ad hoc network. For example, every single node might perform the detection locally. In contrast, it is conceivable that all nodes share information, which might enhance the detection performance. As a further alternative, a single dedicated node analyzes the whole network autonomously and notifies the other nodes upon detection of jamming. In our scenario, this particular approach is of special interest, since the computing power and energy resources of most wireless nodes are very limited. That is, the jamming detection for the whole local wireless network should be performed by the nodes placed on vehicles preferably, where sufficient energy is available.

1.2 Related Work

In the literature, we find extensive studies concerning jamming of wireless communication systems. Many investigations try to reduce the jammer impact by the use of physical layer technologies such as spread spectrum or frequency hopping [3], [4], [5]. However, these techniques are not always applicable due to various reasons including device complexity, additional energy consumption or spectral efficiency. Therefore, the detection of jamming is still a critical issue.

In [6], Wood et al. propose a distributed algorithm to recognize jammed regions in a wireless sensor network. As the metric for the jamming detection, they define a utility threshold which is based on the ability of a certain node to communicate in both send and receive directions. Although the concept of a distributed algorithm proves promising, it does not solve the underlying problem of discriminating jamming from poor network performance caused by natural influences.

Xu et al. proposed several jamming models in [7] targeting the physical layer of wireless communication systems. They also discussed different measurements which may serve as the basis for the detection of jamming attacks. It was shown that each of these measurements by itself is not sufficient to reliably classify the presence of jamming attacks. Through empirical experiments, they showed that the combination of several measurements allows the detection of the jamming attacks they proposed. However, it has to be noted that those experiments were all based on a static scenario, meaning that the effects of mobility were ignored.

Along with the studies adressing denial of service attacks on the physical layer, research has been made also on upper layer attacks and attack detection, for example [8], [9] and [10]. Wood et al. recapitulate known denial of service attacks and possible defense strategies on the physical, link, network and transport layer in [11]. However, since the upper layer attack mechanisms highly depend on the deployed link layer protocols, those results are less related to our work.

1.3 Project Goals and Contributions

The goal of this thesis is the development and assessment of algorithms for the detection of denial of service attacks in wireless communication networks. Based on the underlying theory of wireless jamming detection studied in the literature, we go a step further and study how to practically implement jamming detection algorithms in networks of mobile nodes. For this purpose, we rely on different performance indicators available on commodity wireless LAN devices. On the basis of these metrics, we formulate measurement-based models of different network behaviour. This includes both static and mobile network scenarios either with or without the influence of jamming. We consider jamming which significantly reduces the network performance, but which does not completely interrupt the communication. We then evaluate different implementations of the developed algorithms, where the detection is performed either transmitter-based, receiver-based, by a dedicated node, or as a distributed process between the cooperating nodes in the network. The evaluation of extensive measurements and experimental results sheds light onto the performance and overhead of these different jamming detection approaches.
Along with the analysis of the self-contained, dedicated and distributed detection algorithms, further contributions and differentiation to the preceding research in jamming detection are

- the use of realistic unicast (UDP and TCP) instead of broadcast traffic
- the use of a physical link using multiple modulation schemes and data rates adaptively
- the evaluation of mobility effects on the detection performance
- the evaluation of the detection performance in a scenario with regular background network traffic
- the analysis of the effectiveness of jamming attacks and their detection depending on the positioning of the jammer in the network.

1.4 Report Overview

This section gives a short overview of the report's structure. In chapter 2, we discuss some technical aspects of the IEEE 802.11-2007 WLAN standard which influence our implementation of the deployed jamming detection algorithms. Readers who are familiar with the standard may want to skip this chapter. Chapter 3 shows several strategies for denial of service attacks to the medium access in wireless networks. The subsequent chapter 4 explains the developed algorithms for the detection of such attacks. In chapter 5, our implementation of the detection algorithms are presented along with the details of the experimental testbed setup. In chapter 6, we evaluate the performance of the implemented jamming detection system. Chapter 7 discusses the strengths and weaknesses of the developed algorithms and it contains recommendations for countermeasures against jamming attacks. In chapter 8, we conclude our work and finally, in chapter 9, we provide an outlook to further research which might be done based on the insights of this present work. Along with some additional materials, the assignment and time schedule of the thesis as well as a CD-ROM containing the used software tools and scripts are enclosed in the appendix.

2 Important Aspects of the WLAN Standard (IEEE Std 802.11-2007)

Our jamming detection models rely on carrier sense multiple access networks with collision avoidance mechanism (CSMA/CD). For the evaluation of our models, we use IEEE Std 802.11-2007 [1], known as WLAN, as a reference. For this reason, in this chapter we provide an overview of the most important aspects of WLAN.

Wireless LAN or WLAN is a commonly used communication standard in wireless computer networks today. It is defined by the IEEE[1] as the 802.11 standard. Since its first publication in 1997, several enhancements regarding data throughput and security have been added. However, the basic functionality has remained the same since then. The standard aims to provide a reliable channel which, on the link layer, looks like a common Ethernet link known as IEEE 802.3. As such, it makes no difference for software programs whether they communicate via wireless or wired local area networks, as long as the used communication protocol compatible with the Ethernet MAC layer. In practice, this wired-like behaviour is not trivial to achieve. It is obvious that there arise additional difficulties when the communication takes place over the air since this medium is shared among several networks and other applications in the same frequency band. In contrast, Ethernet uses a wire accessible for a closed group of subscribers only. Therefore, wireless LANs need a mechanism which logically separates all available networks while at the same time coordinating the access of all networks to the medium as fairly as possible. Another important point is the coverage area, which is clearly defined when using a wired infrastructure, but rather vague in case of a wireless network.

In the following, we present the relevant aspects in terms of security and avail-

[1] Institute of Electrical and Electronics Engineers, http://www.ieee.org

ability of wireless LAN networks. Furthermore, we focus on some pecularities of the pysical layer (PHY), which are of particular importance when studying denial of service attacks to wireless LANs. We also look at available security mechanisms and show why they are not able to prevent several attacks against the availability of the network.

Readers who are familiar with the IEEE 802.11-2007 standard may skip these explanations and continue at chapter 3.

2.1 Functional Principles and Comparison with Wired LANs

In this section, we give an overview of the functional principles of a wireless LAN. It is not the intention to look into every detail of the over 1200 pages long standardization document, but to focus on the aspects which are relevant for the detection of jamming. First of all, we show some basic problems which occur only in wireless networks, and we discuss how they are approached in the standard.

2.1.1 Problems Specifically Related to Wireless Networks

In wireless communication, several problems occur independently on the actual used system. Air as the medium used for the data transfer may be used by several communication systems with an unlimited number of subscribers at the same time, and those systems may not be explicitely designed for the concurrent and coordinated use with other systems in the same frequency band. Therefore, every communication system should be able to handle low signal-to-noise scenarios and also packet collisions due to interference from other systems. Furthermore, varying distance between the transmitting and receiving stations could lead to a loss of connection between the two. This should be detectable by the transmitter, since it makes no sense to generate traffic while the receiver is out of reach.

A well known problem related to wireless networks is the "hidden node problem", illustrated in figure 2.1. It occurs, when three or more wireless nodes are arranged so that two of the stations are outside the coverage area of each other, but both inside the coverage area of a third station. In our example, the red and blue node may not see each other, but they can both communicate with the yellow node. We assume that the protocol used by the three nodes implements a collision avoidance mechanism based on the energy measured on the channel. This means that before starting the transmission of a packet, the channel is sensed for ongoing transmissions. As soon as the channel is free, the packet is sent. Although this behaviour helps to greatly minimize collisions on the channel, it is not sufficient in our scenario. If, for example, a data transmission is in progress between the red and yellow node, the packets sent by the yellow node might not be detected by the blue node. Therefore, the blue node assumes that the channel is free and

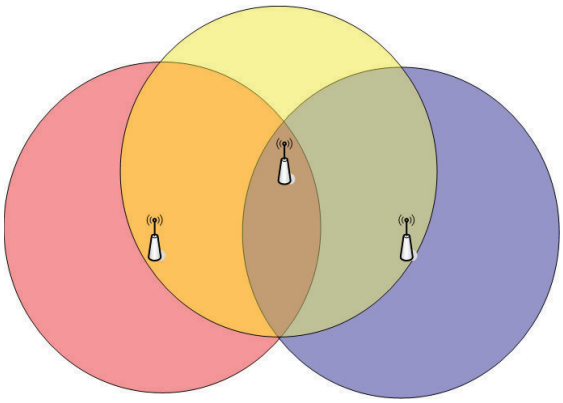

Figure 2.1: Hidden node problem. Source: by author

starts its transmission to the yellow node. This finally leads to packet collisions at the yellow node.

In the case of WLAN, this specific problem is approached by the use of a RTS/CTS handshake before every data frame transmission. The idea is that the station which initiates the data transfer asks the receiver whether it may start the transmission (Ready-To-Send, RTS). If the channel is free in the receiver's region, the receiver responds with an acknowledgement (Clear-To-Send, CTS). The RTS and CTS frames inform other stations about the duration of the following data transmission which makes them stay in quiet mode. However, since this handshake has to be done before the transmission of every single data packet, it generates a certain overhead traffic and slows down the effective throughput, which does not make sense under most operating conditions. Therefore, the RTS/CTS mechanism is an optional feature in the WLAN standard and not widely used.

In summary, the following problems are specific to wireless networks:

- interference from other applications using the same frequency band leading to low signal-to-noise ratio
- low signal quality due to large distance between transmitter and receiver or due to multipath fading (e.g. in buildings)
- hidden node problem, leading to packet collisions

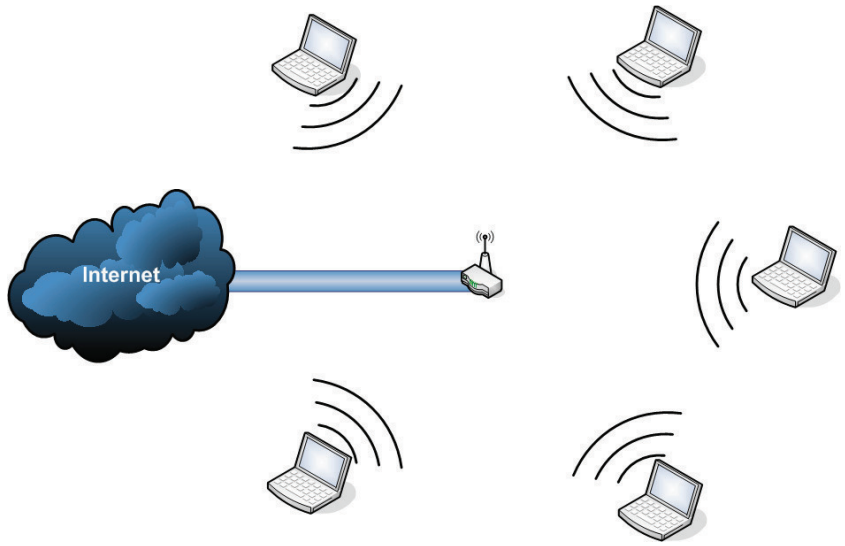

Figure 2.2: Wireless LAN based on the infrastructure mode. Source: by author

2.1.2 Wireless Networking Principles

There are two basic principles for the setup of a wirless LAN. In many cases where WLANs are deployed, the "Infrastructure Mode" is used. An extension to this first mode even allows for a roaming mechanism across several infrastructure networks. The second principle is known as the "Ad Hoc Mode", which is probably less frequently used in case of WLANs, but quite often in combination with other wireless techniques like sensor networks. Below we will explain the two principles "Infrastructure" and "Ad Hoc" more in detail.

Infrastructure Mode

In the infrastructure mode, there is one dedicated node serving as the central access point (AP) for all the other nodes in the same network. Such a network, consisting of an access point and one or more clients, is called a Basic Service Set (BSS). The network traffic may be routed via the access point either to one of the other nodes inside the local network or outside to the wide area network (WAN / internet). For clients in reach of each other, it is possible to communicate directly with each other. In the infrastructure mode, the access point may also undertake additional tasks like

- bridging a local wired LAN with wireless LAN

- wireless service advertising
- authentication and association of wireless clients
- assignment of IP adresses to the devices in the LAN/WLAN (DHCP)
- DNS-relay and/or assignment of DNS adresses
- network address translation (NAT) between WAN and LAN/WLAN
- providing security features including firewall and access restrictions (in- and outbound direction)

Most of those tasks are common to routers in wired LANs as well and are therefore not further discussed. However, we will present the following two concepts more in detail, since they are specific for wireless LANs:

Wireless Service Advertising In the infrastructure mode, the access point periodically sends beacon packets which contain information about the wireless network it is serving. These packets contain the Service Set Identifier (SSID), which is the ID of the corresponding WLAN, as well as some information about the physical layer like the used channel and supported data rates. Additional fields of the beacon packet are used for the announcement of quality of service (QoS) and for the management of larger wireless networks with several access points (see [1], pages 80/81 for more details). Wireless nodes which receive those beacons use the gained information to contact one of those networks. However, if the SSID of a WLAN is known, it is possible to contact the corresponding network even if no beacons were received before.

Authentication and Association When a wireless client tries to connect to a wireless LAN, it has to authenticate first. It is important to note that only the client has to authenticate against the access point but not vice versa. There are two ways how this authentication may be performed. The "Open System Authentication" is a simple handshake where the client tells it's MAC address to the access point. In the following, the access point replies with an authentication response. The authentication may be refused if for example the client's MAC

address is explicitely blacklisted in the configuration of the access point. The second authentication mechanism is called "Shared Key Authentication". There is again a handshake between client and access point, but the client additionally has to know a key shared beyond all authorized subscribers. This key is used in order to reply with the encrypted version of a challenge previously sent to the client by the access point.

As soon as the authentication has completed, the client may try to associate with the access point, which is comparable to plugging in the cable in case of a wired LAN. For this purpose, the client sends an association request to the access point and receives a reply with a status code denoting either successful or unsuccessful association. One possible reason for the access point to refuse the association to a client could be that a defined limit of connected clients has exceeded. In case of a successful association, the client is now able to transmit and receive frames to and from the access point and all other nodes associated with the same BSS.

If we consider a large WLAN with several access points, the concept of association gains in importance. As a wireless client may be associated with only one access point at a time, it is ensured that frames to and from the client are always handled by the correct access point. For more details about the authentication and association procedures refer to [1], pages 35 et seqq.

Ad Hoc Mode

The ad hoc mode allows for the configuration of wireless LANs without a fixed infrastructure. In fact, two or more wireless stations may easily build their own Independent Basic Service Set (IBSS) by just defining a common SSID and, in order to enable the transfer of data among each other, manually assigning fixed IP adresses to the devices.

Unlike in the infrastructure mode, there is no centralized organization of the clients in the ad hoc mode and all nodes are equal in their functionality. This requires some adaptations in the way new clients get in touch with an existing network, compared to the case where we have a dedicated access point. The main differences concern the service advertising and authentication/association concepts as shown in the following.

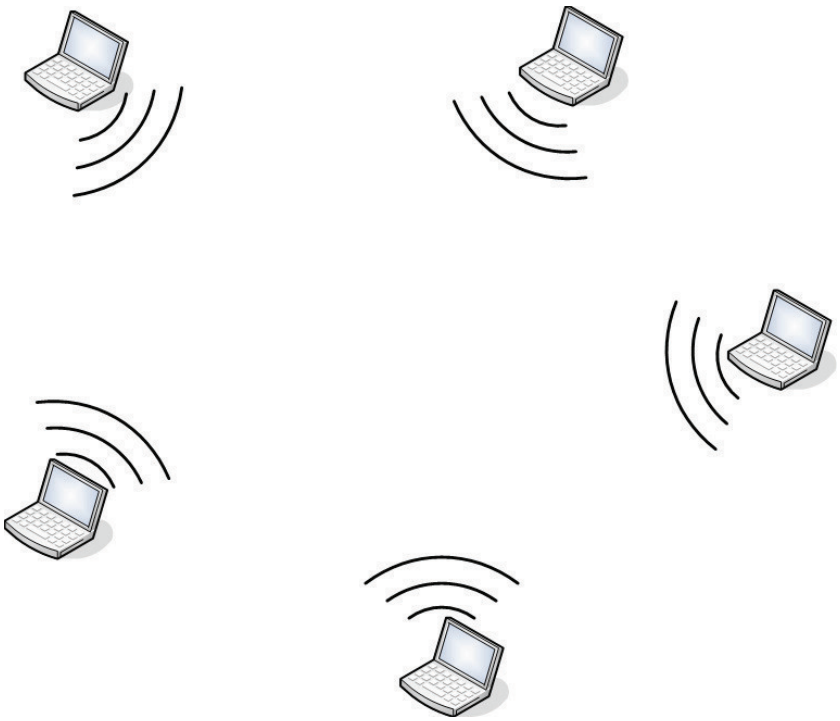

Figure 2.3: Wireless LAN based on the ad hoc mode. Source: by author

Wireless Service Advertising In ad hoc WLANs, the beacons are not sent by one station only, instead the beacon initiation rotates among all stations belonging to the same IBSS.

Authentication and Association The first two wireless stations building a new ad hoc network must agree on a security policy. Clients which want to join the IBSS later have to agree with this policy. There is, like in an infrastructure WLAN, the need for an authentication mechanism. "Open System Authentication" as well as "Shared Key Authentication" are available for ad hoc networks as well. However, they work slightly different, since the authentication has to be done with every station separately. In ad hoc networks, there is no need for an association mechanism, since the communication is always performed directly between two participating clients. For further details refer to the IEEE standard directly.

2.1.3 Coordination Functionalities and Data Transfer in Wireless LANs

If we imagine the communication in a wireless environment, three main aspects are of special interest:

- how may the subscribers avoid collisions resulting from concurrent data transfer?
- how can the subscribers verify that their messages are received by the intended receiver?
- what types of different frames are needed for the communication and how are they built?

Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA)

The wireless LAN standard includes a CSMA scheme as its collision avoidance mechanism. This procedure helps to reduce packet collisions as the stations which want to send data packets listen to the channel prior to their transmissions. The decision whether the channel is busy or idle is based on two distinctive principles. The first one physically measures the energy present on the channel and decides based on a threshold value. The second principle is provided by the frame-structure. As we will discuss in the next section, every frame contains a field which specifies the time needed for the frame transmission. Additionally, if the RTS/CTS procedure is used as described in section 2.1.1, the channel may be reserved for the duration of the data and acknowledgement transmissions. All stations waiting for an empty transmission slot will then remain silent for this time interval.

If several stations are in waiting state, it is obvious that they would all start transmitting at the same time, namely as soon as the medium is idle and no reservation is available. To reduce the collision probability at the beginning of idle time slots, the stations wait another random backoff time while still sensing the medium. If during this backoff time the channel gets busy, the station remains silent and waits for the medium to become idle again, otherwise the transmission starts.

Positive Acknowledgement and Retransmissions

In contrast to the wired LAN where CSMA with Collision Detection (CSMA/CD) is widely used, this is not intended in the wireless standard. For collisions to be detectable by the transmitting device, it has to be capable of a full-duplex operation mode. This means that there needs to be a separate device for the transmission and reception part respectively. According to the IEEE 802.11-2007 standard, the same device should be used for both the transmission and reception of frames, which means that WLAN devices are only half-duplex capable. This decision was made for reasons of economy due to the higher complexity and price of WLAN chipsets compared with LAN chipsets.

A workaround for this problem is included in the wireless LAN standard by using a positive acknowledgement mechanism with the possibility of retransmissions. Whenever station B receives a data frame from station A, station B replies with an acknowledgement frame (ACK). If A receives this ACK within a short time period after the transmission of the initial data packet, the following data frame is sent. If no ACK is received, this might be either because the data packet was lost or corrupted and therefore B did not send an acknowledgement, or, t is possible that station B got the data frame, but the ACK was lost. In the end, for station A this makes no difference and the data frame has to be sent again. These retransmitted frames obtain a special marking in their MAC header.

MAC Frame Format

From wired LANs, network packets of different types and usages are known. For example, there are packets which enable the resolution of an IP adress to the MAC adress of the corresponding network interface using the Address Resolution Protocol (ARP), other packets may indicate errors due to missing hosts or routes using the Internet Control Message Protocol (ICMP). For the purpose of data transfer, the Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) with their specific packet structure are widely used.

The wireless LAN standard offers the required functionality for using the exact same protocols and packet types as they are used in wired LANs. The difference lies in the design of the lower MAC layer, which requires some additional features. As we saw in section 2.1.2, a mechanism for the service advertisement is required.

Furthermore we must keep in mind that collisions may occur without the sender to notice it. This implicates the need for a mechanism which allows the sender to recognize collisions by the use of acknowledgements. Other mechanisms like authentication and association also demand for special features of the MAC layer.

Figures 2.4 and 2.5 show the structure of a general MAC frame as it is used in wireless and wired LANs respectively. Both frames show several address fields containing the destination and source MAC address, a field containing the user data ("Frame Body" and "MAC Client Data" respectively) as well as a Frame Check Sequence (FCS). The additional fields in the 802.3-frame are used by the receiving device for the synchronization with the frame's timing ("Preamble") and to indicate the start of a frame (Start Frame Delimiter, "SFD"). The "Length/Type" field indicates either the length of the payload data or the nature of the MAC client protocol. A "PAD" field is only present if the length of the "Data" field is very small, it may be interpreted as a filler. The same applies to the "Extension" field following the frame check sequence.
Although wireless devices need some synchronisation information as well, this is resolved on a special layer as described later in section 2.2.1. Since wireless MAC frames are used for additional purposes which are not necessary in wired LANs, they have some additional fields compared to the frames in wired LANs:

- Frame Control: specifies the type of frame, whether the frame has been retransmitted and whether it contains an encrypted payload
- Duration/ID: denotes the duration in milliseconds used for the transmission of the actual frame and corresponding acknowledgement frames
- Addresses 1 - 4: together with destination and source MAC address, also the SSID is represented by a MAC-like 48-bit field; the fourth address field may be used if the frame is being relayed e.g. by the access point
- Sequence Control: this field contains a sequence number which helps to identify multiple receptions of the same frame due to lost acknowledgements
- QoS Control: if Quality of Service shall be used, this field contains the appropriate parameters

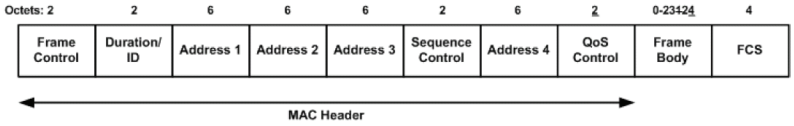

Figure 2.4: General MAC frame format in wireless LANs (according to IEEE 802.11-2007). Source: [1], page 60

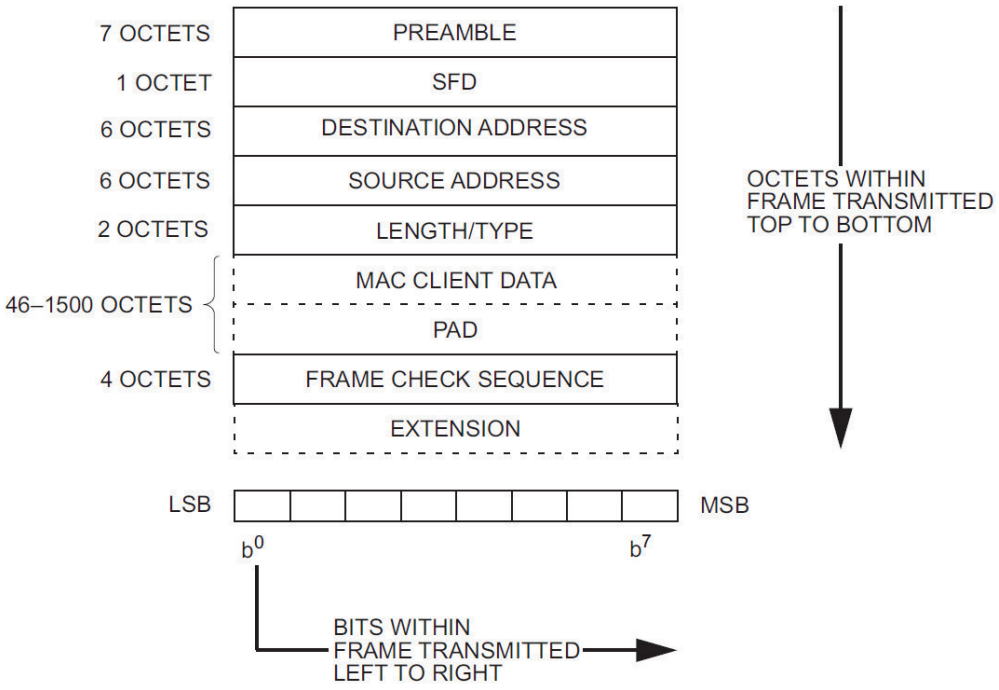

Figure 2.5: General MAC frame format in wired LANs (according to IEEE 802.3). Source: [2], page 49

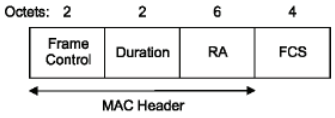

Figure 2.6: ACK frame format. Source: [1], page 74

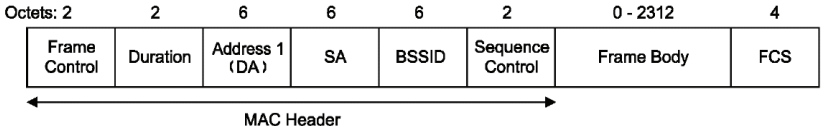

Figure 2.7: Management Frame Format. Source: [1], page 79

Upper layer packets like ARP, ICMP, TCP (including the TCP acknowledgements) and UDP are all transported inside 802.11 data frames.

Along with those data frames, there exist many specialized control and management frame subtypes which have different frame formats depending on their needings.

Examples for control frames:

- ACK
- RTS / CTS

Figure 2.6 shows the format of the acknowledgement frames. The "RA" field contains the MAC address of the station which initiated the original frame.

Examples for management frames:

- Beacon
- Authentication / Deauthentication
- Association / Deassociation

The frame format which is the same for all management frames, independent of the frame subtype, is shown in figure 2.7. Depending on the subtype of management frame, the frame body contains different information.

2.2 Physical Layer Highlights

In the previous section, we mainly focussed on the MAC layer of the WLAN standard. Now we want to look at the structure and functionality of the lower layer, which is the physical layer (PHY), and the Physical Layer Convergence Procedure (PLCP).

2.2.1 Physical Layer Convergence Procedure (PLCP)

The MAC layer of IEEE 802.11-2007 is intended to be independent of the subjacent physical layer. The PLCP defines a method for the mapping of 802.11 MAC frames onto the physical layer, which is used for the transmission. Since there are several different physical layer architectures in use, it is reasonable to define a sublayer which converges their pecularities, which ensures the compatibility. Before the physical layer is discussed more in detail, we present the PLCP frame format as depicted in figure 2.8.

We begin the analysis of its nested structure bottom-up. The whole PLCP frame is called a PLCP Protocol Data Unit (PPDU). This includes all the elements which are required for the transmission of one single MAC frame over the medium.

The next higher level distinguishes between the PLCP and the MAC part of the frame. On the left we have the bits which are sent first, these are the PCLP preamble and header. After this, the MAC Protocol Data Unit (MPDU) is sent, which contains the MAC frame as discussed in section 2.1.3.

The PLCP preamble consists of two fields which are the "SYNC" and "SFD" field. Just like in the case of wired LANs, those fields are used for the synchronisation with the incoming frame. The start frame delimiter is a fixed bit pattern denoting the beginning of the real frame. In contrast to wired LANs, this functionality is not defined as a part of the MAC layer but of the physical layer.

The PLCP header contains information concerning the modulation and duration of the frame:

- Signal: indicates the modulation scheme used for the transmission and reception of the MPDU
- Service: this field is reserved for future use

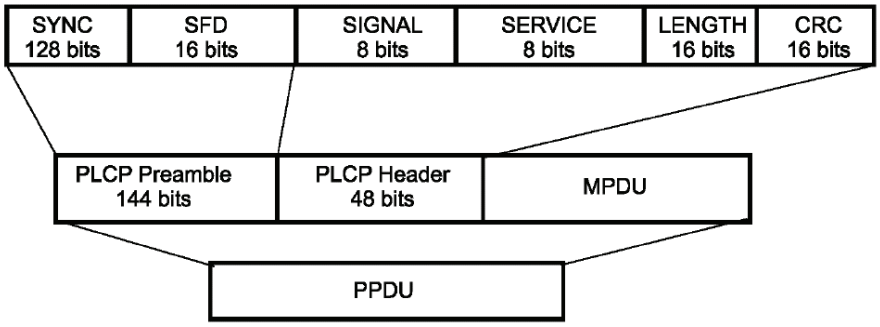

Figure 2.8: PLCP frame format. Source: [1], page 538

- Length: indicates the number of microseconds required to transmit the MPDU
- CRC: protects the Signal, Service and Length fields with a cyclic redundancy check

The frame format shown in figure 2.8 is according to the physical layer based on direct-sequence spread spectrum, however the PLCP frame formats of the other PHYs look similar.

2.2.2 Physical Layer (PHY) Architectures

As mentioned before, there are several architectures defined in the WLAN standard which may be used at the physical layer. In the rest of this thesis, we assume the use of direct-sequence spread spectrum and orthogonal frequency-division multiplexing PHYs. For the sake of completeness and to give an impression of the diversity of the IEEE 802.11-2007 standard, all available physical layer architectures should be mentioned at least. For the detailed specifications of those physical layers refer to chapters 14 to 19 of [1].

- Frequency-Hopping Spread Spectrum (FHSS)
- Direct-Sequence Spread Spectrum (DSSS)
- Infrared
- Orthogonal Frequency-Division Multiplexing (OFDM)

- High Rate Direct-Sequence Spread Spectrum (HR/DSSS)
- Extended Rate PLCP Sublayer (ERP)

2.2.3 Operating Channels in the 2.4 GHz Band

In Europe, most WLANs operate in the unlicensed 2.4 GHz band. This means that everybody may install and operate such wireless networks as long as certain conditions concerning the maximum allowed radiated power are kept. There is a second frequency band in the 5 GHz range which may be used for wireless LAN communications, which is more frequently used in the USA. In the following we will focus on the specifications available for the 2.4 GHz band.

Depending on the geographic-specific regulatory authorities, the exact frequency spectrum licensed for WLANs may differ. In most of Europe (except France and Spain), there are 13 channels available in the range from 2412 MHz to 2472 MHz. Since the channels are separated only by 5 MHz but have a bandwith of 20 MHz each, there results an overlapping of maximum three channels at the same frequency. This may lead to a decrease in the achieved throughput if neighbouring channels are working at high capacity.

The availability of several operating channels enables the stations in a wireless LAN to choose the appropriate channel according to the presence of other networks and potential parasitic coupling from other sources. In fact, there is no channel hopping like in frequency-hopping systems where the operating frequency is changed several times per second. It is rather a last resort if the measured signal quality degrades and packet loss exceeds a certain level.

2.2.4 Data Rates for DSSS and OFDM PHY

The modulation schemes offered by the DSSS and OFDM physical layers allow for different maximum achievable data throughput. One might ask why not always use the scheme offering the highest transmission speed. Nevertheless, it is highly important to have low rate modulation schemes which are more robust to low signal quality at the receiver. The differences in robustness between the modulation schemes mainly come from three characteristics:

- spreading ratio

- symbol constellation
- forward error correction coding rate

In digital modulation systems, data bits are mapped to symbols which may then be represented physically by the emitted electromagnetic waves. DSSS systems use a mechanism to spread x data bits onto y symbols, there $x < y$. This allows for a better decoding of the data bits at the receiver due to the included redundancy. At the same time it is harder to be revealed by adversaries since the spectrum of the modulated signal may not be easily distinguished from white noise if the used spreading sequence is unknown. The downside of the spread spectrum approach is the loss of throughput capacity. By varying the spreading ratio, these effects may be optimized for the desired performance.
The spreading ratio is defined as

$$\frac{\textit{length of the spreading-sequence code}}{\textit{number of bits per symbol}}$$

Due to the nature of the OFDM modulation mechanism, there are additional possibilities for the definition of symbols. OFDM makes use of two orthogonal electromagnetic waves, which may be used to carry information independently. Therefore, it is possible to define more advanced symbol constellations like 16- and 64-QAM (Quadrature Amplitude Modulation). This means that when using 16-QAM, 16 different symbols may be distinguished instead of only two or four when using BPSK (Binary Phase-Shift Keying) or QPSK (Quadrature Phase-Shift Keying) respectively. The large number of distinguishable symbols may then be used for a Forward Error Correction (FEC) by introducing redundancy. Every of the OFDM symbols represents a number of data bits, depending on the symbol constellation choosed, which equates to the upper limit of data throughput per symbol. However, it is possible to define a mapping where x data bits are represented by a sequence of y bits. This mapping is then somehow similar to the spread-sequence mechanism in DSSS.
Definition of the Forward Error Correction Coding Rate (FEC Coding Rate) (cf. [1], page 604):

$$\frac{\textit{number of data bits per OFDM symbol}}{\textit{max. number of bits per OFDM symbol}}$$

Table 2.1 gives an overview of the available physical data rates and the corresponding modulation details for the widely used DSSS and OFDM physical layers.

Rate [Mbps]	Mod. Scheme	Symb. Const.	Spread. Ratio	FEC Coding Rate
1	DSSS [2]	DBPSK	11/1	—
2	DSSS [2]	DQPSK	11/2	—
5.5	DSSS [3]	DQPSK	8/4	—
6	OFDM	BPSK	—	1/2
9	OFDM	BPSK	—	3/4
11	DSSS [3]	QPSK	8/8	—
12	OFDM	QPSK	—	1/2
18	OFDM	QPSK	—	3/4
24	OFDM	16-QAM	—	1/2
36	OFDM	16-QAM	—	3/4
48	OFDM	64-QAM	—	2/3
54	OFDM	64-QAM	—	3/4

Table 2.1: 802.11 PHY rates and corresponding modulation details. Sources: [1], pages 537, 597, 674-678

Finally, we also mention how the devices select an appropriate PHY rate for their transmissions. The IEEE 802.11-2007 standard leaves this open, however it forces some control and management frames to be transmitted at one of the lowest rates. This is to maintain a certain level of compatibility with older devices, since new and higher physical data rates have been added to the standard over time.

There are many different rate switching algorithms in use by device and driver manufacturers. Using different approaches, they all try to maximize the possible data throughput. One simple approach which is used by the cards in our experiments is to start the first transmission of a frame at the highest available rate. If an ACK for this frame may be received at the transmitter, the next frame will again use this highest rate. If no ACK is received, the frame is being retransmitted at the next lower rate and so on. If several frames did not reach the receiver at high rates, it is reasonable to start further transmissions at a lower rate directly. From time to time, the sender could then try to send at a higher rate to see whether the rate might be increased again. In appendix B, we present four well-known rate

[2] DSSS using the 11-chip Barker Code, a particular spreading code. cf. [1], page 567
[3] DSSS using the Complementary Code Keying (CCK). cf. [1], page 674

switching algorithms available in the open source Linux driver implementation "MadWiFi" [4] for example.

[4] http://www.madwifi.org

2.3 Security Features and their Limitations

In section 2.1 we explained in short the mechanism of authentication in wireless LANs. The idea of this procedure is to make sure that only authorized clients may access the network in order to send and receive data frames in the corresponding BSS or IBSS respectively. However, this authentication mechanism provides no additional security like privacy or integrity for the data sent over the channel.

In the following, we will look at existing security mechanisms offered by the IEEE 802.11-2007 standard. We will not go into deep detail, since all three mechanisms may not find a remedy for attacks against the availability of wireless networks, i.e. jamming.

Wired Equivalent Privacy (WEP)

The wired equivalent privacy mechanism was introduced with the original IEEE 802.11 standard in 1997. It aimed to provide the same confidentiality level as it is available in a wired network, meaning that only people knowing the private key could compose and decompose the transmitted frames. Today, this claim must be rejected. Borisov et al. showed the vulnerabilities of WEP in 2001 [12]. They also showed that the use of larger encryption keys (104 bits instead of 40 bits) did not prove to be more secure, since the design flaws of WEP may still be misused by attackers. A major flaw of WEP is its vulnerability against message forgeries and other active attacks (cf. [1], page 169). It must be stated that today WEP basically does not provide any security at all and therefore has been deprecated in the actual version of the IEEE 802.11 standard (cf. [1], page 157).

WiFi Protected Access (WPA)

WPA may be viewed as an enhanced WEP mechanism. It includes a message integrity code (MIC), which makes it much more difficult for attackers to guess the used encryption parameters, since active attacks may not be easily conducted any more (cf. [1], page 169). WPA is the implementation of a subset of the IEEE 802.11i standard amendment [13]. It was made by the Wi-Fi Alliance[5] in 2003, even before 802.11i was released. This amendment was the reaction to the security

[5] a global non-profit industry association, http://www.wi-fi.org

flaws of WEP. One of the main ideas of WPA was to allow the fast spreading and a wide acceptance of the new security measure. This made it necessary to allow existing wireless devices to be upgradeable by a firmware update but had the drawback that no real redesign of the security mechanism could be implemented.

WiFi Protected Access (WPA2)

With the publication of the IEEE 802.11i amendment in 2004, the WiFi Alliance also released a new implementation of its mandatory elements, called WPA2. It introduced new encryption algorithms which required changes to the wireless device hardware. These mechanisms, known as Advanced Encryption Standard (AES), toghether with new key distribution methods (e.g. the four-way handshake, cf. [1], pages 211 et seqq.), are believed secure to date. All devices bearing the WiFi trademark which are sold today are WPA2 certified and therefore comply with the 802.11i amendment. In 2007, this amendment was also integrated into the new IEEE 802.11-2007 standard.

The above described security mechanisms all have in common, that they try to provide a certain level of confidentiality in the data transfer. This is done by encryption of the payload of the data packets, which, if this encryption is strong enough, makes casual eavesdropping impossible. However, not the whole frames sent over a wireless LAN are covered by this encryption, but only specific parts of them. Depending on their type, many frames are not encrypted at all. We will show in the following, to which parts of communication the encryption mechanisms effectively apply to.

Encrypted and Unencrypted Data in Wireless LANs

In section 2.1.3, we presented the format of the MAC frames. It is important to note what portion of the frames is being encrypted if any of the above mentioned security mechanisms is applied. According to IEEE 802.11-2007, no encryption is applied to any of the control and management frames. The only exception are the authentication frames if shared key authentication is used. However, the new encryption methods are based upon the open system authentication. The secure authentication and access control is then made by a handshake mechanism which

uses encrypted data frames instead. This said, it is clear that only the data frames are encrypted. But even this conclusion is rather misleading, since not the whole content of the data frame is encrypted. Only the frame body is encrypted according to either WEP, WPA or WPA2. From this follows that even when using the best available encryption technology available in IEEE 802.11-2007, it is easily possible to overhear the communication and to gain quite a lot of information about the ongoing data transfers. What might be of interest are the MAC adresses of the subscribed devices or the used SSID, which may be found unencrypted in every data frame. Another important security hole arises from the unencrypted management frames, making it easy to compose deauthentication and deassociation frames. By the injection of only one such spoofed frame, adversaries may disconnect several or even all wireless clients from the access point in infrastructure networks. In ad hoc networks, one spoofed deauthentication frame would be necessary to disconnect two clients pairwise. Like this, very efficient attacking models exist today against the availability of WLANs. Since also the control packets are sent unencrypted, intelligent attackers may also utilize this information for their interests.

3 Jamming Techniques

Wherever data transfers take place, there may exist individuals or groups which are interested in overhearing or disturbing this communication. Depending on the security measures taken in the communication system, attackers find it hard to decrypt the cyphertext of the sent messages. Nevertheless, they may be able to make it difficult or even impossible for the legitimate participants of the wireless network to access the medium.

This chapter provides an overview of several strategies how adversaries may conduct denial of service attacks.

3.1 Jamming Strategies

A jamming strategy describes the way an attacker disturbs the medium. Besides the time-based strategies, where the jamming signal is active only in specific time intervals, there are more advanced jamming schemes possible which make use of knowledge about the physical and link layer specifications of the targeted system. Based on the selected strategy, the effective jamming is then performed by emitting an appropriate radio frequency signal. This could be noise or modulated signals.

Figure 3.1 shows several jamming strategies arranged in a tree structure. Such diagrams, called attack trees, were introduced by Bruce Schneier in [14]. They show how specific threats may be reached by the corresponding attacks. In our case, the attack tree shows different jamming strategies which may lead to a failure of communication in a wireless system.

Xu et al. describe several jamming strategies in [7], from where we adopt certain expressions of our jamming terminology.

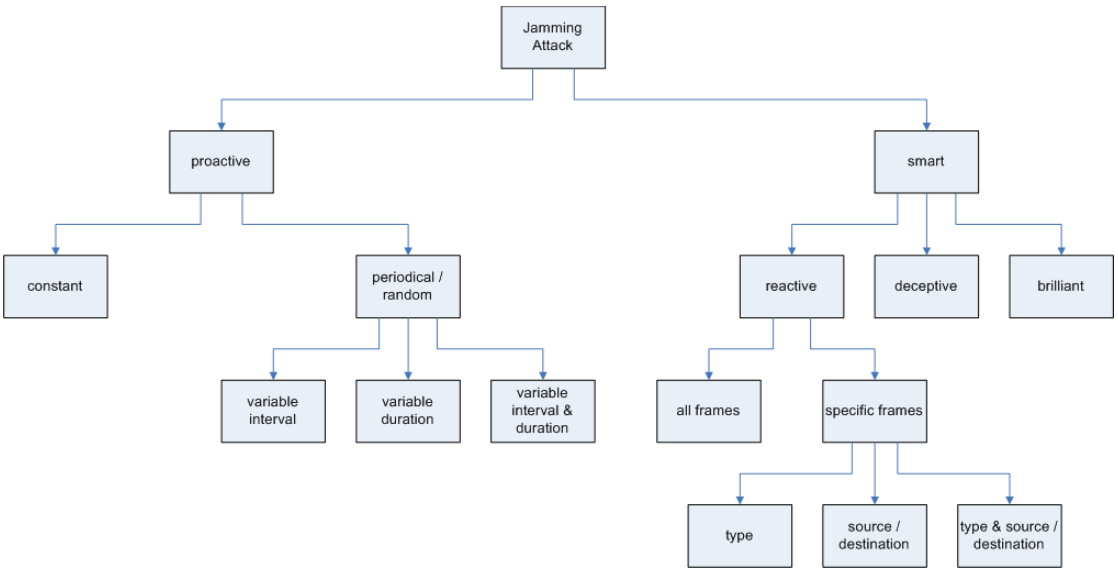

Figure 3.1: Attack tree: Different jamming strategies used in denial of service attacks. Source: by author

3.1.1 Proactive Jamming

The left half of the jamming attack tree shown in 3.1 is characterized as proactive jamming. This means, that the jammer emits a signal irrespective of the regular network traffic.

Constant Jamming A constant jammer continuously emits a signal meaning that there are no silent time intervals in its transmission.

Periodical / Random Jamming In contrast to the constant jammer, a periodical jammer suspends its transmission during a specified time in regular intervals. A modified version is the random jammer, which uses either a random duration, a random interval or both.

3.1.2 Smart Jamming

Smart jammers use a certain a priori knowledge of the used communication system in order to optimize their attacks. Several smart jamming strategies are shown in the right half of the attack tree in figure 3.1. As attacks of this type highly depend on the used communication system, there are an infinite number of possible

jamming strategies. In the following, we mention three general concepts of smart jamming.

Reactive Jamming Reactive jamming requires the sensing of the channel. As soon as a transmission is detected, the jammer starts its transmission.

A more advanced form of reactive jamming includes the analysis of the detected regular data stream. The jamming is then applied systematically to frames from or to specific nodes or to frames of a certain type.

Deceptive Jamming Deceptive jamming denotes attacks where false messages are sent to the channel with the objective of disturbing the organization of the network. In case of WLAN, this could be spoofed management or control frames for example. This way, also higher layer vulnerabilities may be easily exploited in order to run denial of service attacks. Examples for deceptive jamming are deauthentication and deassociation attacks as mentioned in section 2.3.

Brilliant Jamming Poisel mentions another form of smart jamming in [5], which he calls brilliant jamming. Such brilliant jammers attempt to change specific bit patterns of the frames. However, this requires a very high timing precision and significant a priori knowledge of the target signal structure.

3.2 Investigated Jamming Strategies

In our work, we focus on proactive jamming strategies which may be realized with commercially available signal generators. When all MAC frames including the administration frames are encrypted like in our envisioned scenario, this remains the only effective jamming technique.

If specific frames are to be jammed, it is necessary that the attacker analyzes all detected frames very quickly, meaning that the jamming must start before the entire frame reaches the receiver. This requires very low processing latency and it is difficult to achieve in communication systems running in the frequency band of 2.4 GHz or 5 GHz and offering data rates of up to 54 Mbps like it is the case in WLANs. Even though it is possible to build such jammers, specialized and

expensive equipment like software-radios[1] together with technically demanding programming and testing are required for its implementation.

[1] http://en.wikipedia.org/wiki/Software_radio

4 Detection Algorithms

In this chapter, we introduce the deployed algorithms for the detection of jamming in wireless communication systems. We also show the requirements applicable to the used system, which are implicated by the specification of the algorithms. The chapter is organized as follows: First, the basic idea of the detection mechanism are explained. This includes the used metrics and the decision algorithm itself. Then we discuss four implementations of the decision algorithm which are transmitter-based, receiver-based, dedicated and cooperative detection respectively.

4.1 Detector Model

Poor conditions in a wireless network are observed at the connected nodes in different ways. They might be unable to decode arriving frames correctly due to low signal strength or a high interference level. On the other hand, they might not get any response from the nodes where they send data to, or no data may be sent at all.

Xu et al. showed in [7] that it is not sufficient to base the "jammed / not jammed" decision on one single metric. One of their propositions is the correlation of signal strength (Received Signal Strength Indication, RSSI) and packet delivery ratio (PDR) at the receiver. The RSSI denotes the signal energy measured during the reception of a frame preamble while the packet delivery ratio is defined as

$$\frac{\textit{number of received and successfully decoded frames at the receiver}}{\textit{number of frames sent by the transmitter}}$$

In addition to the signal strength and packet delivery ratio metrics, we propose to treat frames sent at different physical rates separately and to use an additional noise metric. The latter describes the level of energy measured on the channel,

which is not identifiable as regular data traffic. Such interference may come from wireless systems using other physical layer specifications or from electronic devices which are not fully shielded. Some jamming activities are also registered by the interference metric.

Since these metrics are widely independent of the actual system, the deployed jamming algorithms are not specific to IEEE 802.11-2007.

4.1.1 Detection Based on RSSI vs. PDR and PHY rate

Let us consider a scenario where node A sends to node B. If the path loss between the two nodes is small, then frames arrive with a high signal strength at node B. In general, the packet delivery ratio and the used physical rate should also be high. Now if the path loss between A and B increases, the signal strength will decrease at the receiver. Assuming that the noise level remains constant, when the signal-to-noise ratio (SNR) decreases, frame decoding at node B becomes error prone, resulting in a decrease of PDR and PHY rate. Hence, there is a clear correlation between signal strength, packet delivery ratio and physical rate under normal operating conditions, meaning without jamming.

If we consider now a jammed scenario, our observations are different in parts. In the case of a large path loss between the transmitting node A and receiving node B, the signal strength at B remains unchanged at a low level. The PDR and PHY rate do not drop significantly, since they were low even without the presence of jamming. This makes it hard to distinguish between a high path loss scenario without jamming from the same scenario but with the presence of jamming. However, this discrimination is not relevant since the only way to increase the network performance significantly is to decrease the path loss between A and B or, if applicable, to reroute the traffic via a third node C, located between the other two nodes.

When the path loss between A and B is small, the detection of jamming proves easier. The receiving node B then observes a high signal strength while the packet delivery ratio and physical rate are lower than expected.

Figures 4.1 and 4.2 show the different RSSI versus PDR and PHY rate patterns of two consecutive experiments with the same underlying node mobility. The differences are remarkable: The experiment without the presence of jamming is

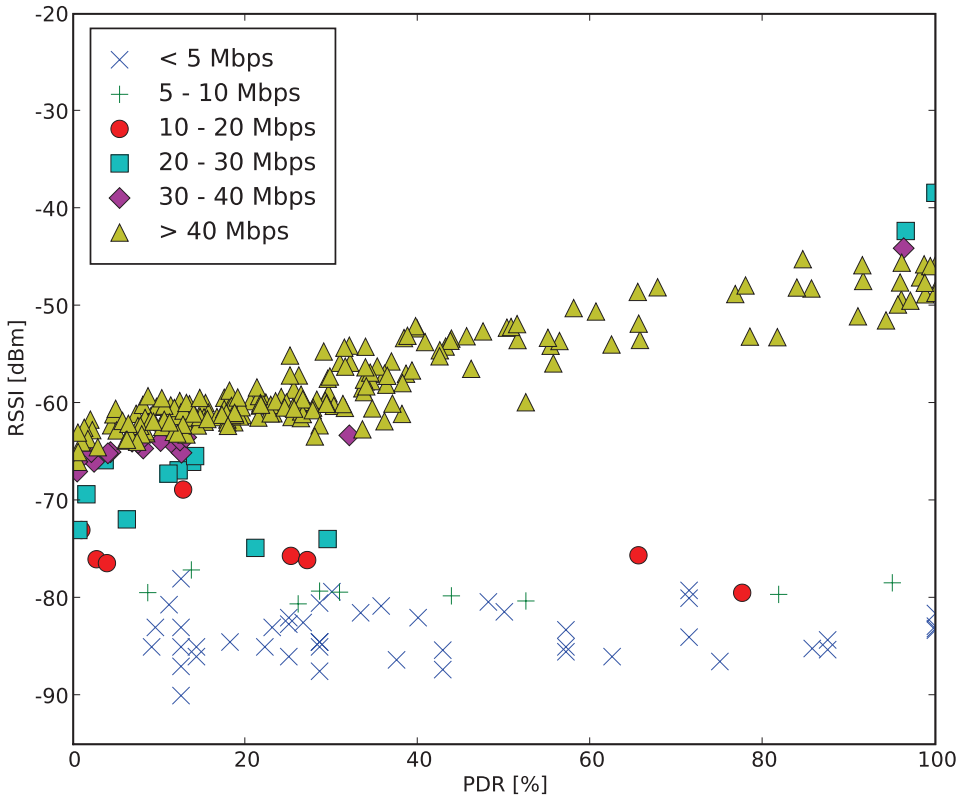

Figure 4.1: RSSI vs. PDR and PHY rate in mobile scenario without jamming

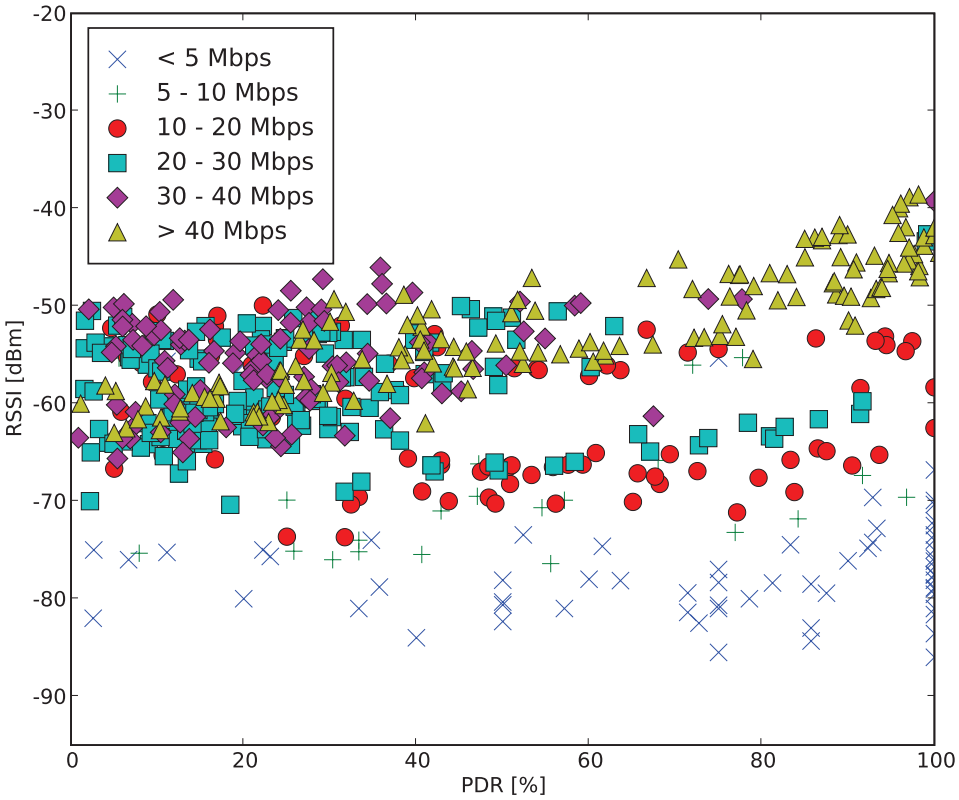

Figure 4.2: RSSI vs. PDR and PHY rate in mobile scenario with frame jamming

represented by figure 4.1, where we can see a certain correlation between RSSI and PDR for frames sent at the same physical rate. By contrast, this correlation is less distinctive in the case where frame jamming is disturbing the communication.

In order to formulate a reference model for the operation under normal conditions, we conduct measurements in an environment free of jamming. This could for example be achieved at manufacturing time and stored on the radio devices or later. For all physical rates supported by the device, a separate set of measured values is required. This may be seen also in figure 4.1, as the relationship between RSSI and PDR additionally depends on the physical rate used. The reason for this lies in the different modulation schemes as described in section 2.2.4. We achieve this measurements by moving a node such that it traverses different path loss conditions and the corresponding physical transmission rates.

4.1.2 Detection Based on Noise

In our detection algorithms, we have one more metric which we did not use in the above considerations yet. The noise metric may be regarded as a very strong measure for the link quality. However, in real-world environment we may not always conclude that there is a jammer present, if noise reaches a certain level. Other regular communication using the same frequency band but other physical layer properties influence the noise metric and could lead to false positive decisions. This means that it is inappropriate to define a fixed value as the noise threshold at manufacturing. If the noise metric is used as an indicator for jamming, it is crucial that the level of regular interference is known precisely.

4.1.3 Decision Algorithm

Our decision algorithm models the *existence* of a jammer as a Bernoulli process, where the existence variable, E, can take on the two values 0 and 1, where $E = 0$ indicates the absence of jamming and $E = 1$ denotes its presence. If the influence of the jammer is below a certain sensitivity level, the jammer is considered inexistent.

The decision is based on hypothesis testing. This is a method of using experimental data to make statistical decisions. We formulate the null hypothesis as the absence of jamming. Available sample data is then tested against this null hy-

pothesis. If the sample data is compatible with the null hypothesis, we accept the null hypothesis, meaning no jammer is present. If the sample data is incompatible with the null hypothesis, a jammer is assumed to be present.

We define two conditions which need to be fulfilled for the null hypothesis to be accepted:

1. the noise level is lower than *noiseThreshold*
2. the measured RSSI value is within the range

$$E_{RSSI}(PDR, PHYrate) \pm rssiDiffThreshold$$

noiseThreshold denotes the maximum level of interference which is acceptable in a non-jammed scenario. This parameter is set dynamically according to known friendly influences from other radio devices as discussed above.

$E_{RSSI}(PDR, PHYrate)$ is the expected RSSI value corresponding to the measured PDR for frames sent at a certain physical rate. This function is retrieved from the reference measurements.

$rssiDiffThreshold$ denotes the maximum acceptable difference between the measured and the expected RSSI value in a non-jammed scenario. This parameter has to be determined empirically in order to optimize the false positive and false negative rates of the decision algorithm respectively.

According to the above definition of the null hypothesis, we consider frames with RSSI values higher *and* lower than a certain interval as an indication for jamming. However, for the highest and lowest available PHY rates we propose the following additional rules, since at these rates the rate switching algorithm can not further increase or decrease the rate respectively:

- **for the lowest and highest available PHY rates, frames with RSSI values lower than expected are classified as "not jammed".**

  lowest PHY rates: If frames are received at lower RSSI than expected, this is an indication for very good network conditions, since otherwise these frames would not have been detected at all.

highest PHY rates: If frames are successfully transmitted at the higest available rate even if the RSSI is not as high as expected according to the reference measurements, this is an indication for good network conditions.

- **for all PHY rates except the lowest, frames with a PDR higher than $maxPDR$ are classified as "not jammed".**

  We observed only weak correlation between PDR and RSSI for high values of the PDR. This comes from the fact that the rate switching algorithm selects higher rates only after successful delivery of several frames at the lower rate. If for example the RSSI is increased due to mobility effects, this also augments the PDR immediately. However, the rate might stay low for the next few frames due to the delay of the rate switching algorithm. This leads to situations where we observe high RSSI and PDR together with low PHY rates.

  very low PHY rates: Since under normal network conditions the lowest PHY rates are only used along with very low RSSI values, it is unlikely that high RSSI and PDR values will ever occur together with the lowest PHY rates.

  $maxPDR$: This boundary depends on the used communication system and needs to be determined empirically.

The definition of "very low PHY rates" depends on the physical layer specification of the used system. The underlying idea is to apply the adapted treatment only to those rates which use the most robust modulation scheme at the cost of a very low maximum throughput.

The decision algorithm may not be performed on the basis of one single frame arriving at the receiver. For the calculation of a reliable PDR value, several received frames are necessary. The number of frames is specified in the parameter *frameQuantity*. We propose to use the same amount of frames for the calculation of an averaged RSSI and noise value as well. The latter helps to reduce the variability of the RSSI especially in mobile environments. Since the algorithm is performed on the basis of averaged information over several frames, the average physical rate must be calculated as well. The number of frames used for the cal-

culation of PDR and average RSSI, noise and PHY rate values determines the readiness of the algorithm, but also affects its accuracy. An appropriate tradeoff between these two has to be found empirically.

4.2 Detection Strategies

For the detection of jamming attacks, several practical implementations are possible. One approach is to perform the detection on the active nodes during their own transmissions. Since these nodes have a different view on the data flow depending on whether they act in the role of the transmitter or receiver, we define two separate algorithms for both cases. We call them "transmitter-based detection" and "receiver-based detection" respectively. However, since typically more than two nodes participate in a network, there are more possibilities where the detection algorithms should run.

The "dedicated jamming detection" is useful in scenarios where the power consumption and device complexity of most of the participating nodes should be low. The detection is then performed by only one or a few nodes having enough resources available. If we recall the scenario introduced in section 1.1.2, the nodes installed on the tanks and other vehicles could be dedicated to the detection of jamming.

Finally, the development of a "cooperative jamming detection" algorithm is motivated by the expected increase of detection performance compared to the standalone detection mechanisms, since a broader view of the network is available.

In the following, each of the four detection strategies "transmitter-based", "receiver-based", "dedicated" and "cooperative detection" are introduced in detail.

4.2.1 Transmitter-Based Detection

In a wireless ad hoc network, the participating nodes communicate directly with other nodes by sending and receiving frames. This means that every node may be in the role of the transmitter and the receiver at the same time. For the illustration of the different detection approaches, let us consider an ad hoc network with node A sending to node B. This scenario is shown in figure 4.3. The additional nodes in this scenario are within range of A and B, as well as the jammer interfering with the regular communication.

This section describes how the jamming detection is performed on the transmitter, i.e. node A. In order to apply the decision algorithm described in the previous section, the transmitter has to determine the four metrics PDR, RSSI, PHY rate and noise corresponding to the communication with node B. While at the noise level at the transmitter is derived from every frame arriving at node A, it is to note that the PDR and RSSI must be calculated on the basis of the communication between A and B only. If node A transfers data to other nodes at the same time, the information from both data streams has to be treated separately. Node A might then run two instances of the jamming detection algorithm in parallel.

Thanks to the positive acknowledgements from node B, the transmitter calculates the corresponding packet delivery ratio as

$$PDR_A = \frac{number\ of\ received\ ACKs\ from\ node\ B}{number\ of\ sent\ data\ frames\ to\ node\ B}$$

It is important to note that node A accounts for every single sent data frame, meaning that retransmissions are included.

The RSSI and PHY rate values are obtained from the ACK frames at node A as well. As mentioned already, the noise level is read from the ACKs or every other frame arriving at A.

After a certain amount of data frames has been sent, A applies the decision algorithm based on the gained statistical data according to section 4.1.3.

The transmitter-based jamming detection offers a realistic perspective on the communication channel between transmitter and receiver. One important ben-

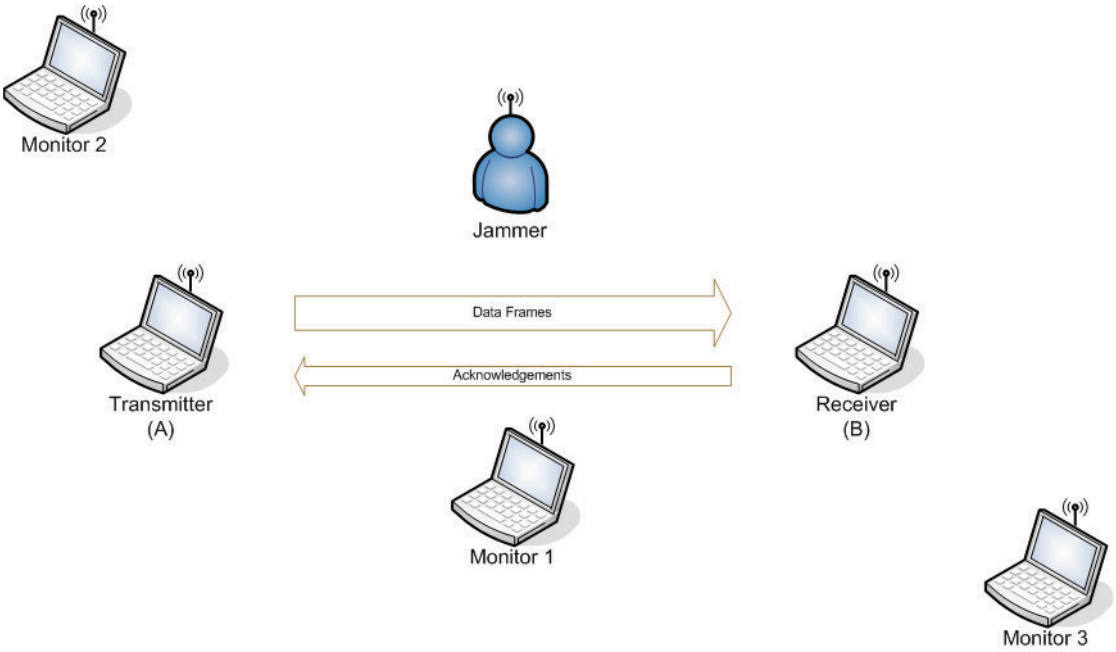

Figure 4.3: Ad hoc network scenario with one data flow between transmitter and receiver. The nodes depicted as monitors may participate in the jamming detection as well. A jammer interferes with the regular communication. Source: by author

efit of this approach is that it may come to a decision even if no frames arrive at the receiver at all. However, since the detection is based on the incoming acknoledgement frames only, it is not possible to make a statement about the nature of occuring interference at the receiver side.

4.2.2 Receiver-Based Detection

Instead of detecting the presence of jamming at the transmitter, we now consider the detection at the receiver. Again, we have the same scenario as illustrated in figure 4.3. The main difference between receiver-based and transmitter-based detection lies in the computation of the packet delivery ratio. While in transmitter-based detection, the transmitter knows the exact number of sent data frames including all retransmissions, this is a priori not known at the receiver since several frames might get lost during transmission. Therefore, it is necessary that the data frames contain additional information which enables the receiver to determine the total number of sent frames.

We achieve this by adding a sequence number to every single data frame. In fact, there is a sequence number assigned to every frame according to the WLAN standard. However, this sequence number remains constant in all retransmissions of the same frame. Since retransmissions are triggered by the WLAN hardware itself (cf. [1], page 291), it is not practical to implement a unique sequence number for every transmitted frame in an upper layer protocol. What could be done nevertheless is the conversion of the present sequence number field of the MAC frame format to be used as a unique counter for each frame including retransmissions. Then, the mechanism for the detection of multiple receptions of the same frame could be implemented either by introducing a separate field to the MAC frame format or by moving this functionality to the upper layers. The latter is technically feasible, since many upper layer protocols implement their own sequence numbers, which makes it easy to identify duplicates.

Using sequence numbers, we calculate the PDR at the receiving node B as

$$PDR_B = \frac{\textit{number of received data frames from } A \textit{ passing the frame check}}{\textit{calculated number of frames transmitted by } A}$$

where the frame check is based on the frame check sequence available in the MAC frame format.

The gathering of RSSI, physical rate and noise level is similar to the transmitter-based approach, except that RSSI and physical rate are collected from A's data frames. The noise level is again read from all arriving frames at node B. The decision algorithm is triggered after the reception of a certain amount of received

data frames or in regular time intervals.

As long as some frames arrive at the receiver, the receiver-based detecion approach allows for a accurate measurement of the communication channel. If no frames are received at the receiver due to extensive jamming, this is detectable by considering the noise metric. However, jamming that prevents the transmitter to access the medium may not be detected at the receiver.

4.2.3 Dedicated Detection

As a third type of jamming detection, we consider the dedicated detection approach. We take the scenario from figure 4.3 again, and let one or more monitor nodes conduct the jamming detection based on their individual view of the network conditions. Since the monitors do not actively participate in the data transfer between nodes A and B, they must try to record the values of the detection metrics in a different way than it is made in the transmitter-based and receiver-based approach respectively. It is important to note that also in the dedicated approach, the detecting nodes must treat each of several concurrent data streams separately. This is the only way for them to calculate the proper PDR values for each of those data flows.

If we put ourselves in the position of one of the monitor nodes, we realize that in principle there are several sources of information we may base our decision on. At first, we receive data frames from node A, containing the sequence numbers which count frame transmission including all the retransmission, as described in the previous section. Since we do probably not receive all of the frames sent by A, this enables us to calculate a packet delivery ratio between A and ourselves:

$$PDR^{*}_{MON} = \frac{number\ of\ received\ data\ frames\ from\ A\ passing\ the\ frame\ check}{calculated\ number\ of\ frames\ transmitted\ by\ A}$$

Finally, we may also calculate a PDR based on the number of sent frames at A and the number of ACKs we received from B:

$$PDR_{MON} = \frac{number\ of\ received\ ACK\ frames\ from\ B}{calculated\ number\ of\ frames\ transmitted\ by\ A}$$

In fact, we do not use the PDR^{*}_{MON} value for the following reason: The selection of the used physical rate for a frame transmission is made on the transmitting node, based on the performance of the communication between transmitter and receiver. Under normal network conditions, there is a certain correlation of RSSI, PDR and physical rate at the receiver. At our monitor nodes, this correlation is not observed for example if the path loss between A and B is much larger than the path loss between A and the dedicated monitor node. Due to the large path loss, a low PHY rate is used even in the absence of jamming. However, the monitor observes

a high RSSI along with the low PHY rate and therefore assumes by mistake that jamming is present.

The problem of wrong RSSI - PHY rate correlation on the monitor is less serious since we use the PDR_{MON} as the measure of the packet delivery ratio. This is due to the fact that acknowledgement frames are always transmitted at a mandatory rate which is the same or lower than the rate of the previous data packet (cf. [1], page 281). The mandatory physical data rates defined in the WLAN standard are 24 Mbps at most, which is lower than the maximum rates achievable for data frames. As a consequence, the dependency from the physical rate is reduced in the algorithm and less false positive decisions are expected.

In case of dedicated detection, the RSSI and PHY rate are read from the acknowledgement frames arriving from the receiver, i.e. node B. As always, the noise level is taken from arbitrary frames arriving at the monitor.

Based on the gathered statistics over several ACK frames, the monitor then applies the decision algorithm. Finally, the node dedicated to the jamming detection announces his decision to the other participating nodes in a broadcast frame. This broadcasting is then repeated whenever the decision changes in future.

The dedicated approach eases the jamming detection a lot, since not every node needs to perform the detection on its own. However, due to its remote location, a dedicated node may be influenced by other or additional sources of interference, which do not impact the transmitter and receiver of the actual data flow. Therefore, the dedicated detection performance is expected to be lower than transmitter-based and receiver-based respectively.

4.2.4 Cooperative Detection

The fourth detection approach is a combination of the former three methods. The idea is to share the information at all nodes among each other and to make a decision based on this broader view. This means that every participating node in the ad hoc network gathers its own information quadrupels containing RSSI, PDR, noise and PHY rate. Since there may be several concurrent data transmissions, these quadrupels are calculated based on the role a node takes on in every of those data flows. According to this, a certain node for example calculates some quadrupels based on the transmitter-based and receiver-based mechanism respectively, depending on whether it is the transmitter or the receiver of this specific data flow. If a node is neather transmitter nor receiver of an observed data flow, it applies the mechanism used in the dedicated detection for the recording of the corresponding information.

In order to share the information among all nodes, every node periodically emits a broadcast frame containing the last few quadrupels based on his observations. The periodicity is chosen so that the generated broadcast traffic does not degenerate and causes congestion in the network itself. This means that the more nodes participate in the network and the higher the average channel usage, the longer the broadcast period must be. This assumes that jamming is not that high so that the broadcast packets are not jammed. Otherwise, this approach degenerates in one of the previous three.

Based on the self-generated quadrupels and those received from the other nodes, every node applies the decision algorithm as described in section 4.1.3.

Using a cooperative detection approach, the amount of information about the ongoing data flows and possible jamming influences is drastically increased. However, its main drawbacks are clearly the generated additional network load and its energy consumption as well as a high device complexity since every participating node has to conduct several measurements and calculations itself. Depending on the actual increase of detection performance by the use of cooperative detection, this drawback could be acceptable however.

5 Experimental Implementation

This chapter provides a deep view into the setup of the experimental testbed and the actual implementation of the developed jamming detection algorithms. The use of commodity wireless network interface cards in our test environment causes some limitations regarding the available operating modes and the device compatibility with certain software tools. As shown in the previous chapter, specific per-frame information is needed in order to apply the detection algorithms. Depending on the actual chipset built in a WLAN card, corresponding device drivers must be used. In practice, it comes apparent that several drivers report some metrics in different ways, or they do not provide some of the needed information at all. Another difficulty comes from the fact that in normal operating mode, WLAN cards do not report all the frames they receive, but only those which are relevant for the upper layer protocols. Although this behaviour may be appropriate for cards solely used as communication devices, it makes it difficult to use them for the detection of jamming.

5.1 Limitations of WLAN Hardware

Since WLAN cards and the corresponding device drivers are primarily designed for the use as communication devices according to the IEEE 802.11 standard, it is not trivial to use the same hardware for certain measurements of the communication channel.

5.1.1 Operating Modes

First of all, it is difficult to use a commodity WLAN card for the network communication while at the same time collecting information for the jamming detection. This limitation derives from the fact that in the normal operating mode, WLAN cards report only the relevant data to the operating system. Management and control frames as well as frames belonging to other basic service sets are not reported at all. Data frames associated with the same BSS are translated so that they appear to the operating system as normal ethernet frames. In this translation, however, all the wireless-specific information provided by the PLCP sublayer is being lost. For example, RSSI and noise level may not be retrieved on a per-frame basis in this operating mode. Furthermore, since control frames are hidden, it is also not possible to observe and count acknowledgement frames.

Different operating systems allow WLAN devices to use a special operating mode, called "Monitor Mode". If a card operates in this mode, it reports every detected frame including all corresponding information, even for corrupted frames showing a wrong frame check sequence. However, since this mode is intended for passive use only, some cards can not actively participate in the communication as long as they are in monitor mode. Since there are no specifications for normal or monitor mode contained in the IEEE 802.11-2007 standard, there are several different implementations available today, depending on the exact chipset and operating system.

5.1.2 Link Metrics

Another difficulty arises from the way, how several link metrics are reported by the WLAN cards. The IEEE 802.11-2007 standard defines only rough service primitives describing what information should be reported by the device (cf. [1], clauses

17.5.5, 18.4.5 and 19.9.5). This leaves room for manufacturers to realize their own implementations. However, it complicates the experimental implementation of our jamming detection algorithms, since they are highly dependent on values reported directly by the device.

RSSI

The energy measured during the PLCP preamble of an arriving frame is expressed as an 8-bit integer value which allows for 256 levels. Some device manufacturers use only 100 levels and express the RSSI as a percentage. Other manufacturers define tables for the mapping between certain RSSI values and the corresponding absolute power levels in Watts or dBm. Even others use formulas which allow the conversion between RSSI values and dBm.

Our WLAN cards are based on Atheros chipsets, where the power level in dBm may be derieved from the RSSI value by subtracting 95. As the maximum RSSI value defined by Atheros is 60, this gives a dBm range of -35 dBm at 100% and -95 dBm at 0% of the detectable signal strength [15].

Noise

While the IEEE 802.11-2007 standard intends that WLAN cards report a certain signal quality measure for DSSS PHYs[1] only, reporting of the interfering radio frequency energy is not mandatory.

The Atheros based WLAN cards used in our experiments report the sum of the interfering energy in the same band plus the noise floor (-95 dBm) inside the radio as the noise value.

PHY Rate

The physical rates are reported on a per frame level according to [1]. The rates are expressed in Mbps.

[1]The incoming bitstream is compared at the PLCP layer with the used spreading sequence. The signal quality is the percentage rate of correct bits per sequence or, in other words, the pseudo noise correlation strength (cf. [1], page 561). This measure is not available in OFDM PHYs.

PDR

The packet delivery ratio may not be expressed on a per frame level and it is not intended to be used as a link metric in the most recent WLAN standard. Therefore, it has to be computed manually based on the number of observed frames of a specific type as well as on their origin and destination. Since the computation of the PDR differs among the deployed jamming detection algorithms, it will be explained in more detailed in the corresponding sub-clause of the next section.

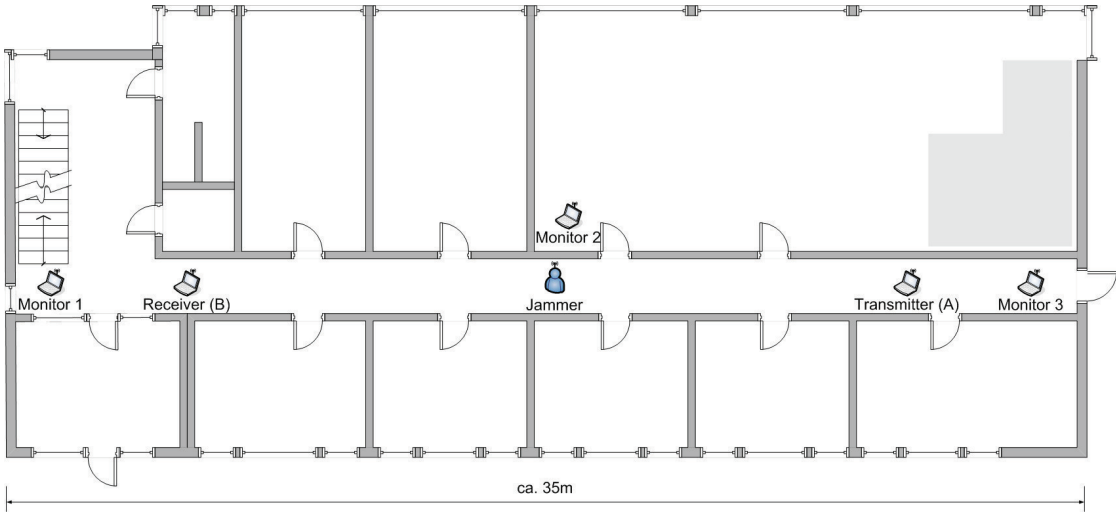

Figure 5.1: Floor plan and basic configuration of the indoor wireless ad hoc network. Source: by author

5.2 Testbed Setup

The core contribution of this thesis is a comprehensive investigation carried out in a real-world system. First of all, we analyzed the influence of jamming on the network performance with respect to the position and transmission power of the jammer. At a later stage, we performed an empirical evaluation of the developed jamming detection algorithms. For this purpose, an indoor ad hoc network was built using IEEE 802.11-2007 compliant wireless network interface cards. Figure 5.1 shows the floor plan of the building where we conducted the experiments. Several network nodes as well as the jammer are depicted in a basic configuration, where the jammer is positioned between transmitter and receiver. Additional monitor nodes are placed near the transmitter, receiver and the jammer. This basic configuration was then modified in order to meet the specific demands of every experiment, i.e., the position of the nodes is not fixed.

5.2.1 Hardware

Due to the limitations of our commodity WLAN cards (cf. section 5.1), it is not sufficient to have one single network interface card per node if regular data transfer should be possible at the same time as running the jamming detection algorithms.

Figure 5.2: ZyXEL ZyAir G-110 card used at transmitter. Source: courtesy of Studerus AG, Switzerland

Moreover, cards which allow for the connection of an external antenna are needed at the transmitter, which we will explain in section 5.2.3. Since WLAN cards are not able to generate arbitrary radio signals but only frames according to the underlying standard, a special device is needed for the implementation of the jamming attacks. The detailed configuration of each node types and the jammer will be discussed in sections 5.2.3 - 5.2.6. In the following we just mention the used hardware itself.

Laptops

For the implementation of the network nodes we used laptops by Dell and Compaq running Windows XP. All laptops have a 32-bit PC card slot which we used to insert the WLAN cards. There are no further special requirements to the laptops.

ZyXEL ZyAIR G-110 Wireless Network Interface Card

This ZyXEL card is a 32-bit CardBus device, compliant with the IEEE 802.11-2007 standard, which allows for the connection of an external antenna (cf. [16]). The card is based on a Prism54 chipset by Intersil[2]. Figure 5.2 shows the WLAN card and figure 5.3 shows an antenna similar to the one used in our experiments.

[2]http://www.intersil.com

Figure 5.3: Omnidirectional antenna as used at the transitter and jammer. The exact type and manufacturer are unknown. Source: courtesy of Easy-Tecs GmbH, Germany

Figure 5.4: Netgear WG511T card used at transmitter, receiver and monitors. Source: courtesy of Netgear Switzerland GmbH

Netgear WG511T Wireless Network Interface Card

This Netgear card shown in figure 5.4 is a 32-bit CardBus device as well (cf. [17]). Like the ZyXEL card, the Netgear card is compliant with the IEEE 802.11-2007 standard, but it has a built-in patch antenna only. The card is based on a AR5001 chipset by Atheros[3].

R&S Vector Signal Generator SMU200A

For the generation of jamming signals, the vector signal generator by Rohde & Schwarz[4] shown in figure 5.5 was used (cf. [18]). It offers the generation of arbitrary waveforms in the frequency range from 100 kHz to 6 GHz. By the installed software, several predefined and user configurable signal types may be selected such as noise, Bluetooth, WLAN, GSM etc. Within WLAN the modulation scheme,

[3]http://www.atheros.com
[4]http://www.rohde-schwarz.com

Figure 5.5: R&S Vector Signal Generator SMU200A, used to generate the jamming signals. Source: courtesy of Roschi Rohde & Schwarz AG, Switzerland

frame type and payload may be chosen arbitrarily. Along with the transmission of WLAN frames in fixed time intervals it is also possible to send in an unframed mode, meaning that random bits are sent using the selected modulation. The SMU200A software contains a pseudorandom number generator which may be used for the generation of random payload and random bit sequences.

5.2.2 Software

In our experimental implementation we distinguish between data acquisition and data processing. Data acquisition is performed on the participating nodes during the experiment, meaning that RSSI, PHY rate and noise values are recorded on a per frame basis. Data processing follows after the experimental phase has ended. It includes the offline calculation of the PDR values and the execution of the decision algorithm based on the recorded data. The separation of the data acquisition and data processing tasks allows us to study the impact of different parameters on the outcome of the jamming detection. For these two tasks, we use specialized software tools as described below.

Another important task which requires the use of a software tool, is the generation of network traffic. The open source project "Iperf"[5] is a command line tool which allows for the easy generation of both UDP and TCP traffic with user-defined characteristics. Since Iperf is widely-used and available for the Windows operating system as well, we decided to use it as our traffic generator.

[5]http://sourceforge.net/projects/iperf

Data Acquisition

For the task of data acquisition, several operating systems and software products have been evaluated. Since newer versions may eliminate weaknesses of the tested products in future, there might be alternatives to the choices we have made.

The network analyzing tool "Omnipeek" by WildPackets[6] proved as a good choice in order to read out the required values per frame. The software suite runs on Windows XP and includes special device drivers for several WLAN chipsets by different manufacturers. Using these drivers, the Netgear cards report RSSI, PHY rate and noise values as expected. There is only limited support for the Intersil chipset of the ZyXEL cards, allowing them to report RSSI and PHY rate only.

Data Processing

The task of data processing includes two phases. In the first phase, the PDR and averaged RSSI, noise and PHY rate values over a series of frames are calculated from the result of the data acquisition by Omnipeek. The second phase of data processing finally applies the jamming detection algorithms described in section 4.1.3. For the data processing, we developed two scripts written in Perl and Python respectively.

The Perl script simply transforms the exported Omnipeek data to a new raw data file format. During this transformation, all data which is not relevant for the jamming detection is removed. This separate preprocessing eases the adaption to upcoming versions of Omnipeek or to other software products using different export file formats.

As soon as the raw data files are available, the decision algorithm written in Python is initated. It is important to note that in order to run the decision algorithm, we need the information from two separate raw data files. This is, the raw data file of the node where the algorithm should be run, and additionally the raw data file containing the timestamps of all data frames sent by the transmitter. This workaround is necessary for the computation of the PDR values, since we assume according to section 4.2.2 that every data frame has its own sequence number, which is unique even among all retransmissions of the same frame. As this

[6] www.wildpackets.com

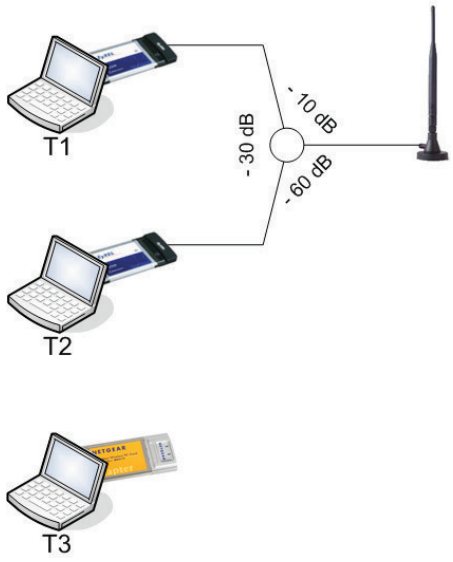

Figure 5.6: Experimental implementation of the transmitter node. Source: by author

assumption is not true in our testbed (frame sequence numbers remain the same among retransmissions of the same frame), the receiver and monitor nodes would not be able to calculate the correct PDR values. Therefore, we have to provide this information to the decision algorithm manually, by including the transmitter's frame timestamps.

For the evaluation of the cooperative detection, it is possible to pass the raw data files from all participating nodes to the decision algorithm in parallel. In this case, prior to the execution of the real decision algorithm, all the data tuples consisting of timestamp, RSSI, PDR, PHY rate and noise are sorted in chronological order. There are no further changes to the decision algorithm needed for the cooperative detection.

The subsequent sections explain the implementation of our transmitter, receiver and monitor models and show how they practically perform the jamming detection. Further details about the implementation and handling of the data processing scripts may be found in appendix A.4.

5.2.3 Transmitter Model

Compared with the other node types, the implementation of the transmitting node is the most complex. This becomes clear if we consider the tasks which are up to the transmitter in our setup:

- recording the timestamp of every sent data frame by the transmitter
- recording detailed information about the incoming ACK frames from the receiver

To make clear the reason why we need three separate WLAN cards, we may look at figure 5.6 where the implementation of the transmitting node is shown. The transmission unit consists of three laptops denoted as *T1*, *T2* and *T3*. While *T1* and *T2* use a ZyXEL card each, *T3* is equipped with a Netgear card. The laptops *T1* and *T2* are assigned to generate data traffic and to record the corresponding statistics. In the following, we assume *T1* to emit the data frames generated using Iperf. Since Omnipeek is unable to capture retransmitted frames if run directly on *T1*, we need the second laptop *T2*, which is in monitor mode and records the timestamp of every sent data frame by *T1*. Since external influences like jamming could disturb the reception at *T2*, a direct connection between *T1* and *T2* is preferable. Therefore we use the ZyXEL cards which allow the wired connection of the two cards in *T1* and *T2*. By the use of a radio frequency splitter we may also connect the external antenna. The splitter is inserted in the way that its attenuation helps to further decrease the impact of jamming to the frame detection at *T2*. As shown in figure 5.6, the splitter has different attenuation on the three loops. While we use the smallest (-10 dB) between *T1* and the antenna, the largest attenuation (-60 dB) is applied between antenna and *T2*. The attenuation between *T1* and *T2* finally is small enough (-30 dB) to guarantee a proper reception of *T1*'s data frames at *T2*.

The third WLAN card is used to capture the acknowledgement frames from the receiver. A separate device for this purpose is needed for two reasons. First, there are no device drivers available which allow the ZyXEL cards to fully cooperate with Omnipeek. Instead of a real noise level, a constant value is reported due to this incompatibility. However, there is a second motivation for the use of the additional Netgear card: Even if the noise levels reported by the ZyXEL

card *T2* were correct, they must have been handled with care due to the strong attenuation between the antenna and *T2*. It is likely that several acknowledgements would have been lost due to this attenuation yet. By the use of the separate Netgear card, these difficulties may be overcome. As a further advantage of this approach, we may apply the same device-dependent parameters in the decision algorithm for the transmitter as we use at the receiver and all the monitor nodes. The device-dependent parameters concern the expected distribution of points in the RSSI versus PDR diagram in the absence of jamming as discussed in section 4.1.1.

Since we use more than one WLAN card at the transmitter, it is important that they are positioned very close to each other. In several experiments we verified that the transmitting and receiving conditions of all cards are then almost equal.

To sum up, we assign the following tasks to each component at the transmitter:

- *T1*: transmission of data frames
- *T2*: recording statistics about the transmitted data frames by *T1*
- *T3*: recording statistics about the incoming acknowledgement frames

The information which is retrieved by *T2*, namely the timestamps of all the sent data frames, is used by all nodes for the offline computation of their local PDR values. We call this information the "reference frame data". In section 4.2.2, we described what changes would be needed in order that every node could determine this information itself.

For the send-site jamming detection, the data retrieved by *T3* is passed to the decision algorithm along with the reference frame data provided by *T2*.

5.2.4 Receiver Model

The implementation of the receiving node consists of the two laptops *R1* and *R2* as illustrated in figure 5.7. The following tasks are assigned to *R1* and *R2* respectively:

- *R1*: reception of data frames and transmission of acknowledgement frames

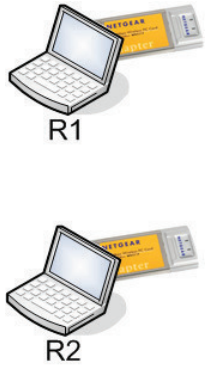

Figure 5.7: Experimental implementation of the receiver node. Source: by author

Figure 5.8: Experimental implementation of the monitor nodes. Source: by author

- *R2*: recording statistics about the incoming data frames

The laptop *R1* is running Iperf in the server mode, meaning that the traffic from *T1* is beeing acknowledged at the upper layer protocol if needed (e.g. TCP data flow). *R2* runs Omnipeek in the monitor mode and captures the incoming data frames from *T1*. Since both cards are positioned next to each other, we make sure that their transmitting and receiving conditions are almost equal.

For the receive-site jamming detection, the data retrieved by *R2* is passed to the decision algorithm along with the reference frame data provided by *T2*.

5.2.5 Monitor Model

The monitor nodes are implemented by one single laptop and a Netgear WLAN card. This is sufficient, since their role is just passively analyzing the ongoing data transfer. Each of those Laptops *M* running Omnipeek record statistics about the acknowledgement frames received from *R1*.

For the monitor based jamming detection, the data retrieved by *M* is passed to the decision algorithm along with the reference frame data provided by *T2*.

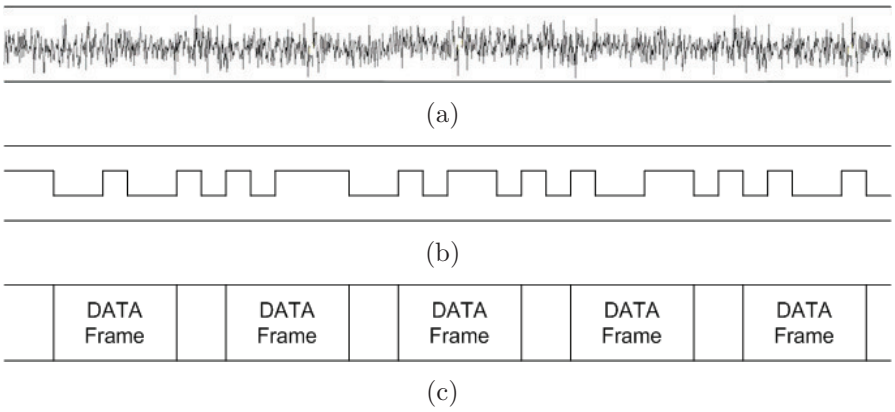

Figure 5.9: Jamming models used in the experiments: (a) noise jamming, (b) bit jamming, (c) frame Jamming. Source: by author

5.2.6 Jamming Models

In the following, we discuss the properties of the jamming models used in our experiments. Using the SMU200A vector signal generator by Rohde & Schwarz, three jammers were defined and used to evaluate the performance of the detection algorithms. According to the classification in chapter 3, we implemented three constant jammers based on the noise, bit and frame mechanisms respectively.

Noise Jamming

The channel bandwith used by the targeted system is jammed with noise energy. This raises the level of background noise at the receiver and makes it difficult to detect frames correctly. In other words, the signal-to-noise ratio at the receiver is decreased.

White noise with a bandwidth of 20 MHz (= channel width of WLAN) is generated around the center frequency of the used WLAN channel. The jamming signal is emitted at a constant power level.

Bit Jamming

Jamming using the same frequency and modulation scheme as the targeted system seriously decreases the network performance as the devices try to detect a known pattern in the bitstream allowing them to synchronize. Since this modulated signal

may not be filtered out like white noise, it decreases the signal-to-noise ratio at the receiver and occupies the channel heavily.

A pseudorandom binary sequence (PRBS 9[7]) is used in order to generate a random signal on the physical layer. We use the 11 Mbps data rate with complementary code keying as the scheme for the transmission of bits. The bitstream is emitted at a constant power level.

Frame Jamming

Jamming using frames according to the targeted system is hard to detect, since the jamming signal is masked as regular frames. Its impact goes beyond minimizing the signal-to-noise ratio. Due to the unfairness of the jammer, the channel may be occupied over long periods of time. Depending on the system, this might be achieved with very low energy consumption by periodically announcing long duration frames which makes the participating to remain in silent mode this amount of time.

We generate data frames containing a pseudo random payload of 1024 bytes. For the transmission on the physical layer, we use the 11 Mbps data rate with complementary code keying as in the case of bit jamming. The frames are sent in regular intervals where the silence between two frames is 0.1 milliseconds. The transmission of one frame with the specified properties takes about 25 milliseconds.

[7] cf. http://en.wikipedia.org/wiki/PRBS

6 Performance Evaluation

This chapter evaluates the performance of the proposed jamming detection algorithms. For this purpose, extensive investigations have been made using our experimental indoor testbed, based on the WLAN standard. The configuration of this ad hoc network was shown in figure 5.1. Based on this setup, experiments were conducted in order to show how several factors influence the performance of the jamming detection. Together with the detection of different jammer models, we investigated the following effects to the detection performance:

- the jammer's position
- different upper layer data flows (UDP, TCP)
- mobility in the network
- background traffic

In order to define reasonable experimental test cases, prior analysis was made regarding the impact of jamming attacks to a real-world network. Along with the effectiveness of different jamming mechanisms, we also investigated on the relation between the position of a jammer and its impact on the network performance.

6.1 Impact of Jamming on the Network Performance

The impact of jamming attacks on the network performance varies among the different jamming models we deployed. As our results show, not only the jamming model, but also the position of the jammer influences the maximum achievable throughput of regular nodes communicating in the network.

6.1.1 Effectiveness of Different Jamming Mechanisms

We start by analyzing a scenario with one UDP data flow between a transmitter and a receiver, the impact of noise jamming, bit jamming and frame jamming. The nominal data rate of the UDP stream was set to 20 Mbps, since this equals approximately the maximum achievable throughput in a 54-Mbps WLAN network. However, since the achieved real throughput in this scenario did not exceed 18 Mbps even in the absence of jamming, we take this as the 100% mark. The distance between transmitter and receiver was 32 meters without any obstacles in the line of sight. The jammer was positioned between the transmitter and receiver, at a distance of 22 meters from the transmitter and 10 meters from the receiver respectively. Figure 6.1 shows the achieved effective throughput in the presence of different jamming at varying power levels. The Signal-to-Interference-plus-Noise Ratio (SINR) on the horizontal axis describes the difference between the received signal strength and the measured interfering energy including internal termal noise at the receiver. Just like the SNR, also the SINR is expressed in dB as well. From the figure, it becomes clear that a noise jammer must jam with significant higher power to cause the same decrease in network performance compared with the bit and frame jammers. This is explainable by the ability of WLAN cards to filter out white noise, which is not possible for bit and frame jamming. Since the spoofed frames of the frame jammer contain a random duration field, the regular transmitter is prevented from accessing the channel over long time intervals, which additionally decreases the network throughput. From the figure, it may be seen that incoming jamming signals do not affect the network performance, if the SINR reaches 40 dB. On the other side, a SINR at the receiver of less than 25.5 dB in case of frame jamming and less than 12.4 dB in case of bit jamming makes the

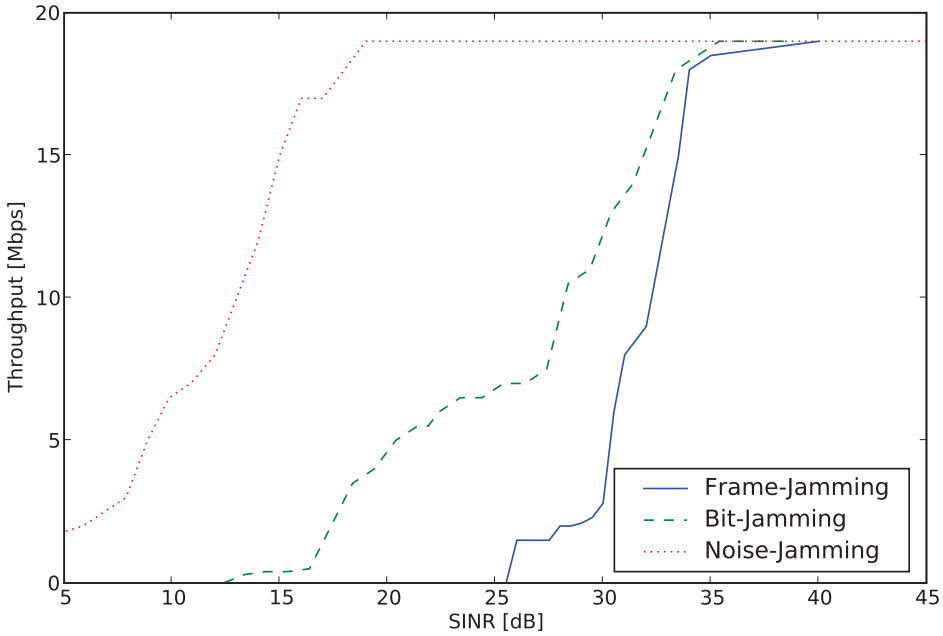

Figure 6.1: SINR vs. throughput for different jamming mechanisms.

communication impossible. Due to limited output power of the signal generator, we were not able to fully break the connection using noise jamming.

6.1.2 Jammer Position vs. Network Performance

We took a large serie of measurements in order to investigate the impact of the jammer position on the network performance. These experiments were made separately in the presence of noise jamming, bit jamming and frame jamming. The transmitter and receiver were placed at a distance of 33 meters without any obstacles in between, and a UDP stream of nominal 20 Mbps was sent in one direction. The jammer was moved in discrete intervals from the transmitter to the receiver, while applying each one of the three mechanisms at 5 distinct power levels consecutively. The effective throughput measured at the receiver is shown in figures 6.2, 6.3 and 6.4.

As expected, the plots for the three jamming mechanisms differ. However, they have a basic structure in common. At first, it may be seen that with increasing

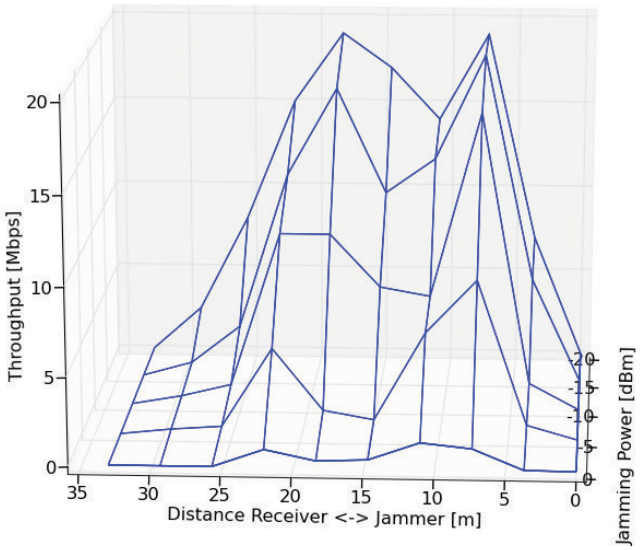

Figure 6.2: Influence of noise jamming depending on the jammer's location. The Receiver and transmitter are fixed at 0 and 35 meters respectively.

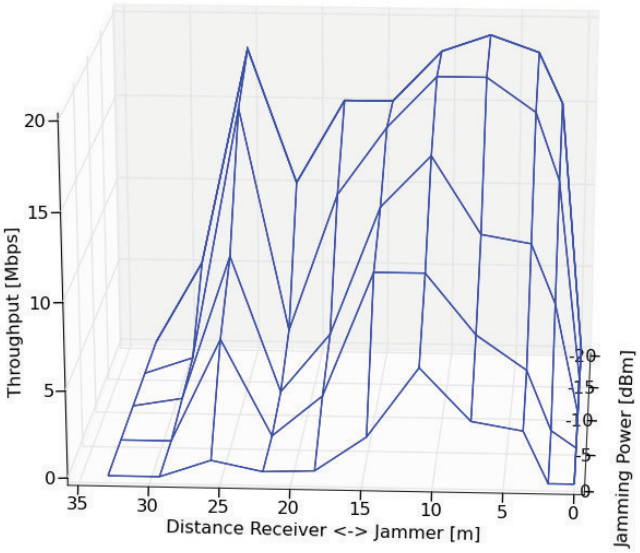

Figure 6.3: Influence of bit jamming depending on the jammer's location. The Receiver and transmitter are fixed at 0 and 35 meters respectively.

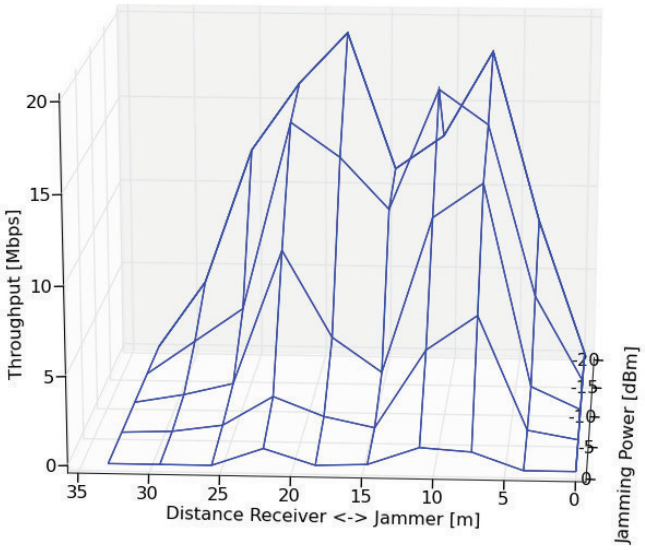

Figure 6.4: Influence of frame jamming depending on the jammer's location. The Receiver and transmitter are fixed at 0 and 35 meters respectively.

jamming power, the throughput decreases. This corresponds to the behaviour discussed above. Another commonality among the jamming mechanisms is the qualitative distribution of their impact on the network performance, depending on the relative position between transmitter and receiver. If the jammer is close to the transmitter or the receiver, the communication degrades completely. As expected, the throughput increases with increasing distance between the jammer and either of the network nodes. However, the throughput decreases again when the jammer is positioned around 22 meters away from the receiver in case of noise and bit jamming, and 15 meters in case of frame jamming. In the following, we call this particular place *position 15/22*.

This somewhat surprising behaviour can be explained with the rate switching algorithm used at our transmitter WLAN card: With the first try, a frame is sent with the highest PHY rate, and all necessary retransmissions with the next lower rate. For the next frame, the card tries to send at the higest rate again. If a certain amount of frames was reliably received at low rates but not on higher rates, the following transmissions will start directly at this low rate. If packets are lost at

all rates equally, the rate switching algorithm regularly tries the higher rates as well. If we consider our experiments with the jammer at different positions, we can conclude the following:

Jammer next to the transmitter Due to the high energy of the jammer, the channel appears busy all the time. No frames are sent by the transmitter.

Jammer near the transmitter The energy of the jammer is low enough, so that frames are sent. Most frames arrive at the receiver, therefore a high PHY rate is chosen.

Jammer at *position 15/22* The jamming interference at the receiver enables the receiver to detect frames at lower rates reliably while frames sent at higher rates are lost more often. As a result, the rate switches to those lower rates. This guarantees the connection, but at the price of a lower throughput.

Jammer near the receiver The SINR at the receiver is now smaller and the detection is not reliable at low rates anymore. This lets the transmitter try the higher PHY rates again, which effectively allow higher throughput under certain circumstances, as shown in the previous section. The effect is a less stable, but faster connection. Thanks to the positive acknowledgement system of WLAN, upper layer protocols do not notice the decreased PDR on the physical layer.

Jammer next to the receiver The sensitivity of the receiver card is disturbed due to the high energy coming from the jammer. No frames are detected.

6.2 Jamming Detection Performance

All the experiments discussed below were conducted using our testbed as shown in figure 5.1. A UDP data stream at a nominal rate of 20 Mbps was used where not stated differently. The ad hoc network ran on channel 11 (2.462 GHz), which was not used by other networks or communication systems during our experiments. Except for the measurements regarding the effect of mobility, all network nodes as well as the jammer are fixed. The transmitter and receiver are placed at a distance of 28 meters from each other. Monitor 1 and 3 are near the receiver and transmitter respectively and monitor 2 is placed centered between receiver and transmitter.

Based on our investigations regarding the effectiveness of jamming, for the evaluation of the jamming detection performance we adjust the jamming power so that a significant impact on the network performance is achieved. This means that the maximum throughput is decreased by approximately 80% due to the jamming..

Parameters According to the explanations in section 4.1.3, several parameters of the decision algorithm need to be specified. The empirical values used in our experiments are:

- $frameQuantity = 100$
- $noiseThresh = -85\ dBm$
- $maxPDR = 90\%$

As the reference model for the RSSI vs. PDR and PHY rate - classifier, we use linear approximations for six ranges of the PHY rate. They were retrieved based on prior measurements by linear regression. This is a popular form of describing empirical data in a closed formula (cf. [19], pages 9, 466 et seqq.). The six ranges of averaged PHY rates are:

- PHY rate $<$ 5 Mbps
- 5 Mbps $<=$ PHY rate $<$ 10 Mbps
- 10 Mbps $<=$ PHY rate $<$ 20 Mbps

- 20 Mbps <= PHY rate < 30 Mbps
- 30 Mbps <= PHY rate < 40 Mbps
- 40 Mbps <= PHY rate

We present most of our results in table form below. The first column specifies the algorithm, which achieved a certain detection accuracy.

Correct Decisions The "correct decisions" are defined as

$$\frac{\textit{number or correct decisions during the experiment}}{\textit{total number of decisions during the experiment}}$$

Since each node bases the decisions on its own information, the total number of decisions per experiment varies among the nodes. This complicates the comparison of the detection performance. If we consider for example the fictive nodes X and Y, having correct decisions of $\frac{20}{21}$ and $\frac{18}{19}$, we would say that based on their available information, the detection performance is equal. However, if we calculate the corresponding ratios, we find a difference of 0.5% ($\frac{20}{21} = 95.2\%$; $\frac{18}{19} = 94.7\%$). Therefore, we always include the correct decisions and not just the rate, whenever we compare results among each other.

CER & CER Parameter The Crossover Error Rate is a well-known measure for the assessment of a classifier. There are two types of errors which occur in the jamming detection. If the classifier decides "jamming" in a situation where no jamming is active, this is called a false-positive error. On the other hand, if the presence of jamming is not detected, this is a false-negative error. By varying the CER parameter, in our case this is $rssiDiffThreshold$ (cf. section 4.1.3), the algorithms may be tuned in order to produce more or less of these errors. The crossover error rate is defined as the error ratio at the intersection of the false-positive and false-negative error curves.

The duration of the measurements was chosen in order to generate representative results. We found that the 95% confidence interval is reasonable small ($< 0.1\%$) when the statical measurements are conducted over one minute and the measurements with mobility over at least four minutes. This means that of 100

similar measurements, 95 would produce a detection performance which is within this confidence interval around the empirical value found.

6.2.1 Effect of Jammer Position

Tables 6.1 - 6.3 show the jamming detection performance when the frame jammer is placed at the centre between transmitter and receiver, near the receiver and near the transmitter respectively. We clearly see that the transmitter-based, receiver-based and dedicated detection perform on an equal high accuracy level. Every single detection makes only one false classification during one minute of measuring. The performance of a cooperative detection among the transmitter and receiver, and the cooperative detection among all nodes are given as well. Since the optimal CER parameter differs among the nodes, there are additional false classifications when a common parameter is used for all the data, as it is the case in the cooperative detection. As we can see, the CER parameter in the cooperative case is smaller than the value used at the transmitter-based, but higher than the value used at the receiver. On the other hand, the availability of additional data may also lead to a better numerical performance as in experiment 1, using the cooperative detection among all nodes. The measured CER is lower than every CER of the self-contained approaches, simply due to the fact that the amount of available data is larger.

The results for the noise and bit jamming are summarized as follows: If the amount of "normal" radio frequency interference present in a specific location is known, then an appropriate noise threshold may be defined. Since the regular interference might change over time, this parameter should be somewhat adaptive. I a noise threshold is defined, the detection of noise and bit jamming is achievable with a zero error quota. If no noise threshold is defined, the jamming may still be detected based on the RSSI vs. PDR and PHY rate criterion. In this latter case, the impact of noise and bit jamming is comparable to that of frame jamming, namely preventing the transmitter from accessing the channel and lowering the signal-to-noise at the receiver. This then results in a lower PDR and PHY rate at the given RSSI level, just like in the case of frame jamming.

Since in our scenario, all nodes are within the transmission area of the jammer

and the noise parameter is set according to the known regular interference level, we achieve a perfect jamming detection at all nodes.

Algorithm	correct decisions	CER	CER parameter
Transmitter-based	50 / 51	1.96 %	13.390 dB
Receiver-based	50 / 51	1.96 %	14.740 dB
Dedicated (Monitor 1)	73 / 74	1.35 %	19.415 dB
Dedicated (Monitor 2)	35 / 36	2.78 %	8.328 dB
Dedicated (Monitor 3)	52 / 53	1.89 %	6.699 dB
Cooperative	247 / 265	0.76 %	12.066 dB

Table 6.1: Experiment 1: Jammer between transmitter and receiver. Frame jamming, UDP data traffic @ 20 Mbps, duration: 1'00.

Algorithm	correct decisions	CER	CER parameter
Transmitter-based	43 / 44	2.273 %	5.384 dB
Receiver-based	53 / 54	1.85 %	3.510 dB
Dedicated (Monitor 1)	25 / 26	3.85 %	4.948 dB
Dedicated (Monitor 2)	33 / 34	2.94 %	2.147 dB
Dedicated (Monitor 3)	39 / 40	2.50 %	4.958 dB
Cooperative	191 / 198	3.54 %	4.816 dB

Table 6.2: Experiment 2: Jammer near the receiver. Frame jamming, UDP data traffic @ 20 Mbps, duration: 1'00.

6.2.2 Effect of TCP vs. UDP

In order to check whether the performance of our detection algorithms depends on the characteristic of the used upper layer protocol, we also used a TCP instead of UDP data flow. The results are given in table 6.4, and they show no significant difference to experiment 1. Both of these measurements were conducted with the frame jammer placed centered between transmitter and receiver. In experiment 4 using TCP, we have less data for the detection as compared with the UDP experiment, as seen from the "correct decision" column. This comes from the fact that the TCP protocol implements a back-off mechanism which interrupts the data flow for a certain amount of time after extensive packet losses are detected. Since the number of total decisions is lower in case of TCP, the CER values are slightly

Algorithm	correct decisions	CER	CER parameter
Transmitter-based	39 / 40	2.50 %	16.553 dB
Receiver-based	33 / 34	2.94 %	13.338 dB
Dedicated (Monitor 1)	35 / 36	2.78 %	17.512 dB
Dedicated (Monitor 2)	32 / 33	3.03 %	18.310 dB
Dedicated (Monitor 3)	29 / 30	3.33 %	21.900 dB
Cooperative	172 / 173	0.58 %	14.768 dB

Table 6.3: Experiment 3: Jammer near the transmitter. Frame jamming, UDP data traffic @ 20 Mbps, duration: 1'00.

increased. However, if we look at the number of correct decisions, we notice that there is still just one false classification over the duration of one minute. The higher CER in the cooperative detection among all nodes is explained by the same arguments as in the case of UDP.

For the noise and bit jammer, the results are the same when using TCP or UDP, meaning that the detection performance is at 100% for all nodes within the jammer's transmission area.

From this, we conclude that the used upper layer protocol does not influence the performance of our jamming detection algorithms.

Algorithm	correct decisions	CER	CER parameter
Transmitter-based	26 / 27	3.70 %	8.148 dB
Receiver-based	32 / 33	3.03 %	6.822 dB
Dedicated (Monitor 1)	30 / 31	3.22 %	6.851 dB
Dedicated (Monitor 2)	27 / 28	3.57 %	3.291 dB
Dedicated (Monitor 3)	27 / 28	3.57 %	5.835 dB
Cooperative	143 / 147	5.88 %	6.555 dB

Table 6.4: Experiment 4: Jammer between transmitter and receiver. Frame jamming, TCP data traffic, duration: 1'00.

6.2.3 Effect of Mobility

Several experiments were conducted to investigate the effect of mobile network nodes to the jamming detection performance. The setup was the same as in the

case of experiment 1, except that the receiver was not fixed. Instead, the receiver was carried around at walking pace in the building, covering the whole transmission area of the transmitting node. We tested this scenario using frame, bit and noise jamming. The results are given in tables 6.5, 6.6 and 6.7 respectively. The scenario

Algorithm	correct decisions	CER	CER parameter
Transmitter-based	260 / 308	15.58 %	18.445 dB
Receiver-based	402 / 432	6.94 %	7.035 dB
Dedicated (Monitor 1)	268 / 357	24.93 %	4.595 dB
Dedicated (Monitor 2)	307 / 362	15.19 %	15.402 dB
Dedicated (Monitor 3)	358 / 421	14.96 %	13.485 dB
Cooperative	1519 / 1880	19.20 %	11.899 dB

Table 6.5: Experiment 5: Mobile receiver. Frame jamming, UDP data traffic @ 20 Mbps, duration: 9'30.

with a mobile receiver and frame jamming shows an increased crossover error rate at all nodes. This general trend comes from the increased packet loss due to the mobility, even in the absence of jamming. Therefore, it is more difficult to distinguish between jammed and unjammed environments. Reasons for the lower network performance in case of mobility are sudden drops of the received signal strength due to varying fading effects, for example caused by obstacles coming in the line of sight.

However, we also realize significant differences among the performance of the detection algorithms in the mobile scenario with frame jamming.

While the receiver-based detection reaches a CER of 6.94%, the CER of the transmitter-based detection increases to 15.58%. If we consider the performance of the dedicated detection at the three monitor nodes, we realize that the performance is better, the nearer the node is positioned to the transmitter.

We conclude that mobility makes the detection of frame jamming more difficult. The best performance is achieved at the receiver, while the transmitter-based detection is significantly less accurate. The performance of the dedicated detection depends on the position of the node, meaning that for nodes far away from the transmitter the performance is decreased.

For mobile scenarios, it is clearly shown that the cooperative detection among

all nodes gives a bad performance in general. To achieve a better performance using the cooperative jamming detection algorithm, it would be necessary to discriminate the monitoring nodes which do not report relevant information due to their location relative to the actual data stream.

The case of noise and bit jamming is basically the same as in the static measurements. However, since the receiver is carried out of the jammer's range, not all detected frames show a noise level which is higher than the defined threshold value. Therefore, some false classifications occur at these locations. Furthermore, since in our setup this falls together with a very low received signal strength from the transmitter, the detection is difficult anyway as long as the receiver is far away from the transmitter.

As all nodes except the receiver have a 100% rate of correct decisions in case of bit and noise jamming, tables 6.6 and 6.7 contain only the results of the receiver-based as well as the cooperative detection.

Since most of the time during the experiments, the receiver is inside the transmission range of the jammer, the performance of the receiver-based detection is higher than in the case of frame jamming. Since the information available by the transmitter and the monitor nodes allows for an assured classification, in case of noise and bit jamming the performance of the cooperative algorithm is higher than the self-contained detection at the receiver.

Algorithm	correct decisions	CER	CER parameter
Receiver-based	528 / 552	4.35 %	10.704 dB
Cooperative	2505 / 2544	1.53 %	25.013 dB

Table 6.6: Experiment 6: Mobile receiver. Bit jamming, UDP data traffic @ 20 Mbps, duration: 4'35.

Algorithm	correct decisions	CER	CER parameter
Receiver-based	450 / 459	1.96 %	16.200 dB
Cooperative	1693 / 1702	0.53 %	15.696 dB

Table 6.7: Experiment 7: Mobile receiver. Noise jamming, UDP data traffic @ 20 Mbps, duration: 4'35.

6.2.4 Effect of Background Traffic

Since in real-world applications it is rather unrealistic that there is only one ongoing data transfer at a time, we investigated the effect of "background traffic" in the network. We differentiate between background traffic which is in the same collision domain from background traffic generated by hidden nodes. In case of background traffic in the same collision domain, all participating nodes are able to minimize the number of occuring collisions using their collision avoidance mechanism. Background traffic generated by hidden terminals means that the transmitting nodes are not aware of each other, which causes collisions at the receiver. We found that background traffic generated in the same collision domain does not derogate the performance of our jamming detection algorithms. Therefore, we achieve about the same performance as in the static scenario without background traffic, i.e. table 6.1.

However, the differentiation between background traffic generated by hidden terminals, which however should be classified as non-jamming, and frame jamming is more difficult. Figure 6.5 shows the detection performance for this scenario. The background traffic consists of a UDP stream of 2 Mbps, 10 Mbps and 20 Mbps respectively. Along with the transmitter-based and the receiver-based detection, also dedicated detection at monitor 2 was performed. Additionally, the cooperative detection among all three nodes was investigated as well.

We notice, that as long as the background traffic is at lower throughput rates than 10 Mbps, the performance differs not significantly from the case where no background traffic is present (experiment 1, shown at 0 Mbps). For higher amounts of background traffic, the probability of false classifications is increased strongly. To conduct these measurements, extensive testing was required in order to set up two separate networks, both transmitting at such high rates, while collisions occur at the receiver. Therefore, we believe that this particular scenario is of less practical relevance and represents a bound for the worst-case scenario.

The scenario of uncoordinated background traffic together with noise and bit jamming was not tested experimentally, but it is expected that the detection performance would not be affected by the background traffic according to our previous explanations.

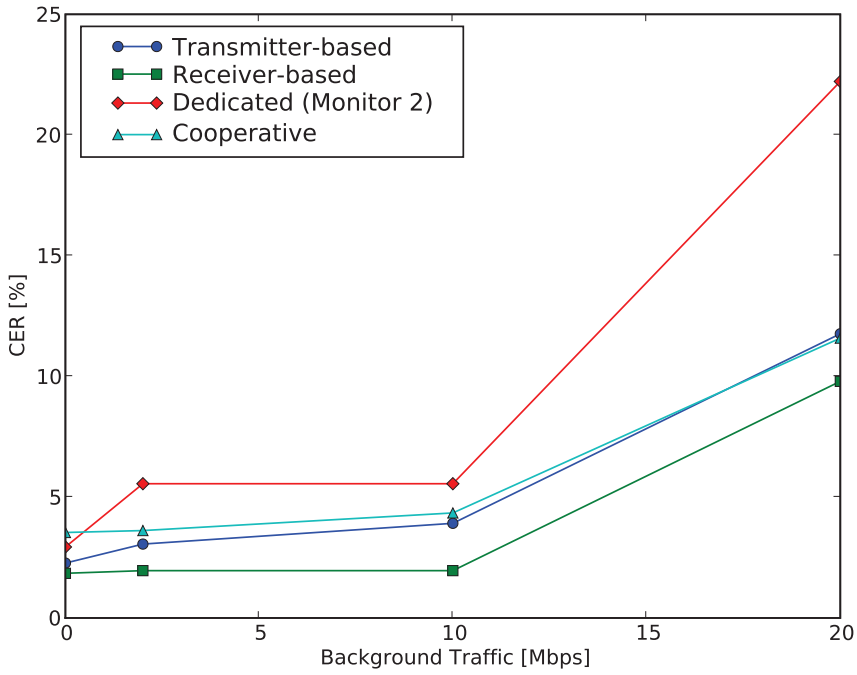

Figure 6.5: Experiment 8: Frame jamming and background traffic from hidden node at varying rates in the receiver area.

7 Discussion

In this chapter, the strengths and weaknesses of the developed jamming detection algorithms are discussed. Based on the gained experience during the experimental evaluation, we also make some recommendations for possible countermeasures against jamming attacks.

7.1 Strengths and Weaknesses of the Developed Algorithms

According to the evaluation in the previous chapter, the jamming detection algorithms proposed in this work allow the differentiation between jamming and natural effects, which may impact the performance of wireless communication systems. In particular, our algorithms exhibit the following strengths:

Common Metrics By using widely-used metrics which are readily available in different wireless communication systems, the algorithms may be easily adopted to other platforms.

Low Use of Resources Since no extensive signal processing is involved, the algorithms are appropriate for the usage in low-energy and low-computing power systems.

Detection of Various Jamming Types Thanks to the non-specific approach of the detection algorithms, the detection of various jamming mechanisms is possible. This might include even more advanced jamming techniques like smart jamming as well, since many of those techniques target the packet delivery ratio of a network

flow, as the proactive jammers do. However, this needs to be proved by further investigations.

Mobile Networks and Network Load Results show that the jamming detection is possible in mobile networks as well. In this case, we recommend the use of the receiver-based algorithm. In networks showing high regular traffic load, the discrimination of jamming is still possible.

Dedicated Detection The algorithms allow for the jamming detection at dedicated nodes. Our results show a high detection performance especially when there is only limited mobility in the network.

On the other hand, due to the nature of the developed algorithms, they show some weaknesses which should not be hidden:

Minimum Throughput A certain amount of received frames is needed in order to apply the algorithms. This means that a complete failure of the communication may not be detected by the proposed jamming detection algorithms. This comes from the fact, that for the computation of the PDR metric several received frames are needed. However, by using an additional signal quality metric as will be mentioned in chapter 9, or by using hardware detection techniques, this drawback could be minimized.

Unfairness not Detectable Attackers which generate high traffic load, but according to the used standards, might highly decrease the network performance. Furthermore, if they use a modified collision avoidance mechanism, they might access the channel in an unfair way, meaning that they could occupy the channel most of the time. Since in this case the amount of collisions is not increased by the jammer, this may not be detected by our algorithms. In order to detect such unfairness attacks, the nodes must have a deep knowledge of the amount of traffic which is expected on the channel under normal circumstances.

7.2 Recommendations for Jammed Environments

As soon as jamming attacks are detected, appropriate countermeasures should be taken by the participating regular network nodes. Since this is a separate problem which is not covered by this work, we just sketch some ideas:

Increased Transmission Power By increasing the transmission power, a higher signal-to-noise ratio may be achieved at the receiver. However, depending on the available energy resources on the regular nodes compared to the jammer, this might be not a sufficient countermeasure.

Change of Location If applicable, the affected nodes of a network could move away from the jammer. However, this might only be a temporary answer to the problem of jamming.

Higher PHY Rates As we showed in section 6.1.1, in jammed environments it might be worth trying to increase the physical rate. This leads to a higher packet loss, however, the resulting throughput might be increased in total.

We investigated how different physical rates, i.e. modulation schemes, perform in the presence of jamming. In fact, the lower rate schemes are optimized so that the communication is still possible in high path loss or non line of sight scenarios. However, it is not clear, whether the behaviour of selecting lower rates in a jammed scenario is appropriate as well.

A serie of four experiments was conducted in order to analyze the performance of all available PHY rates from 1 Mbps up to 54 Mbps. For the first nonjammed experiment, the distance between the transmitter and the receiver were increased, until the rate switching algorithm reached the lowest rate. Then, the PHY rate was manually increased while recording the achieved throughput. Figure 7.1 shows that the throughput was increased only marginal by using the 2 Mbps and 5.5 Mbps rates. At all other rates, no connection could be established. For the corresponding measurements in presence of jamming, the transmitter and receiver were positioned in a distance of 33 meters. The jammer was positioned between the two, at a distance of 18 meters from the transmitter and 15 meters from the receiver. In this setup and with any of the jammers present, the rate switching algorithm

selected the lowest rate, showing a throughput of approximately 800 kbps. If the PHY rate was increased manually, the best throughput was achieved at 11 Mbps for frame and noise jamming and at 5.5 Mbps for bit jamming. In case of frame and noise jamming, a major increase of 6.1 Mbps and 5.2 Mbps was achieved by selecting the 11 Mbps rate instead of 1 Mbps. Using bit jamming, the net increase was only 700 kbps between the 1 Mbps and 5.5 Mbps rates.

From these results we conclude, that in case of jamming, it might not always be the best strategy to select lower PHY rates, even if several frames are lost at higher rates. The physical differences between the DSSS and OFDM rates, which are also reflected by figure 7.1, show that it is reasonable to use the lower and more robust DSSS rates. However, the lowest DSSS rate is not always the best choice in jammed areas.

Adjust Busy Threshold Since certain jamming mechanisms are up to keep the channel busy, the regular nodes might want to adjust their "channel busy" threshold dynamically.

Ignore Spoofed Frames Spoofed frames from jammers might contain misleading "length" fields, preventing regular nodes from accessing the channel for long time intervals. Therefore, if frames are identified as malicious, they should be ignored by the collision avoidance mechanism. However, the main problem would be probably the differentiation between malicious and regular frames.

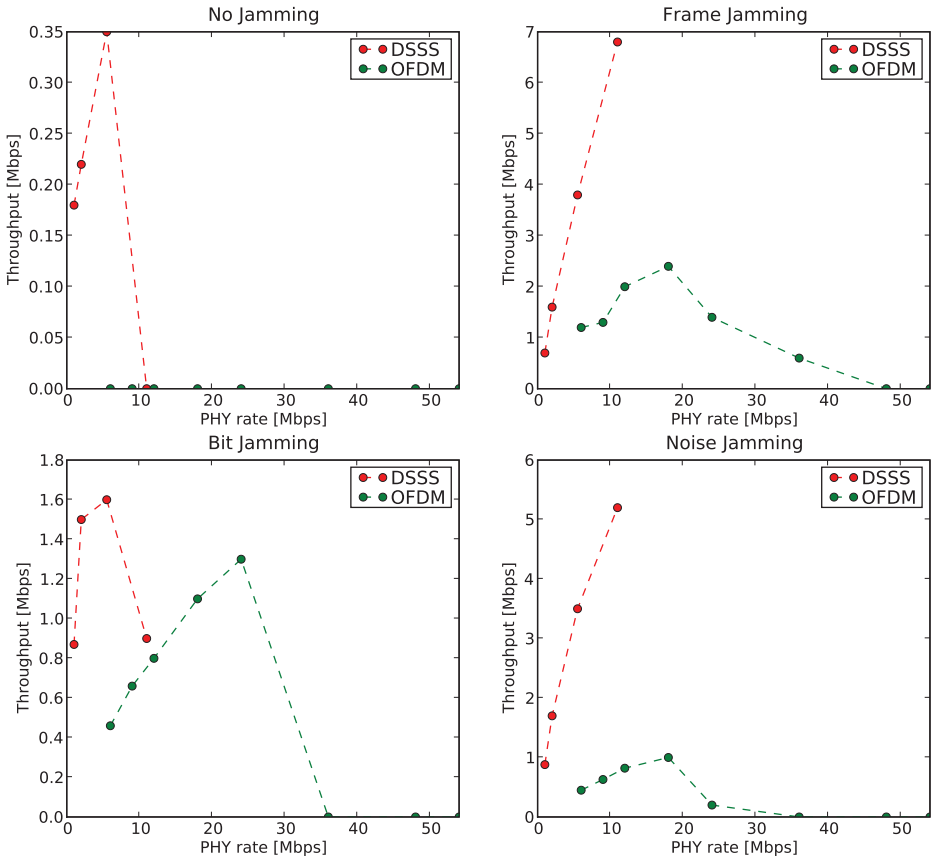

Figure 7.1: Performance of different PHY rates in case of jamming compared with a long-distance scenario.

8 Conclusion

To the best of our knowledge, jamming detection on the physical layer and under real-world conditions has not been widely studied in the literature so far. In this work, we contribute new insights from extensive experimental investigations based on an indoor WLAN ad hoc network. By the use of multiple modulation schemes and data rates on the physical layer and unicast data traffic on the transport layer, we explore several realistic network scenarios.

The developed jamming detection algorithms are based on the correlation between received signal strength, packet delivery ratio and the physical rate of an observed data flow. Additionally, the measured level of non-WLAN interference is considered as well. Since these are all well-known and widely-used metrics, the developed algorithms are applicable to a variety of wireless communication systems.

Depending on the basic conditions of every investigated scenario, the four algorithms achieved high detection performances. Three different jamming types were applied for the evaluation of the algorithms, namely noise, bit and frame jamming. However, we expect that other low-layer jamming mechanisms can be detected as well, as long as the decreased network performance is achieved by reducing the packet delivery ratio.

Jammer Position We found that in networks without mobile nodes, all participating nodes are able to reliably detect an active jammer. The transmitter-based, receiver-based, dedicated and cooperative detection approach achieved almost perfect performance. Therefore, under static conditions, we recommend to use the dedicated jamming detection algorithm, which allows to save resources on the remaining nodes.

Transport Layer Protocol By conducting measurements using UDP and TCP data streams, we investigated the influence of the traffic characteristics on the detection performance. Due to the nature of our algorithms, which work independently of the absolute number of sent frames, no significant differences were found.

Mobility Mobile nodes in the network make the detection of jamming more difficult. Fading effects lead to an increased natural packet loss, which must be distinguished from the packet loss due to the jammer influence. However, we found that in mobile scenarios, a high detection performance is achieved at the receiver. The performance of the transmitter-based and dedicated algorithms was significantly lower, but still remarkable. Depending on the actual application of the network, the dedicated detection is the right choice for mobile networks as well. This holds especially if only low energy and computing power are available at most nodes. In practice, by tuning the algorithm and at the cost of a higher false positive error rate, the number of false negatives may be further decreased.

Using the information from all nodes, which is the concept of the cooperative approach, did not show better performance than the other algorithms. This comes from the fact that in our algorithm, the information from all nodes is treated identically, which proved to be not adequate. Depending on their location, the information from several nodes might be less relevant and should be therefore weighted accordingly. Therefore, further investigations regarding an improved cooperative jamming detection algorithm are needed.

Background Traffic Background traffic generated in the same collision domain, meaning that it is detectable at the transmitter, does not influence the jamming detection performance. This is given by the collision avoidance mechanism of WLAN, which allows for low collisions even at a high network load, as long as no hidden nodes are transmitting. If the background traffic is generated at a hidden node, meaning that it may not be detected at the transmitter but at the receiver, there are collisions occuring, which lower the detection performance. We found a noticeable decrease as soon as the background traffic exceeds about 50% of the maximum channel throughput for the case of hidden nodes. However, this is a worst-case scenario which might not be encountered frequently in real applications.

9 Outlook

In the following, we present several suggestions for further investigations which might be carried out on the basis of our work.

Other Jamming Attacks It could be investigated, how our algorithms perform in the detection of other jamming attacks. For example, certain smart jamming mechanisms could be tested against our detection algorithms. Since smart jammers intend to save energy resources and try to hide themselves, it would be instructive to investigate how the detection performance is affected by such jamming, targeting only specific frames.

Other Platforms The developed algorithms could be implemented and tested on different wireless platforms, for example on sensor motes. Thanks to their low-level programming concept, the practical implementation is expected to be easier than when using WLAN cards. This would also allow for detailed analysis of the additional energy consumption caused by the execution of the jamming detection algorithms, which could be an important criterion in battery-driven environments.

Signal Quality Metric In WLAN systems, the bit error rate metric could be used instead of the PDR, when frames are sent at the lower DSSS PHY rates (cf. section 5.1.2). Since this is a measure for the number of bits which are altered during a frame transmission, this allows for per-frame execution of our algorithms. However, it is not clear if the correlation between the bit error rate and the signal strength is similar to the one we used, i.e. PDR versus signal strength. The main advantage of this approach would be the ability to apply the algorithms even at very low frame reception rates, since a single frame would be enough to make a jamming / no jamming decision.

Advanced Cooperative Detection Using additional information from positioning systems (e.g. GPS), the cooperative detection might be improved and maps of jammed an non-jammed areas could be drawn. This means that the information coming from different nodes would be weighted according to its local relevance. As a consequence, the rate of false decisions is expected to drop.

Bibliography

[1] Ieee standard for information technology - telecommunications and information exchange between systems - local and metropolitan area networks - specific requirements. part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications, June 2007.

[2] Ieee standard for information technology - telecommunications and information exchange between systems - local and metropolitan area networks - specific requirements. part 3: Carrier sense multiple access with collision detection (csma/cd) access method and physical layer specifications, December 2005.

[3] D. L. Adamy. *EW 101: A First Course in Electronic Warfare*. Artech House, 2001.

[4] D. L. Adamy. *EW 102: A Second Course in Electronic Warfare*. Artech House, 2004.

[5] R. A. Poisel. *Modern Communications Jamming Principles and Techniques*. Artech House, 2004.

[6] A. Wood, J. Stankovic, and S. Son. Jam: A jammed-area mapping service for sensor networks. *24th IEEE Real-Time Systems Symposium*, pages 286–297, 2003.

[7] W. Xu, W. Trappe, Y. Zhang, and Wood T. The feasibility of launching and detecting jamming attacks in wireless networks. *Proceedings of ACM MobiHoc*, pages 46–57, May 2005.

[8] P. Kyasanur and N. Vaidya. Detection and handling of mac layer misbehavior

in wireless networks. *Proceedings of the 2003 IEEE Internationsal Conference on Dependable Systems and Networks*, pages 173–182, 2003.

[9] J. Bellardo and S. Savage. 802.11 denial-of-service attacks: Real vulnerabilities and practical solutions. *Proceedings of the USENIX Security Symposium*, pages 15–28, 2003.

[10] G. Noubir and G. Lin. Low-power dos attacks in data wireless lans and countermeasures. *SIGMOBILE Mob. Comput. Commun. Rev.*, 7(3):29–30, 2003.

[11] A. Wood and A. Stankovic. Denial of service in sensor networks. *IEEE Computer*, 35(10):54–62, October 2002.

[12] N. Borisov, I. Goldberg, and D. Wagner. Intercepting mobile communications: the insecurity of 802.11. In *MobiCom '01: Proceedings of the 7th annual international conference on Mobile computing and networking*, pages 180–189, New York, NY, USA, 2001. ACM.

[13] Ieee standard for information technology - telecommunications and information exchange between systems - local and metropolitan area networks - specific requirements. part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications. amendment 6: Medium access control (mac) security enhancements, July 2004.

[14] B. Schneier. Attack trees. *Dr Dobb's Journal*, 24(12), December 1999.

[15] J. Bardwell. *Converting Signal Strength Percentage to dBm Values*. WildPackets, November 2002.

[16] ZyXEL. *ZyXEL ZyAir G-110. Datasheet*, 2004.

[17] Netgear. *Netgear WG511T. Datasheet*, 2007.

[18] Rohde & Schwarz. *Rohde & Schwarz R&S SMU200A Vector Signal Generator. Operating Manual.*, November 2008.

[19] R. Duda, P. Hart, and D. Stork. *Pattern Classification*. John Wiley & Sons, 2 edition, 2001.

Appendices

A Implementation Details

This section contains additional details regarding hardware and software used in the experimental implementation.

A.1 Hardware and Driver Details

WLAN Cards Those WLAN cards used for the network communication were run with their normal drivers, which are enclosed in the respective folder on the CD-ROM. The cards used for the information retrieval were run with the Atheros Driver version 4.2.2.8 by WildPackets, which is part of the Omnipeek suite and enclosed on the CD-ROM as well.

Vector Signal Generator SMU 200A For the noise jamming, white noise with a bandwith of 20 MHz around the center frequency was used. The bit and frame jamming was generated using the following preferences:

- Standard: 802.11b
- Physical Layer Mode: CCK
- Simulation Mode: Framed (Frame Jamming) / Unframed (Bit Jamming)
- Predefined Frames: Data
- Sequence Length: 1 Frame
- Idle Time: 0.100 ms
- PPDU Configuration:
  - PLCP P+H Format: Long PLCP

  - PSDU Bit Rate: 11 Mbps
  - Data Length: 1024 bytes
  - PSDU Data Source: PRBS 9
  - Scrambler: On
  - Service Field Clock Bit: Locked

A.2 Software Details

All the used software tools are enclosed on the CD-ROM. These are:

- JPerf 2.0.0 (includes Iperf)
- Omnipeek for Windows 5.0
- ActivePerl 5.1
- Python 2.5
- Matplotlib 0.91 (library, used for the generation of graphics)
- Numpy 1.2.1 (library, used for calculations inside the python script)

A.3 Data Acquisition Details

In order to export the required frame details from Omnipeek, and for the compatibility with our perl script, the following fields must be visible in the Omnipeek packet overview in the specified order:

- Packet
- Source
- Destination
- BSSID
- Flags

- Signal dBm
- Noise dBm
- Data Rate
- Size
- Relative Time
- Protocol
- Summary
- Expert

The selected frames should be exported using the command `File -> Save Selected Packets...` and selecting the "Packet List (.csv)" file type.

Using the `getRawData.pl` script, the data is transformed in our specific format. Under the Windows operating system, the required commands are:

```
C:\Users\UserXY>perl getRawData.pl fileName.csv
```

This generates the file `fileName.dat` which contains only the needed frame data.

A.4 Data Processing Details

The data processing is done by the `dataProcessing.py` script. The script includes extensive comments, therefore we mention only its structure.

At the beginning, several user-defined parameters may be specified. This includes the name of the experiment, its duration etc. If the script should be tested with the enclosed example files, "experimentID" should be defined as "coopA1".

The further parameters which may be specified concern the $frameQuantity$, $noiseThreshold$, $maxPDR$ and $rssiDiffThreshold$ parameters respectively.

If the script is executed, the required values are read from the specified data files. The PDR and all average values are calculated. Then, these data tuples are

evaluated and classified as jammed or not jammed. Finally, the plots containing the sequences of all parameters as well as the RSSI vs. PDR and PHY rate - plot are generated. Additionally, some statistical data is shown like the number of total decisions, the number of jamming and non-jamming decisions as well as the confidence interval for the specific measurement. If a fixed $rssiDiffThreshold$ was specified, then also the CER is calculated and shown.

B PHY Rate Switching Algorithms

In the following we present some rate switching algorithms in the pursuance of the MadWiFi source code available at http://www.madwifi.org.

Minstrel Rate Control Algorithm This code is takes a wandering minstrel approach. Wander around the different rates, singing wherever you can. And then, look at the performance, and make a choice. Note that the wandering minstrel will always wander in directions where he/she feels he/she will get paid the best for his/her work.

Onoe Rate Control Algorithm Onoe is a credit based RCA where the value of the credit is determined by the frequency of successful, erroneous and retransmissions accumulated during a fixed invocation period of 1000 ms. If less than 10% of the packets need to be retransmitted at a particular rate, Onoe keeps increasing its credit point till the threshold value of 10 is reached. At this point, the current transmission rate is increased to the next available higher rate and the process repeated with credit score of zero. Similar logic holds for deducting the credit score and moving to a lower bit-rate for failed packet transmission/retransmission attempts. However, once a bit-rate has been marked as failure in the previous attempt, Onoe will not attempt to select that bit-rate until 10 seconds have elapsed since the last attempt. Due to the manner in which it operates, Onoe is conservative in rate selection and is less sensitive to individual packet failure.

AMRR Rate Control Algorithm AMRR uses Binary Exponential Backoff (BEB) technique to adapt the length (threshold) of the sampling period used to change the values of bit-rate and transmission count parameters. It uses probe packets and depending on their transmission status adaptively changes the threshold value. The adaptation mechanism ensures fewer failed transmission/retransmission and

higher throughput by not switching to a higher rate as specified by the backoff mechanism. In addition to this, the AMRR employs heuristics to capture the short-term variations of the channel by judiciously setting the rate and transmission count parameters.

SampleRate Rate Control Algorithm SampleRate decides on the transmission bit-rate based on the past history of performance; it keeps a record of the number of successive failures, the number of successful transmits and the total transmission time along with the destination for that bit-rate. Stale samples are removed based on a EWMA windowing mechanism. If in the sampling process, no successful acknowledgment is received or the number of packets sent is multiple of 10 on a specific link, it transmits the packet with the highest rate which has not failed 4 successive times. Other than that it transmits packets at the rate which has the lowest average transmission time.

C Assignment

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Eidgenössisches Departement für Verteidigung,
Bevölkerungsschutz und Sport VBS

armasuisse
Wissenschaft und Technologie W+T

Master of Science Thesis Assignment
for
Markus Schafroth (D-ITET, ETH Zürich)

Main advisor:	Dr. Vincent Lenders, armasuisse, S+T
Advisor ETH:	Dr. Franck Legendre, ETH Zürich, TIK
Supervisor:	Prof. Dr. Bernhard Plattner, ETH Zürich, TIK

Start date:	September 8th 2008
End date:	March 13th 2008

Detecting Denial of Service Attacks in Wireless Ad Hoc and Sensor Networks

Introduction

Wireless ad hoc and sensor networks are increasingly being used in military operations. These networks offer cheap and rapidly deployable connectivity for communication and sensing purposes. However, these networks are often deployed in hostile environments and the broadcast nature of the wireless medium exposes them to targeted attacks. In particular, those networks are vulnerable to denial of service (DoS) attacks at the link level by attackers emitting interfering radio signals with the purpose of disturbing undergoing transmissions (so called jamming).

Scope

We consider a CSMA wireless mobile ad hoc network (e.g., IEEE 802.11, Zigbee, etc.) in which the participating nodes are protected by means of strong encryption at the link layer. An external attacker can perform a DoS attack by emitting signals which interfere with the wireless signals of the network, leading to a degradation of the network performance. The goal of this thesis is to develop mechanisms to detect such DoS attacks while using performance metrics available at the MAC layer such as the received signal strength indicator (RSSI), the packet delivery ratio, the bit error rate, etc. One of the main challenges is that these measurable metrics are affected by factors like the environment, the mobility, and the traffic load. The developed DoS detection mechanisms should hence be robust and work properly in various settings with different environmental, mobility and traffic conditions. Furthermore, the detection mechanisms must work on node with limited resources like small battery powered sensor nodes (e.g. Iris motes).

Tasks

The tasks of the thesis are:

1. Review of the literature on DoS attacks in wireless networks (e.g., [1,2]).
2. Development and evaluation of different link metrics that can be used to monitor the status of wireless links (e.g., RSSI, packet delivery ratio, bit error rate).
3. Development of non-cooperative and cooperative DoS detection algorithms based on the evaluated link metrics.
4. Implementation of the developed DoS detection algorithms on laptops and/or Iris motes.

5. Assessment of the accuracy of the developed mechanisms by measuring the performance of those under different attacker models and different environmental conditions.
6. Refinement of the developed mechanisms based on the outcome of the measurements and conclusion on the feasibility of detecting individual attacker models.

Deliverables

- At the end of the second week, a detailed time schedule of the thesis must be given and discussed with the main advisor.
- At half time of the semester thesis, a short discussion of 15 minutes with the supervisor and the advisors will take place. The student has to talk about the major aspects of the ongoing work. At this point, the student should already have a preliminary version of the written report, including a table of contents. This preliminary version should be brought along to the short discussion.
- At the end of the thesis, a presentation of 15 minutes must be given at armasuisse and at ETH Zürich. The presentations should give an overview as well as the most important details of the work. If possible, a demonstrator should be presented at this time.
- The final report should be written in English but may be written in German. It must contain a summary written in both English and German, the assignment and the time schedule. Its structure should include an introduction, an analysis of related work, and a complete documentation of all used software tools. Four written copies of the final report must be delivered to the main advisor.

References

[1] Anthony D. Wood and John A Stankovic, "Denial of Service in Sensor Networks", IEEE Computer, 35(10):54-62, October 2002.
[2] WenYuan Xu, Wade Trappe, Yanyong Zhang and Timothy Wood, "The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks", ACM Mobihoc, Urbana-Champaign, Illinois, USA, May 2005.

armasuisse
Science and Technology
Command and Control Lab

Dr. Vincent Lenders

Berne, September 8th 2008

D Time Schedule

Master of Science Thesis

Markus Schafroth (D-ITET, ETH Zürich)

Jamming Detection in Wireless Ad Hoc Networks

Time Schedule

Start date: September 8th 2008
End date: March 13th 2008

Weekly schedule

Week	Tasks
37 (2008)	Literature review
38	Literature review
39	Development of link metrics
40	Development of link metrics Tools evaluation
41	Tools evaluation
42	Development of DoS attack mechanisms
43	Development of non-cooperative jamming detection algorithms (1)
44	Development of algorithms (1)
45	Implementation of algorithms (1)
46	Implementation of algorithms (1)
47	Testbed installation and testing
48	Testbed installation and testing
49	Measuring and analysis of algorithms (1)
50	Measuring and analysis of algorithms (1)
51	Refinement of algorithms (1) Mid-term presentation 1 at ETH Zürich
52	*** Christmas *** (Buffer time)
1 (2009)	Development of dedicated and cooperative jamming detection algorithms (2)
2	Implementation of algorithms (2)

3	Measuring and analysis of algorithms (2)
4	Measuring and analysis of algorithms (2)
5	Final experiments
6	Final experiments Mid-term presentation 2 at ETH Zürich
7	Final experiments
8	Buffer time Writing final report
9	Writing final report
10	Writing final report Preparation of presentations
11	Finishing the final report Presentations at armasuisse and ETH Zürich

Milestones

Week 48 (2008)	First results gained from testbed
Week 2 (2009)	Completed the implementation of the detection algorithms
Week 7	Final experimental data available
Week 10	Draft of final report completed
Week 11	Final report delivered

E CD-ROM Contents

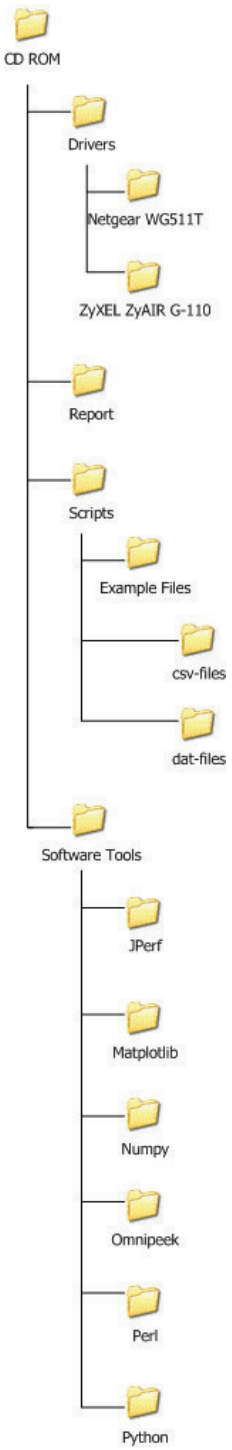

Figure E.1: CD-ROM contents