RZ 3094 (# 93140) 12/28/98
Engineering & Technology 70 pages

Research Report

Experimental Performance Study of the Advanced Infrared MAC Protocol

Dewis Bernasconi

IBM Research Division
Zurich Research Laboratory
8803 Rüschlikon
Switzerland

IBM

Research Division
Almaden • Austin • Beijing • Haifa • T.J. Watson • Tokyo • Zurich

Experimental Performance Study of the Advanced Infrared MAC Protocol

Dewis Bernasconi

IBM Research Division, Zurich Research Laboratory, 8803 Rüschlikon, Switzerland

Abstract

Most mobile computers today are equipped with wireless infrared ports based on the Infrared Data Association (IrDA) standard. The IrDA specifications undergo several restrictions: the communication is aimed, short-range and point-to-point; the speeds range from 2.4 kbps to 4 Mbps. The recent advances in the speed and in the extension of range and angle of infrared communications, together with the requirement of more sophisticated mobile networking capabilities, have emphasized the need to overcome the IrDA restrictions enhancing the capabilities of the existing standard. This need gave birth to Advanced Infrared (AIr), an IBM project that seeks to define the next generation infrared data communications system. The IrDA community is currently in the process of adopting the proposed AIr communication system as a new standard.

The AIr medium access control (AIr MAC) protocol coordinates the access to the infrared medium between various AIr devices. Its main features are a connectionless data transfer mechanism, a Request-To-Send/Clear-To-Send (RTS/CTS) collision avoidance protocol for multipoint-to-multipoint connectivity and the dynamic data transfer rates.

This report proposes an experimental study of the AIr MAC protocol performance. Four prototyped AIr stations have been operated in different scenarios to investigate how some protocol options (reservation termination, data frames acknowledgement, backoff algorithm and slot synchronization mechanism, RTS body format) have an effect on some selected performance evaluation parameters (throughput, data frame loss/drop and packet success rate, channel access delay), particularly under critical coverage conditions and imperfect transceiver parity. It turns out that channel reciprocity (transceiver parity), backoff window adjustment policy and RTS body rate are key factors for the protocol performance.

Table of Contents

Figures

1. Introduction

This introductory chapter defines the problem statement, the objectives and the scope of this diploma work thesis. The organization of this report is also presented.

1.1 Problem Statement

Most mobile computers today are equipped with wireless infrared ports based on the IrDA (Infrared Data Association) standard, which enables aimed short-range point-to-point communication at speeds from 2.4 Kbps up to 4 Mbps. Mobile users, however, require more sophisticated networking capabilities, such as support of Local Area Network (LAN) access and peer-to-peer networking with longer range and unaimed multipoint transmission links. This poses new challenges with respect to devising methods for reliable transmission under adverse channel conditions (e.g., fluctuating background light and variable transmission distances) and for avoiding collisions in the case of multiple links operating within the same room. The Infrared Communication Systems Group at IBM Zurich Research Laboratory is enhancing the capabilities of the existing IrDA standard to achieve these objectives. The pursued work addresses the design and implementation of a new Advanced Infrared (AIr) physical and medium access control layers to establish a foundation for enhanced upper protocol layers.

1.1.1 Starting Point

AIr is an IBM project for enabling multipoint-to-multipoint wireless communication over short distances in an office environment. It represents a significant extension over the existing point-to-point IrDA standard in the sense of providing more robust, longer distance and wide-angle transmission paths. AIr is based on a new physical layer involving four-slot Pulse Position Modulation (4-PPM) with Variable Rates (VR, 4 Mbps – 0.25 Mbps), the new concept of transceiver parity to establish channel reciprocity, and a "Request-To-Send"/"Clear-To-Send" (RTS/CTS) medium access protocol (AIr MAC). In order for the RTS/CTS protocol to work properly it is required that media access information is reliably communicated to all stations capable of interfering (mitigation of hidden terminal problem). This is achieved by repetition-encoding media access information in each frame header with maximum Rate Reduction RR=16 to provide maximum reach, whereas data are repetition-encoded according to the required transmission distance. Transceiver parity ensures channel reciprocity by demanding that each AIr transceiver has spatially matched optical emission/reception characteristic (LEDs and photodiodes). The question arises how the network throughput performance is affected by various AIr MAC implementation options and by manufacturing tolerances (parity) especially under critical coverage conditions, i.e., at the cell boundaries for each data rate.

1.1.2 Goals of This Work

The following goals have been defined, by my industrial supervisor Dr. Gfeller, at the beginning of the industrial thesis work:

1. Install AIr prototype hardware and software components and bring four AIr stations to operating status. Acquire a working knowledge of the system.

2. Design and implement additional hardware/software tool(s) to trace logic events of the AIr MAC finite state machine.

3. Exercise the different options provided by the AIr MAC Super Finite State Machine and analyze performance results.

4. Investigate the effects caused by imperfect transceiver parity and different user scenarios under critical coverage conditions.

5. Compare simulated throughput results (provided by Dr. Schott) with actual measurements.

6. Provide recommendations for optimum usage of the different AIr MAC protocol options.

7. Final report describing the work and summarizing the results.

1.2 Organization of This Report

The rest of this report is split into five chapters and four appendices:

Chapter 2 gives a brief overview of the Infrared Data Association and of its current standards (IrDA 1.0 and IrDA 1.1). This overview is needed to understand the ideas that are behind the Advanced Infrared system, which is an evolution of the existing IrDA standards and tries to remove their limitations.

Chapter 3 presents the Advanced Infrared system, its architecture and its main features. Particularly, descriptions of the physical layer (AIr IrPHY), of the medium access control layer (AIr MAC) and of the tools that have been used to evaluate the performance of AIr MAC are provided. The information collected in this chapter should be detailed enough to allow a full understanding of the performance measurements presented in chapters 4 and 5.

Chapter 4 presents the test scenarios and the results obtained with the *preliminary performance evaluation measurements* (experimental test models S1 to S5). The goal of the preliminary measurements was to investigate how some protocol options (reservation termination path, data acknowledgement, backoff algorithm, size of data bursts) can influence the AIr MAC performance, evaluated in terms of system throughput and frame drop/loss at the involved AIr stations. A theoretical simulation model has also been used to gain a deeper insight in the obtained experimental results. Some recommendations for the protocol options usage have been finally provided.

Chapter 5 presents the test scenarios and the results obtained with the *interference measurements* (experimental test models S6 to S8). The goal of the interference measurements was to investigate how a transceiver parity mismatch at potential interferers can cause a degradation of the AIr protocol performance (in terms of packet success ratio), and have been inspired by the study presented in [8]. During the measurements, another problem (the effect of RTS capture during RTS-RTS collisions) has been individuated and investigated. Some recommendations that should be followed to mitigate the bad effects of transceiver disparity and RTS capture have been finally provided.

Chapter 6 summarizes the results of the preliminary and of the interference measurements and draws some conclusions.

Appendix A provides a tutorial which explains how to set up and configure the prototype AIr stations that have been used for the experimental performance measurements.

2. The Infrared Data Association (IrDA)

A brief overview of the current Infrared Data Association (IrDA) standards is presented in this chapter. It should allow a better understanding of the AIr system specification, which is presented in chapter 3.

2.1 IrDA Overview

IrDA is an international organization that creates and promotes interoperable, low cost infrared data interconnection standards that support a walk-up, point-to-point user model. It has been founded in 1993 as a non-profit organization and has over 160 members [1]. In 1997, there were over 20 million laptops, PDAs, phones, camera and printers fitted with IrDA ports.
The key features of IrDA standard are:

- Simple and low cost implementation
- Low power requirement
- Directed, point-to-point connectivity
- Efficient and reliable data transfer

IrDA has specified the physical and protocol layers necessary for any two devices that conform to the IrDA standards to detect each other and exchange data [2]. The initial IrDA 1.0 specification detailed a serial, half-duplex, asynchronous system with transfer rates of 2.4 kbps to 115.2 kbps at a range of up to one meter with a viewing angle of between 15 and 30 degrees. It could be implemented using standard universal asynchronous receiver-transmitters (UARTs), a transceiver and a fairly simple communications stack. More recently, IrDA has extended the physical layer specification to allow data communications at transfer rates up to 4 Mbps (IrDA 1.1). As these higher speeds could not be implemented on UARTs, it took longer to develop and incorporate technical solutions.
The IrDA standard architecture has been summarized in the next section.

2.2 IrDA Architecture

The layers within the IrDA protocol stack can be divided into two groups: required and optional protocols [2]. The required layers of the stack are:

- The physical layer (IrPHY), which specifies optical characteristics, encoding scheme and framing definition for various speeds
- The link access protocol (IrLAP), which establishes the basic connection reliably
- The link management protocol (IrLMP), which multiplexes services and applications on the IrLAP connection (LM-MUX sublayer) and provides a directory or list of services on a device (LM-IAS sublayer).

The supported optional protocols include:

- TinyTP: tiny transport protocol which adds per-channel flow control and is required by most applications
- IrOBEX: object-exchange protocol which enables easy transfer of files and other objects
- IrCOMM: the protocol which offers serial and parallel port emulation, enabling existing applications to use IR
- IrLAN: local area network access which enables walk-up access to notebooks and other mobile devices.

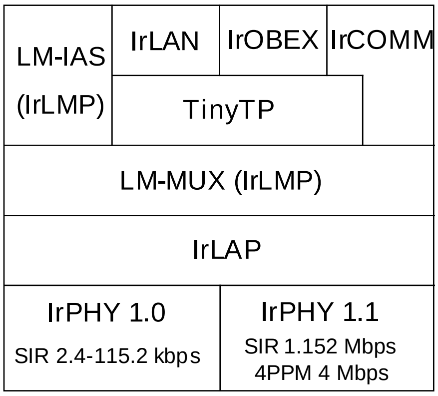

Figure 2.1: IrDA protocol stack

The IrPHY, IrLAP and IrLMP standards form the core of the IrDA architecture, and all are required for a device to be IrDA-compliant.

2.2.1 IrPHY (IrDA Physical Layer)

The IrDA physical layer is split into three distinct data ranges: 2.4 to 115.2 kbps, 1.152 Mbps and 4 Mbps. The links are designed to be used in a line-of-sight, point-and-shoot manner and hence have a modest minimum coverage of one meter, with a ±15 viewing angle. The 2.4 to 115.2 kbps mode is a half-duplex system with the maximum UART based data rate of 115.2 kbps. The 1.152 Mbps and 4 Mbps modes require a few important upgrades and modifications to the transceiver circuitry and driver software needed for these higher speeds. The 4 Mbps mode uses Pulse Position Modulation (PPM) data encoding. High speed modes are transparent to IrLAP and IrLMP with speed options negotiated during normal IrLMP discovery processes.

2.2.2 IrLAP (IrDA Infrared Link Access Protocol)

IrLAP is a link protocol for the serial infrared link and is derived from an existing asynchronous data communications standard (an adaptation of HDLC). It provides the basic link layer connection between a pair of IrDA devices, like connection establishment, data transfer and flow control. It also has additional features as a result of the specific properties of the Infrared medium.

2.2.3 IrLMP (IrDA Link Management Protocol)

IrLMP is layered on top of IrLAP and has two main functions: application and service discovery and multiplexing of application-level connections over single IrLAP connections. The IrLMP layer allows individual service users (applications) to connect and exchange information with similar entities in the peer device, independent of any other service users that may be using the IrLAP connection. The IrLMP layer provides multiple independent channels to the IrLMP layer in the remote device. It also provides a service with which applications can locally register themselves and some significant parameters in an information base. The two main functions provided by IrLMP are split between two sublayers: the Link Management Multiplexer (LM-MUX) provides the facilities for multiplexing application level connections over an IrLAP connection between a pair of devices; the Link Management Information Access Service (LM-IAS) provides the services necessary to allow applications to discover devices and access the information in the information base of a remote device.

3. The Advanced Infrared (AIr) System

This chapter provides a general overview of the Advanced Infrared (AIr) system, its architecture and its features. All the concepts and definitions that are needed to understand the performance evaluation measurements presented in chapters 4 and 5 have been detailed. It should be noted that the AIr specification is in continuous evolution and it is currently in a draft status. Some new issues, that have been introduced in the AIr specification during the evolution of this diploma thesis work and that were not implemented yet in the prototype devices used for the performance study measurements, have been left out of this chapter.

3.1 AIr Overview

The IrDA 1.x specifications undergo several restrictions: the communication is aimed, short-range and point-to-point; the speeds range from 2.4 kbps to 4 Mbps. The recent advances in the speed and in the extension of range and angle of infrared communication, together with the requirement of more sophisticated mobile networking capabilities (such as support of LAN access and peer-to-peer networking with longer range and unaimed multipoint transmission links), have emphasized the need to overcome the IrDA restrictions enhancing the capabilities of the existing standard. This need gave birth to Advanced Infrared (AIr), a research collaboration between IBM Research (IBM ZRL, IBM T.J. Watson Research Center), IBM Microelectronics Division in Toronto and Hewlett-Packard Laboratories that seeks to define the next generation infrared data communications system. The IrDA community is currently in the process of adopting the proposed AIr communication system as a new standard.

3.1.1 AIr Features

AIr is designed to provide the following features:

- Interoperability with IrDA 1.x, while removing all its constraints on range and angular coverage. Communication is no longer constrained to a 15-30 degree half-angle cone that extends 1 to 2 meters from the devices: AIr allows multiple devices to communicate with each other over ranges up to 5-10 meters at data rates between 4 Mbps and 125 kbps. The higher rates will function over shorter ranges while the lower data rates will achieve near diffuse coverage in confined spaces.

- Peer-to-peer connectivity between multiple devices operating in the same space, using a medium access control protocol that allows multiple devices to interact with each other. It is clearly not reasonable that communication between a single pair of AIr devices should prevent other devices within a 5-10 meters range from establishing communication with each other. If two devices are within each other field of view then, subject to the medium access control protocol rules, they may establish communication. This results in an enhanced freedom of movement.

- Hidden terminal interference management, through the use of a burst reservation technique to gain access to the medium. Robust coding of reservation messages is employed to ensure that they reach all devices capable of perturbing the received data. Fig. 3.1 shows the classical hidden terminal problem.

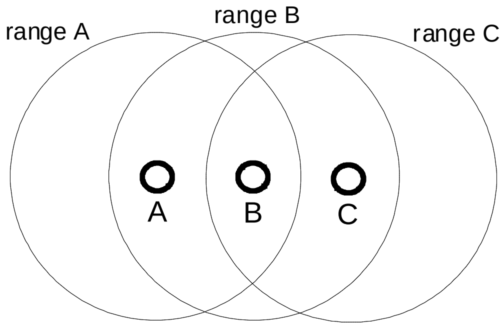

Figure 3.1: The hidden terminal problem

The three wireless stations A, B and C are depicted with their communication ranges (the transmission range overlaps with the reception range). Let's suppose that a CSMA scheme is adopted: every station listens for other transmissions and only transmits if no one else is doing so. C is within range of B but not within range of A. When A is transmitting to B and C senses the medium, it will not hear A and thus falsely conclude that it can transmit. If C start transmitting, it will interfere at B perturbing the frames from A. The problem of a station that is not able to detect a potential competitor for the medium because it is out of range is called the hidden terminal problem. It increases with the extent of the coverage of a wireless technology and with imperfect transceiver parity.

The AIr MAC protocol enables two communicating nodes to reserve the medium and to notify all possible interfering nodes of this reservation. This is done by the transmission of Request-To-Send (RTS) and Clear-To-Send (CTS) reservation messages. The transmission of a RTS by the sender (station A in Fig. 3.1) enables the nodes hidden to the intended recipient to be aware of the communication and not to interfere. The transmission of a CTS by the receiver (station B), causes the nodes hidden from the sender (station C which is hidden from station A) to avoid an interfering transmission. In the case of ideal channel reciprocity between all devices, only the CTS frame would be needed to prevent any interference with the receiver.

- Retention of IrDA upper layers (IrLMP, TinyTP, IrCOMM, IrLAN), so that the transition to AIr is transparent to the existing applications.

3.1.2 AIr Architecture

Fig. 3.2 shows the AIr system architecture.

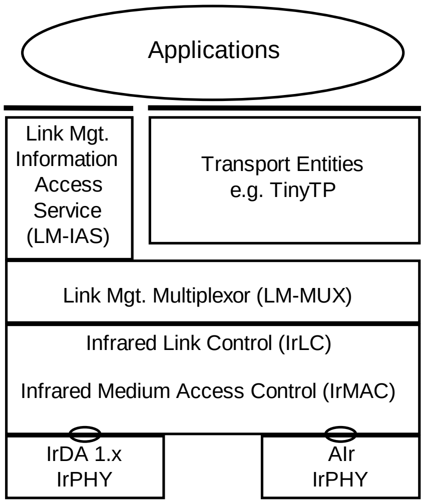

Figure 3.2: AIr system architecture [5]

It consists of the infrared physical layer (AIr IrPHY) and of the link access protocol layer (AIr IrLAP) at his top. AIr IrLAP is split in two sublayers: the link control layer (AIr IrLC) and the medium access control layer (AIr MAC).
AIr IrLC supports connections to multiple devices and services for discovery and data rate negotiation.
AIr MAC (IrMAC) coordinates access to the infrared medium between various AIr devices, allowing a device to cope with the relaxing of restrictions on the angle and the range of the IrDA physical layer. Its main features are a connectionless data transfer mechanism and a RTS/CTS collision avoidance protocol for burst reservation.
AIr IrPHY is a new physical layer featuring variable data rates and transceiver parity for approximate infrared channel reciprocity. It makes use of repetition coding to robustly code reservation information and to allow the dynamic adjustment of communication range and data transfer rate. Both short range high speed communication and long range low speed communication can be achieved.

A more detailed overview of AIr IrPHY and AIr MAC (IrMAC) is provided in the next two sections.

3.2 AIr Physical Layer (AIr IrPHY)

The AIr physical layer (AIr IrPHY) is specified in the draft document [5]. It has been designed with the aim of ensuring the following two properties:

1. The capability to transmit the information contained in the RTS and CTS frames at a distance sufficient to cover the potential interfering range. It is done by repetition encoding frame headers carrying AIr MAC information.
2. An approximate Infrared Channel Reciprocity (ICR) between all AIr stations

The hidden terminal problem can be avoided, or at least significantly mitigated, if these two properties are ensured.

3.2.1 Variable Repetition Encoding (4-PPM/VR)

AIr IrPHY adopts a 4-PPM/VR modulation scheme. Four-slot Pulse Position Modulation (4-PPM) is achieved by defining a data symbol of duration Dt and subsequently subdividing the symbol into a set of 4 equal time slots called “chips”. In the 4-PPM scheme, only one time slot, or chip, per symbol contains a pulse (logical one); the other chips contain no pulse (logical zero). Each chip has a duration of Ct given by Ct = Dt / 4. For AIr IrPHY Dt = 500 ns and Ct = 125 ns (see Fig. 3.3).

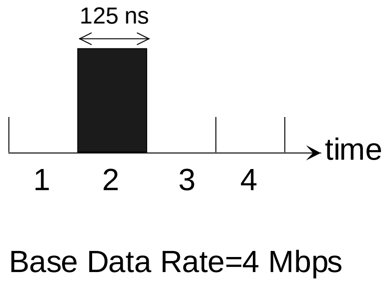

Figure 3.3: Four-slot Pulse Position Modulation

Because there are four unique chip positions within each 4-PPM symbol, four independent symbols exist. We define these four symbols to be the only legal data symbols (DD) allowed in 4-PPM. Each DD represents two bits of payload data, or a single “data bit pair (DBP)”, so that a byte of payload data can be represented by four DDs in sequence. The following table defines the chip pattern representation of the four unique DDs defined for 4-PPM.

Data Bit Pair (DBP)	4-PPM Data Symbol (DD)
00	1000
01	0100
10	0010
11	0001

Logical “1” represents a chip duration when the transmitter is emitting light, while logical “0” represents a chip duration with no light emission. The above definition of the 4-PPM symbol results in the base data rate of 4 Mbps.

In AIr IrPHY, the standard 4-PPM modulation scheme is extended to support repetition encoding. In a string of 4-PPM data symbols DD, every individual symbol is repeated “RR” times, where RR is the rate reduction factor (see Fig. 3.4). For example, a sequence of 4-PPM symbols “A B C” yields the redundant sequence “AAAA BBBB CCCC” with RR=4. The rate reduction factor can be 1,2,4,8 or 16, where RR=1 corresponds to the base data rate of 4 Mbps.

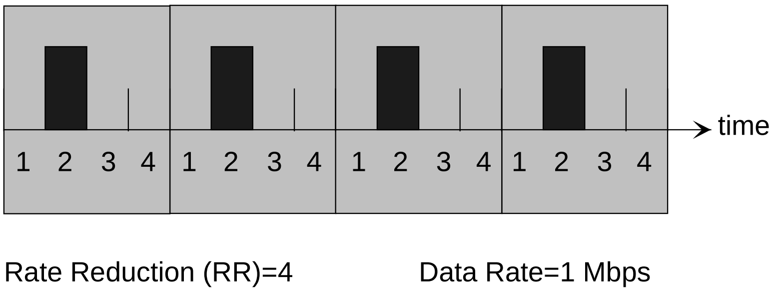

Figure 3.4: Repetition encoding with RR=4

The introduced redundancy in the symbol stream can be exploited with suitable digital signal processing methods in the demodulator circuit and provide a SNR gain of nominally 3 dB (electrical) for each doubling of the repetition rate, at the expense of a reduced data rate of 4/RR Mbps.

Infrared data transmissions that can be reliably received within a certain range will cause interference with a receiver at a much greater range. In other words, a transmitter may subject a receiver to interference even though that receiver cannot reliably decode those transmissions. This situation is highlighted in Fig. 3.5.

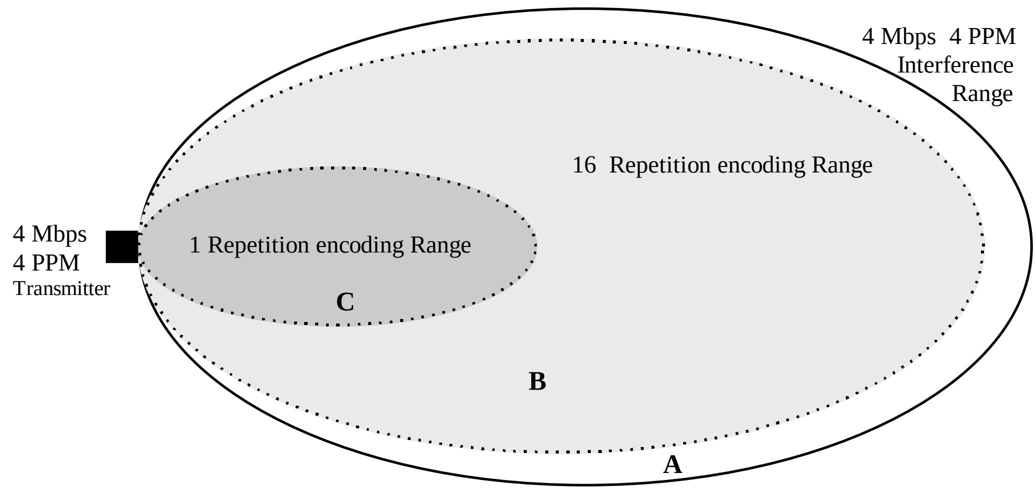

Figure 3.5: Interference/range problem [4]

If a receiver is within region C it will be able to reliably receive data transmissions at RR=1 from the AIr transmitter. However, if the receiver is inside region A but outside region C, then the transmitter can interfere with the receiver even though the receiver cannot decode those transmissions. Region B shows the region of reliable communications range for a 16 repetition encoded data stream.
In order to minimize the interference problem, the control information required by the AIr MAC protocol is transmitted at 16 repetition encoding. This will allow the control information to be communicated over an area almost equal to the interference range.

3.2.2 Infrared Channel Reciprocity and Transceiver Parity

Infrared Channel Reciprocity (ICR) implies that the propagation characteristic between two AIr transceivers is symmetrical, irrespective of the propagation environment. Radio channels usually benefits from channel reciprocity since a common antenna is used for transmission and reception. In contrast, infrared systems use two different elements with different properties: Infrared Emitting Diodes (IREDs) for transmision and Photodiodes (PDs) for reception. Therefore ICR has to be established, using the concept of transceiver parity (PT). The transceiver parity PT is divided in two requirements, optical parity (PO) and system parity (PS).

Optical Parity (PO)

The optical parity requirement specifies an approximate angular match of IRED emission characteristic and PD reception characteristic. Figure 3.6 shows the angular emission and reception characteristics of an optical port in the horizontal plane.

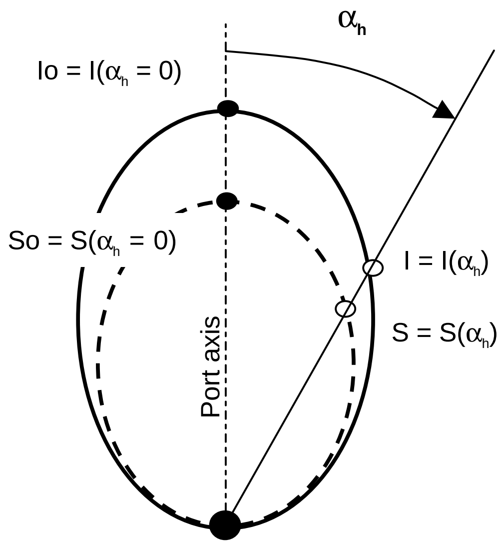

Figure 3.6: Optical port: intensity and sensitivity [5]

The radiant intensity I in [mW/sr] of the optical emitter and the sensitivity S in [µA/mW/cm^2] of the optical detector correspond to measured peak values for optical test pulses with 125 ns width. Optical parity PO is defined as the dimensionless ratio of relative radiant intensity I/Io and relative sensitivity S/So of the optical port. It should be equal to 1:

$$PO = \frac{I / I_0}{S / S_0} = 1 \pm tolerance$$

System Parity (PS)

System parity (PS) requires that the magnitude of signal radiant intensity emitted in the direction of the port axis, and the required magnitude of signal irradiance incident in the direction of the port axis for achieving a given BER is in a constant relationship. The required irradiance is a measure of the receiver system sensitivity. PS is defined as the product of radiant intensity Io in [mW/sr] emitted in the direction of the optical port axis, and required irradiance Eo in [μW/cm^2] incident in the direction of the port axis for achieving BER=10^{-8} at a data rate of 4 Mbps. It should be equal to 300 [mW/sr] [μW/cm^2]:

$$PS = I_0 \times E_0 = 300[mW/sr][\mu W/cm^2] \pm tolerance$$

Transceiver Parity (PT)

The resulting transceiver parity (PT) is defined as the product of the term PO and the term PS:

$$PT = PS \times PO$$

AIr transceivers with drastic deviations of PT could cause a noticeable hidden terminal problem.

3.3 AIr MAC Protocol (IrMAC)

The AIr medium access control (AIr MAC) protocol is specified in the draft document [7]. The following sections resume its key features.

3.3.1 Data Transfer Methods

Three types of data transfer method between AIr devices are defined by AIr MAC, each providing varying degrees of reliability, overhead and functionality:

- Unreserved
- Announced
- Reserved (with or without data acknowledgement)

Unreserved Data Transfer Method

This method is simple and has the lowest overhead, but has to be used sparingly because of its high unreliability. The source device first listens to the channel to see if any other station is transmitting at that moment. If the channel is busy, the source device waits until the channel becomes idle. When it detects an idle channel for a certain period of time, it is allowed to transmit one single data frame. No information is exchanged between the stations to coordinate the access to the channel and the hidden terminal problem cannot be avoided. Turning back to Fig. 3.1, station C could start a transmission while station A is transmitting its data packet to station B, causing a collision at station B (which is able to hear both A and C).

Announced Data Transfer Method

The announced data transfer method increases the reliability and allows multiple frame (burst) transfers at different rates (different RR). However, this method also introduces a protocol overhead and, like the unreserved method, it is not able to deal with the hidden terminal problem. The source device announces that it is going to transmit a data burst by emitting a Request-To-Send (RTS) frame at the common rate, detailing the amount of time during which the device intends to transmit and the rate used for the data transmission. Only the devices within range of the source device hear the RTS and will remain silent during the whole transmission, thus avoiding collisions. After emitting its RTS frame, the source device has to wait a given amount of time before beginning its data transmission. If the source device hears an RTS during this pause, it assumes that there is contention for the medium and no data frame is transmitted. Instead, the contention is resolved with a backoff and retry mechanism. If the source device hears no RTS during the pause, it is allowed to transmit its data burst. At the end of the transmission, the source device emits an End-Of-Burst (EOB) frame to inform all devices within range that the medium is now free.

Reserved Data Transfer Method

The reserved data transfer method is intended to be the primary method for AIr and has been used for all performance study measurements presented in this report. It is the only method that can deal with hidden terminals. All devices within RR=16 range of only the source device, of only the destination device, or of both source and destination devices, are informed that there is a pending data transfer and that they should remain silent to avoid causing interference. The source device sends an RTS to the destination device, detailing the amount of time it needs for the data transfer and the data transfer rate. The destination device responds with a Clear-To-Send (CTS) frame, which also details the length of the data transfer and its rate. All devices within range of source and/or destination will hear this RTS/CTS exchange and remain silent for the reserved period. Upon reception of the CTS frame, the source device is allowed to send its data burst (one or more data frames). After the last data frames, the source device will transmit an EOB frame to which the destination will respond with an End-Of-Burst-Confirmed (EOBC) frame. The EOB/EOBC handshake informs all devices within range of source and/or destination device that the transmission is over. The reserved data transfer method is illustrated in Fig. 3.7.

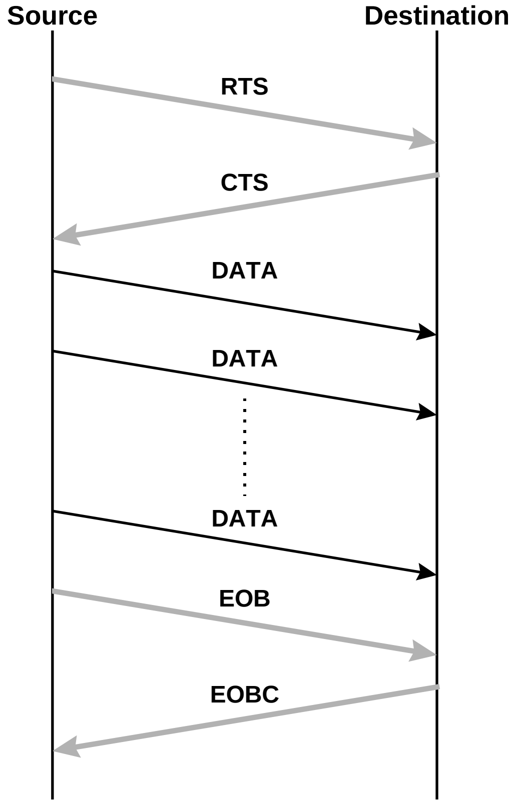

Figure 3.7: Reserved Data Transfer Method (without data acknowledgement) [7]

The data transfer only proceeds following the successful exchange of matching RTS and CTS. A slotted random backoff algorithm is used to resolve contention for the medium manifested by incomplete RTS/CTS exchanges. It will be presented in section 3.3.6.

The reservation can also be terminated without making use of the EOB/EOBC handshake, only relying on the reservation termination (RT) timer. If this reservation termination path, also called RT termination, is adopted, the stations involved in the reservation simply wait until the RT timer has expired and then consider that the reservation is over, without transmitting or receiving any EOB or EOBC frame. The reservation termination method that has been finally proposed in the AIr standard is the EOB/EOBC handshake. However, the RT termination has also been investigated with the performance measurements.

The source device can request that the destination device acknowledge receipt of every data frame sent within the burst (ACK-DATA/ACK handshaking). After sending a so-called ACK-DATA frame, the source waits for the destination device to respond with an Acknowledgment (ACK) frame before sending the next ACK-DATA frame. The destination only sends an ACK frame if the corresponding ACK-DATA frame has been successfully received (i.e., without CRC error). If the source does not receive the ACK frame, it ends prematurely the reservation emitting an EOB frame.

The AIr MAC reserved method also comprises the Reserved Sequenced Data Transfer and the Reserved Groupcasting. However, these two transfer methods are not considered in this report because they are not implemented in the prototype that has been used for the performance measurements.

3.3.2 General Frame Format

The general frame format is shown in Fig. 3.8.

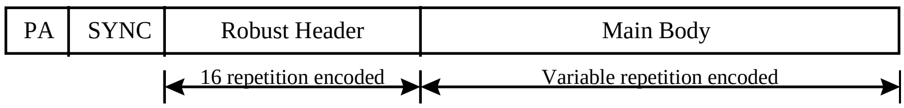

Figure 3.8: General format of an AIr MAC frame [7]

The general AIr MAC frame is composed by four fields: the preamble, the synchronization field, the robust header and the main body:

Preamble (PA)
The preamble field consists of a well-known 256-bits sequence that allows the receiver Phase Locked Loop (PLL) to synchronize with the incoming data stream. It also allows for reliable carrier sensing. Transmission of the preamble requires 64 μs.

Synchronization Field (SYNC)
The synchronization field is used by the receiver to determine the exact start of the robust header, even if it does not hear the whole preamble. Its length is 160 bits. Transmission of the synchronization field requires 40 μs.

Robust Header
The robust header is used to identify the type of the AIr MAC frame being transmitted and convey parameter values associated with the AIr MAC protocol flow. Its fields vary depending on the specific frame type being transmitted but its length is constant (32 bits). The robust header is sent at a repetition rate of RR=16, in order to allow this information to be received at approximately the interference range of the device. Transmission of the robust header requires 128 μs.

Main Body
The main body is variable in length and is not present on the CTS frame. Unlike the robust header the main body of the frame is transmitted at a variable repetition rate.

3.3.3 Specific Frame Formats

There are eighth AIr MAC frame types (three new frame types, S-DATA, S-Poll and SACK, have been defined recently but will not be presented here). The Request-To-Send (RTS), Clear–To-Send (CTS), End-Of-Burst (EOB) and End-Of-Burst-Confirmed (EOBC) frames are used to initiate, contend, qualify and terminate reservations of the medium. The unreserved data (UDATA), unacknowledged data (DATA), acknowledged data (ACK-DATA) and the Acknowledgement (ACK) frames are used in the transfer of service user data between AIr MAC entities. Fig. 3.9 illustrates the different frame formats (the length of the fields is in bit).

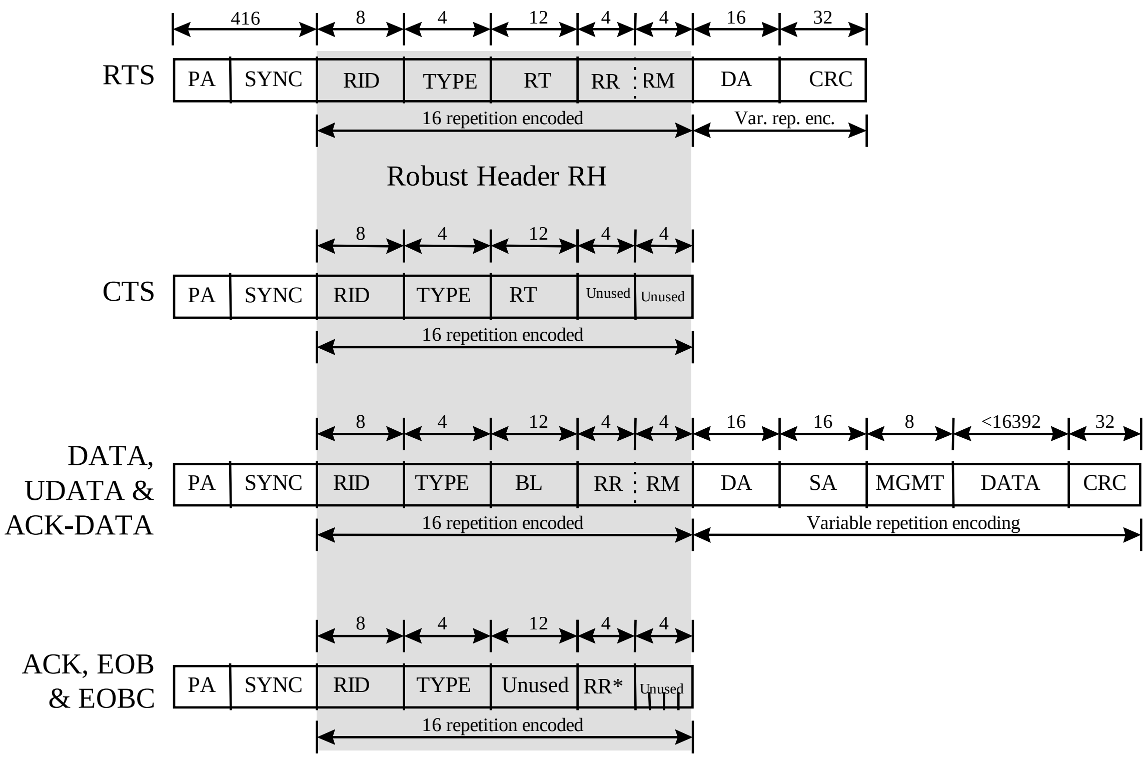

Figure 3.9: AIr MAC frame formats

RID Field

This 8-bit field specifies the Reservation Identifier of a reservation attempt or an on-going reservation. The RID is selected randomly each time a device attempts to reserve the medium. It is used to associate the various AIr MAC frames with a particular reservation and allows the AIr MAC to replace the destination and source addresses in all frames except the RTS frame. The RID acts as a shorthand notation for the source and destination addresses for the duration of a single reservation.

TYPE Field

This is a 4-bit field that defines the type of frame. The types defined so far are Request-to-Send (RTS), Clear-to-Send (CTS), Unreserved Data (UDATA), Data frame (DATA), Acknowledgeable Data (ACK-DATA), Acknowledgement (ACK), End-of-Burst (EOB) and End-of-Burst-Confirmed (EOBC).

BL Field

This 12-bit field defines the length of the main body of the frame in terms of number of bytes. It is used only in UDATA, DATA and ACK-DATA frames to indicate the length starting from the destination address field (DA) up to and including the CRC field, i.e., it indicates the number of bytes in the variable repetition encoded section.

RT Field

This is a 12-bit field that specifies the length of the time being reserved. It is present in both the RTS and CTS frames. Although the termination of the reservation is achieved by observing the EOB and/or EOBC frames, the RT field is used to terminate a reservation in situations where overlapping reservations and collisions prevent this occurring normally. The RT field specifies the time between the start of the initial DATA or ACK-DATA frame and the time up to and including any guard time after the EOBC has finished transmission.

RM Field

This is a 4-bit sequentially encoded field that specifies the base rate and modulation format of the data transmissions. It is placed here to allow for future expansion and technological advancements.

RR Field

This is a 4-bit field and is only present in the RTS, UDATA, DATA and ACK-DATA type frames. It specifies the rate reduction (RR) used for the main body of that frame, i.e. the DA, SA, MGMT, DATA and CRC fields.

RR* Field

This 4-bit field is similar to the RR field in the DATA frames except that it specifies recommendations for the RR in the reverse direction based on an evaluation of the link quality or some other indication. In other words it is a recommendation by the destination device that the source device should use this parameter when communicating with it. The RR* field is only present on the ACK and EOBC frames.

Address Fields (SA & DA)

These fields are the 16-bit Source (SA) and Destination (DA) AIr MAC device addresses of the frame.

MGMT Field

This 8-bit field is used to carry control information in each UDATA, DATA or ACK-DATA type frame.

DATA Field

This field is variable from 0 to 16392 bits (i.e. 2048 + 1 bytes) and contains data provided by the AIr MAC service user.

CRC Field

This is a 32-bit field containing the frame check sequence (FCS) resulting from to the IEEE 802 CRC32 algorithm. The CRC is a calculated payload data dependent field.

3.3.4 AIr MAC Finite State Machine

The AIr MAC Finite State Machine (FSM) specifies the protocol operation, implementing Reserved Data Transmission, Acknowledged Data Transmission, Unreserved Data Transmission, Announced Reservations and IrDA 1.x Transmission. The state table that fully describes the protocol operation is presented in the protocol specification document [7]. The FSM diagram for the reserved data transmission method is illustrated in Fig. 3.10.

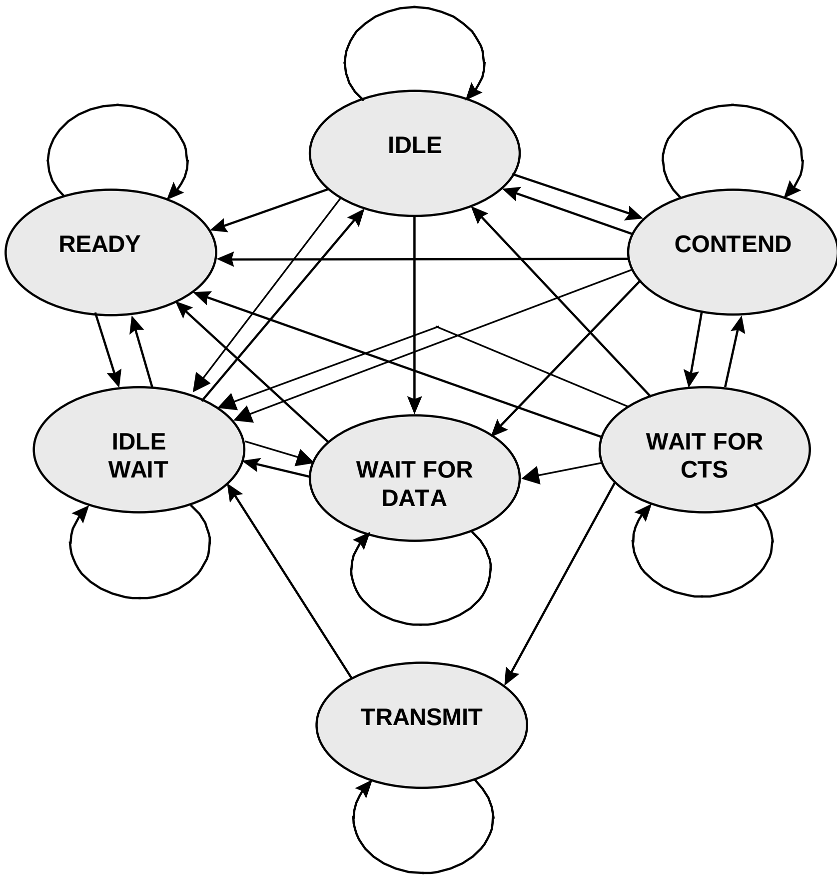

Figure 3.10: Finite State Machine diagram for Reserved Data Transfer Method

There are seven states, IDLE, CONTEND, WAIT FOR CTS, TRANSMIT, WAIT FOR DATA, READY and IDLEWAIT. They can be defined as follows:

- IDLE: The medium is currently idle and the device is waiting to receive or transmit data.
- CONTEND: The device has data to transmit. If the reserved data transfer method is being used, the device has determined the amount of backoff time it is going to wait before emitting its RTS frame, and is now waiting for this time to expire.
- WAIT FOR CTS: The device has transmitted its RTS to a specific device and is now waiting for the CTS response.
- TRANSMIT: The device has reserved the medium: it has received a matching CTS and is now in the process of transmitting a data burst. If it is transmitting ACK-DATA, then it will also process the received ACK frames.
- WAIT FOR DATA: The device has received an RTS and, where appropriate, has replied with a CTS. It is now waiting for the first data frame of the burst.
- READY: The device is either actively receiving data or waiting for a data transmission between two other devices to end.
- IDLEWAIT: This state acts as a synchronization point to allow all devices within the range of a reservation to transition from this state to the IDLE state at the same time (backoff slot synchronization).

3.3.5 AIr MAC Timers

The AIr MAC implementation uses several timers:

- Backoff (or Collision Avoidance, CA) Timer: it is used in the CONTEND state to count down a certain number of backoff slots before transmitting the RTS frame.
- Response Timer: it is started after the transmission of frames that require a response (i.e., RTS, ACK-DATA and EOB). When this timer expires, the station gives up getting the expected response.
- WFData Timer: it is started after receiving an RTS and responding with the CTS. When this timer expires, the station gives up waiting for the reservation to be confirmed with an appropriate DATA or ACK-DATA.
- WFDataN Timer: it is started after the reception of a RTS that is not destined to this station. When this timer expires, the station gives up waiting for the reservation to be confirmed with an appropriate DATA or ACK-DATA.
- MTD (Min Turn Delay) Timer: it loads the minimum amount of time required by the modem after a transmission in order to be able to receive a frame (turnaround time). It is started after receiving a frame that requires a response, to make sure that the response will be received properly.
- RT Timer: it is started by sources, destinations and non-participant devices upon the reception of a CTS frame or a DATA (or ACK-DATA) frame and is used to time the length of any reservation of the medium. It is used as another way of terminating the reservation should the EOB/EOBC be lost.

3.3.6 Backoff Algorithm

The adopted backoff mechanism can be explained as follows. Whenever a station has data to transmit, it randomly throws an integer value between 1 and the current backoff window size (BOWS). This thrown backoff value, also called Collision Avoidance Slot (CAS) value, indicates the number of backoff slots the station has to wait before transmitting its RTS frame. After having thrown the backoff value, the station goes from the IDLE to the CONTEND state, loads the backoff value into the backoff timer and start the backoff timer, which is decremented by one after each backoff slot tick.
The length of a backoff slot is defined as being grater than the transmission time of the RTS plus the turnaround (MTD) time plus the time required to detect the beginning of the returning CTS (the PA and SYNC transmission times). This size is chosen to be large enough to allow stations that are hidden from either the source or the destination device to hear at least one side of the exchange. The worst-case transmission time for an RTS frame (see RTS frame format in Fig. 3.9) is the following:

$$TT_{RTS} = TT_{PA+SYNC} + TT_{RH} + TT_{DA+CRC} = \frac{416 bits}{4 \cdot 10^6 bits/s} + \frac{32 bits \cdot 16}{4 \cdot 10^6 bits/s} + \frac{48 bits \cdot 16}{4 \cdot 10^6 bits/s} = 424 \mu s .$$

The turnaround time (MTD time) has been determined to be MTD = 200 μs. The time required to detect the PA and the SYNC fields is $TT_{PA+SYNC}$ = 104 μs. It follows that the backoff slot time should be:

$$TT_{Slot} \geq TT_{RTS} + MTD + TT_{PA+SYNC} = 424 \mu s + 200 \mu s + 104 \mu s = 728 \mu s$$ (see Fig. 3.11).

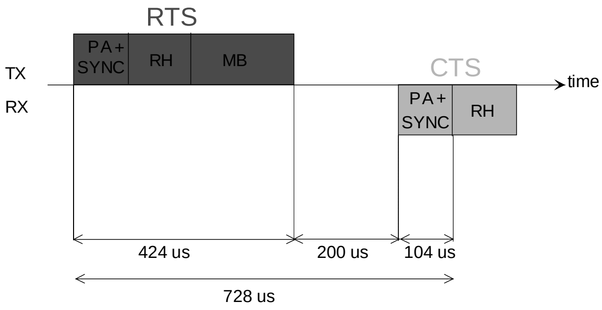

Figure 3.11: Minimum backoff slot time

However, the AIr prototype that has been used for the performance measurements implements an older specification of the protocol, in which the station address length were 32 bits (instead of 16 bits) and the source address was also transmitted in the RTS frame body. For this older protocol specification, the worst case RTS transmission time is TT_{RTS} = 616 μs (TT_{RTS} = 488 μs without CRC) and the backoff slot time should be $TT_{Slot} \geq 920$ μs ($TT_{Slot} \geq 792$ μs without CRC).

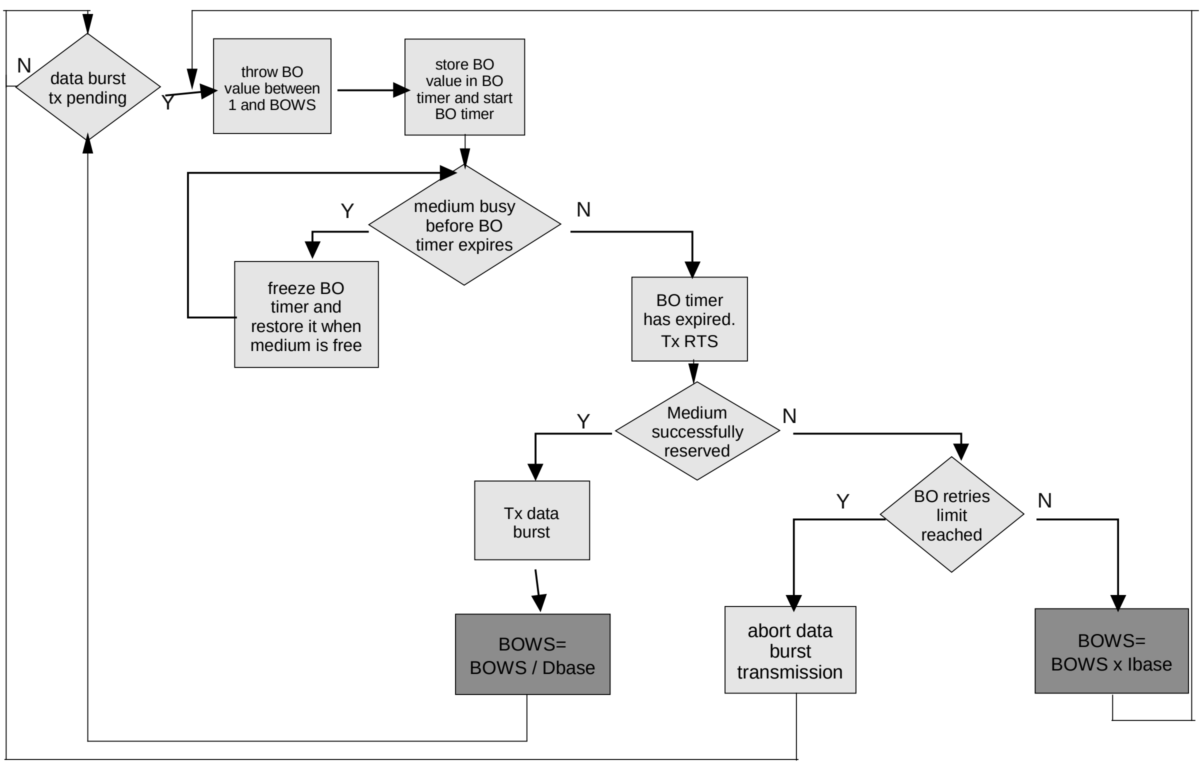

Figure 3.12: Backoff algorithm

If a station senses the medium to be busy while it is in the CONTEND state waiting for the backoff timer to expire, the station freezes the backoff timer. It can then temporarily enter into non-participant mode and wait for a data transmission between two other devices to be over. When the medium becomes available again, the station resumes the backoff from the value which was in the backoff timer at the time it was frozen.

The algorithm that has been chosen for calculation of the backoff window size, $BOWS_i$, of the i'th channel access attempt is the following:

$$BOWS_i = BOWS_{i-1} \cdot \left\{ ID \cdot I_{base} + (1 - ID) \cdot \frac{1}{D_{base}} \right\}$$

where I_{base} is the factorial base for increasing the backoff window size, D_{base} is the factorial base for decreasing the backoff window size and where ID = 0 if the (i-1)'th channel access attempt was successful and ID = 1 if the (i-1)'th channel access attempt was unsuccessful.
In the default backoff algorithm, $I_{base} = D_{base} = 2$: if the last RTS transmission attempt did not result in a successful reservation, BOWS is multiplied by 2; if the last RTS transmission resulted in a successful reservation, BOWS is divided by 2.
Fig. 3.12 illustrates the complete backoff policy.

3.3.7 Backoff Slot Synchronization

Slot synchronization is the notion that all stations backoff slot clock generators are in phase with each other. The backoff slot clock generators are the counters that generate the clock that drives the backoff timer (see Fig. 3.13).
The resynchronization of the backoff slot clock generator at the end of a reservation is possible only at all the non-participant devices that are within RR=16 range of the source device and/or of the destination device.
The slot synchronization occurs during the transition from IDLEWAIT to IDLE or upon reception of EOB or EOBC with non-matching RID in the READY state. Due to the slot synchronization mechanism, RTS-RTS collisions will be always launched with near perfect overlap.

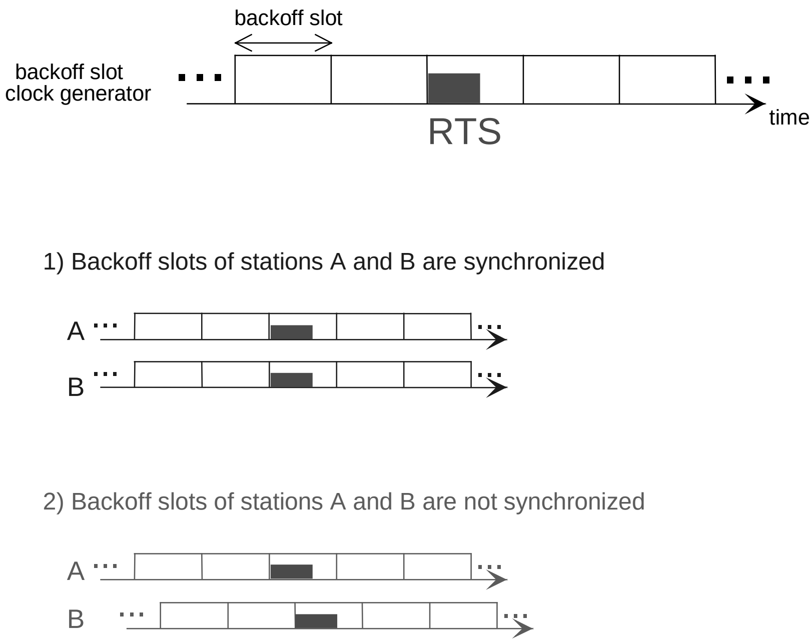

Figure 3.13: Backoff slot synchronization

3.4 Tools for AIr MAC Performance Study

The following tools have been used for the AIr MAC protocol performance study:

1. **Prototype hardware and software:** four AIr stations, equipped with AIr prototype hardware and software, have been used. An AIr station consisted of an IBM ThinkPad 760 on a docking station, a prototype board (with integrated VR-Modem, AIr MAC FSM controller and ISA interface), its prototype software driver and an AIr transceiver. Each station involved in a test can be configured using the software utility that displays all the user configurable test parameters and allows to change them to customize the test. A tutorial which explains how to prepare and configure an AIr station is proposed in Appendix A. At the end of each test some information is displayed by the software utility, e.g., the throughput and frame drop/loss of the station. Some other information can be obtained by connecting a logic analyzer (or a scope) to some prototype board FPGA pins. The logical signal TX_ACTIVE, which is in the high voltage state when the station is in the process of transmitting, has been measured during the tests. The eighth AIr MAC state and event vector pins, which reflect almost all signals connected to the FPGA chips and are there for testing and debugging, have also been connected to a transitional based logic analyzer.

2. **Theoretical model:** Dr Schott has mapped the Super Finite State Machine (version 0.5) of the AIr MAC protocol into a SDL-based protocol specification using the software tool SDT. SDL is the formal Specification and Description Language recommended by the ITU, and SDT is Telelogic's software tool which provides a SDL-to-C development environment for graphically specifying, analyzing, simulating, and validating protocols. The defined SDL/SDT model for the AIr system can comprise up to 12 dynamically attachable AIr nodes. The behavior (MAC layer) of each AIr node is defined by several SDL processes. The physical transmitter and receiver were modeled so that the physical layer service specification is satisfied and transmission errors and delays can be introduced. An AIr MAC service user was defined as a SDL process which can transmit and receive frames using the MAC service primitives. However, it is very difficult to model some physical layer features, as the capture effect.

4. AIr MAC Performance: Preliminary Measurements

The goal of the preliminary measurements is to investigate how the AIr MAC performance (when the Reserved Data Transfer method is used) can be influenced by the following protocol options and parameters:

- Reservation termination path (RT termination or EOB/EOBC handshake)
- Data acknowledgement (use of DATA or ACK-DATA)
- Backoff algorithm (backoff window size adjustment policy, maximum number of backoff retries)
- Size of a burst (number of packets per burst and length of a data packet payload)

4.1 Performance Evaluation Parameters

The following performance parameters, calculated by each station during and/or at the end of a test, have been considered:

1. Transmitter Throughput:

The transmitter throughput (TP) of an AIr station is defined as follows:

$$TP = \frac{N_{Tx}}{T_{LastTx} - T_{FirstTx}}.$$

N_{Tx} is the total number of successfully transmitted data bits. T_{LastTx} and $T_{FirstTx}$ are, respectively, the times when the end of the last data packet and the begin of the first data packet was successfully transmitted.
Even if the AIr transmission rate is 4 Mbps, the throughput of a station will always be lower due to the MAC protocol overhead (time wasted for channel contention and transmission of AIr MAC control frames).
The system (or aggregate) throughput is the sum of the individual throughputs. The throughput of a particular station, say station A, will be denoted as TP_A.

2. Frame Drop

The frame drop (FD) of an AIr station is defined as follows:

$$FD = \frac{N_{TxTotal} - N_{TxSuccess}}{N_{TxTotal}}.$$

$N_{TxTotal}$ is the total number of data packets the station has to transmit to its peer and $N_{TxSuccess}$ is the number of successfully transmitted data packets.
A station discard a data burst when it has not been able to transmit it after the maximum number of backoff retries has been reached. If one or more data burst have been discarded, $N_{TxTotal} > N_{TxSuccess}$ and FD > 0%.
The frame drop of a particular station, say station A, will be denoted as FD_A.

3. Frame Loss:

The frame loss of a station-to-station link (FL) is defined as follows:

$$FL = \frac{N_{TxSuccess} - N_{RxError-free}}{N_{TxSuccess}}.$$

$N_{TxSuccess}$ is the number of data packets successfully emitted by the source station and $N_{RxError\text{-}free}$ is the number of error-free (i.e., CRC is OK) data packets successfully arrived at the destination station.
The frame loss depends on the link quality (without interference) and on the signal-to-interference ratio at the receiving station when interference is present.
The frame loss of a particular link, say link A->B, will be denoted as FL_{AB}.

4.2 Test Parameters

Some common parameters have been chosen and kept unchanged for all experiments. Other parameters have been varied to test the various AIr MAC options.

4.2.1 Constant Test Parameters

The constant test parameters listed in the following table have been used in all experiments:

Number of Transmitted Frames	20006
Packet Payload Size	2048 Bytes
Min Turn Delay Time (MTD)	200 μsec
Backoff Slot Time (T_{slot})	900 μsec
Response Timer	636 μsec
WFDataN Timer	904 μsec
WFData Timer	468 μsec
Minimum Backoff Window Size	8

The following rules and options have also been adopted:

- Each transmitting station sends packets to only one other station. Thus, the SrcStationNumber (SSN) of each station in a test must be unique. The same goes for the TrgStationNumber (TSN).
- RR values are fixed, i.e. no rate selection algorithm is being used. Data packets payloads are always transmitted with fixed RR=1.
- The RTS body contains the destination and the source addresses (no CRC field is present). It is always sent at RR=16. The address size is 32 bits, even if in the new AIr proposal it should be 16 bits.
- The backoff slot synchronization mechanism is always active
- Configuration geometry is static (no mobility).
- Whenever a station has completed the transmission of all its data packets, the whole test should be terminated.

4.2.2 Variable Test Parameters

The following parameters have been varied:

Reservation termination	EOB or RT
Use of Data acknowledgement	DATA or ACK-DATA
Packets Per Burst (PPB)	1 to 7
Backoff Retries (Max_Retries)	1 to 9
BOWS Algorithm parameter I_{base}	1 or 2
BOWS Algorithm parameter D_{base}	1 or 2

The adopted backoff algorithm is presented in section 3.3.6.

If $I_{base} = D_{base} = 2$, the default backoff algorithm is used and the backoff window size (BOWS) is variable. Each station selects a random number between 1 and the current BOWS to determine the number of collision avoidance slots for which it will wait before emitting its RTS. If the last RTS transmission did not result in a successful reservation, BOWS is multiplied by 2. If the last RTS transmission did result in a successful reservation, BOWS is divided by 2. The maximum BOWS is equal to $8 \cdot 2^{Max\_Retries}$.

If $I_{base} = D_{base} = 1$, BOWS is frozen to its minimum value 8, no matter what Max_Retries is. Each station always selects a random number between 1 and 8 to determine the number of collision avoidance slots for which it will wait before emitting its RTS.

The measurements have been done for four different protocol paths:

- RT/NACK: neither data acknowledgements nor EOB/EOBC handshaking are used
- RT/ACK: EOB/EOBC handshaking is not used; data acknowledgements are used
- EOB/NACK: EOB/EOBC handshaking is used; data acknowledgements are not used
- EOB/ACK: both data acknowledgements and EOB/EOBC handshaking are used

4.3 Experimental Test Models

The test models proposed for the experimental performance evaluation are illustrated in Fig. 4.1.

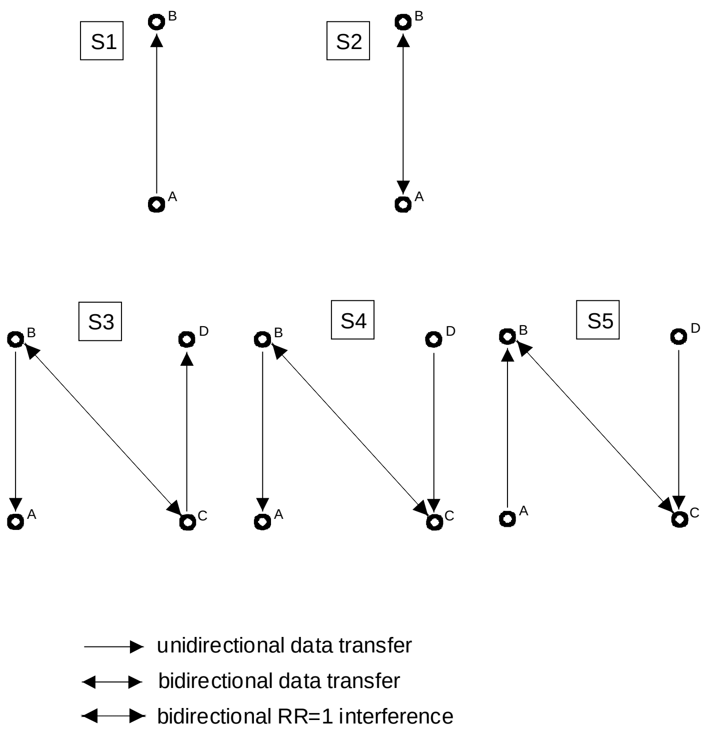

Figure 4.1: Experimental test models S1 to S5

A circle represents a wireless AIr station. A thick solid line represents a logical connection over which two stations exchange data and the arrow(s) indicate(s) the direction of the data flow. The destination station is in the data range (RR=1 range) of the source station, so that data can be received error-free if no interference is present. A thick dashed line links two stations that are within RR=1 range of each other (i.e., data with RR=1 could be transmitted error-free between the two stations) and could interfere with each other during the tests. If neither a thick solid nor a thick dashed line links two stations, then these stations cannot hear each other (they are completely unaware of each other).
The measurements results obtained for the five scenarios S1 to S5 are presented in the next sections.

4.4 Test Model S1

In this simple model, a unidirectional data transmission from station A to station B takes place. There is no contention for the channel: each time the IrLC layer indicates that a reservation is pending, station A selects a random number between 1 and 8 (the minimum BOWS) in order to determine the number of backoff slots for which it will wait before transmitting its RTS. Each RTS transmission will always result in a successful reservation. Due to the lack of RTS collisions, the BOWS always corresponds to the minimum value 8 and the average backoff period is 4.5 slots. The parameters I_{base} and D_{base} have no influence on the measurements because, as stated above, BOWS is never modified (neither incremented nor decremented). Also, $FD_A = FL_A = 0$.
The throughput measurements that have been carried out are in good agreement with the ones obtained in IMD Toronto, and are illustrated in Fig. 4.2.

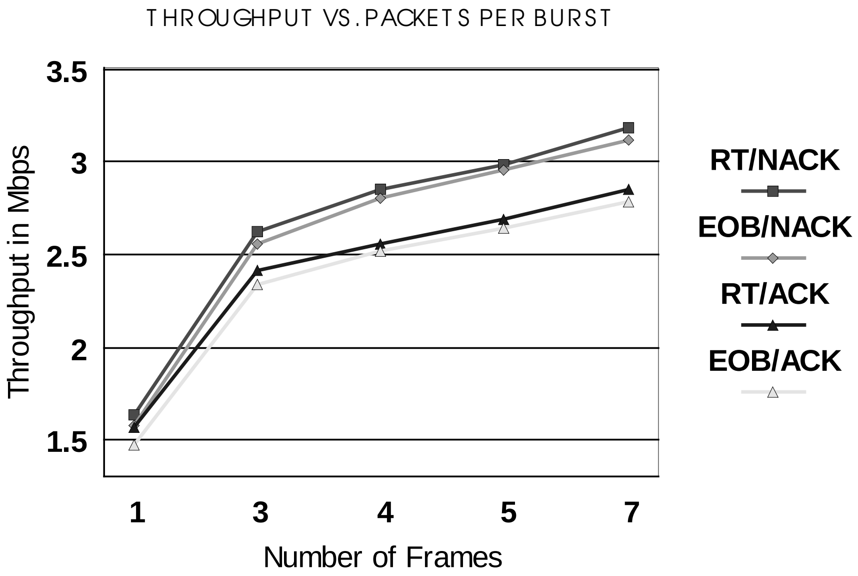

Figure 4.2: Test Model S1: Throughput vs. packets per burst

Comments:

1. For all protocol paths, the throughput increases as the number of packets in a reservation burst (PPB) increases. It is clear that the transmitter can improve its throughput if, after waiting for the backoff period to send a RTS and get the channel reserved, it is allowed to keep it for a long time transmitting a large number of packets in its reservation burst.
2. Both the acknowledgment overhead and the EOB/EOBC handshake overhead lower the throughput. The throughput improvement resulting from using the RT termination instead of the EOB termination is 30-90 kbps. The throughput improvement resulting from using DATA instead of ADATA (for the same reservation termination) increases as PPB increases and is much more significant: with PPB=7 it is about 340 kbps.
3. The RT/NACK path performs better than the others in the considered scenario. In the absence of interferers both the EOB/EOBC handshaking and the data acknowledgments add an unnecessary signaling overhead. The performance difference between RT and EOB/EOBC terminations is very little if compared with the difference between the use of DATA and the use of ACK-DATA.

4.5 Test Model S2

In this model, a bidirectional data transmission between stations A and B takes place (see Fig. 4.3)

Figure 4.3: Test Model S2

This is the most controlled form of contention. This scenario has been tested in IMD Toronto to benchmark the bi-directional throughput of two stations with no interference. The system throughput is expected to exceed the one of test model S1, due to the nature of the backoff algorithm which cuts down on the average time each station has to spend in contention before accessing the channel. The measurements confirmed this expectation.
In a first experiment, Max_Retries has been set to 4 and I_{base} = D_{base} = 2. Consequently, the maximum BOWS will be 128. Fig. 4.4 shows the obtained throughput measurements. Note that one to two data burst drops occurred at each station in this experiment ($FD_A \neq 0\%$ and $FD_B \neq 0\%$). Like in S1, the RT/NACK scheme offers the highest throughput in such an environment. Since the backoff slots of the two stations are synchronized, thanks to the exchange of the RTS/CTS frames, the EOB/EOBC frame exchange is not necessary for synchronization and represents an unnecessary signaling overhead.

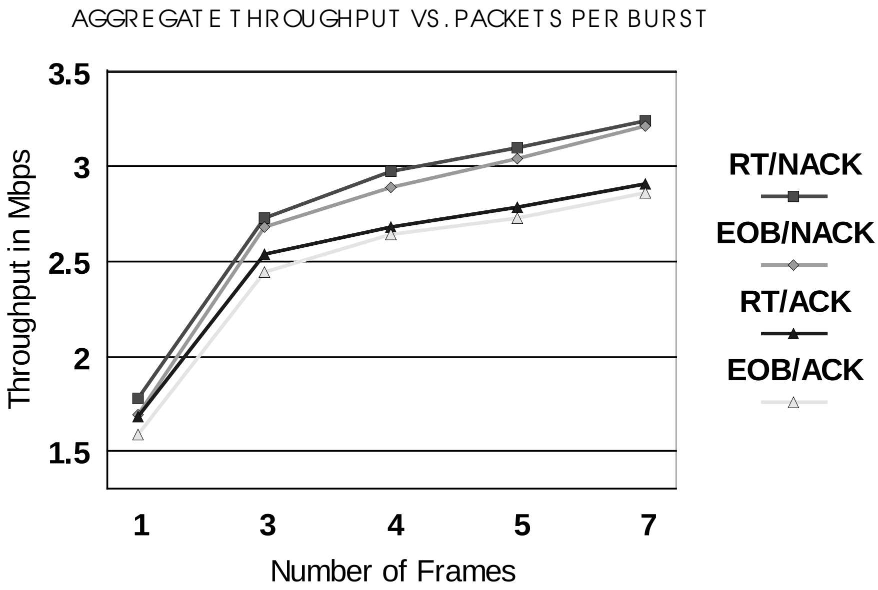

Figure 4.4: Test Model S2: Throughput vs. packets per burst

To get a better understanding on the access scheme and especially the data burst drop effect, Dr. Schott measured the access delay of each burst for each node station with his SDL/SDT model for the AIr system. The access delay is defined as the time difference between the events "RTS_transmitted" and

"Reserved-Data-Pending_raised". In Fig. 4.5 the histogram of the access delay is plotted for each station. The EOB/NACK scheme was used, 7 frames per bursts were transmitted, Max_Retries has been set to 4 and I_{base} = D_{base} = 2. The access histograms for the two stations look similar so the access scheme is fair. However, there are significant access delays up to 1.5 s. This bad effect can be explained by tracing the backoff window size as a function of time (see Fig. 4.6).

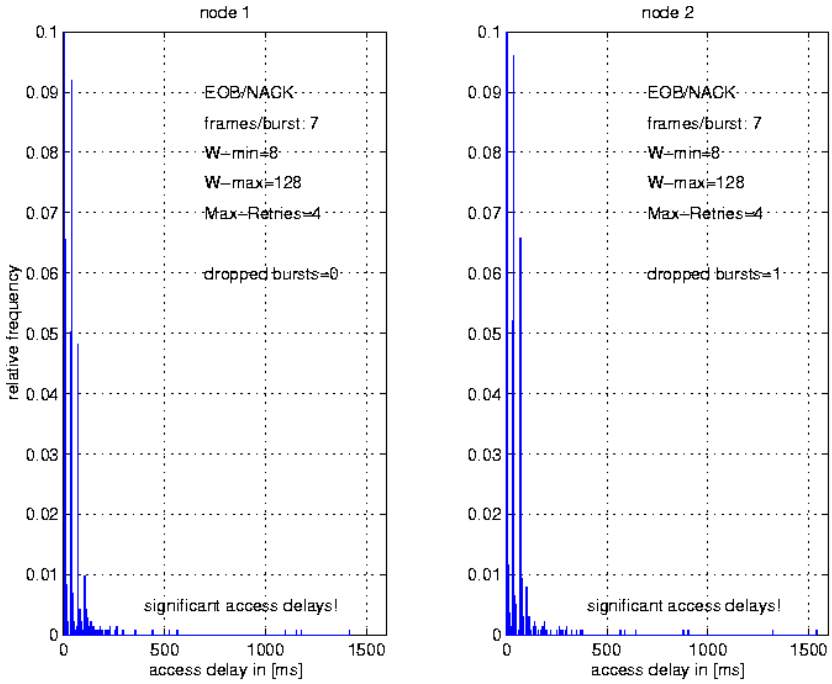

Figure 4.5: Test Model S2: Histogram of access delay

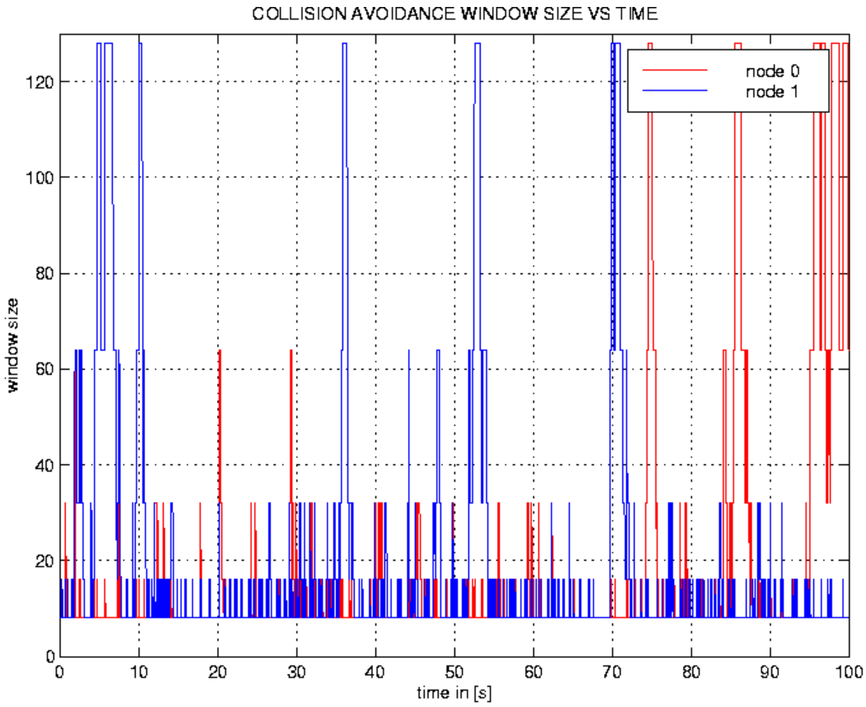

Figure 4.6: Test Model S2: Backoff window size

Figure 4.6 indicates that the maximum value of the backoff window size (128) is reached several times due to collisions at both stations. When the window size is at its maximum value and a collision occurs, the burst is dropped. In addition, when the window size is at its maximum value and a high backoff timer value is selected, it may last a long time until the backoff timer can expire, as shown in Fig. 4.7.

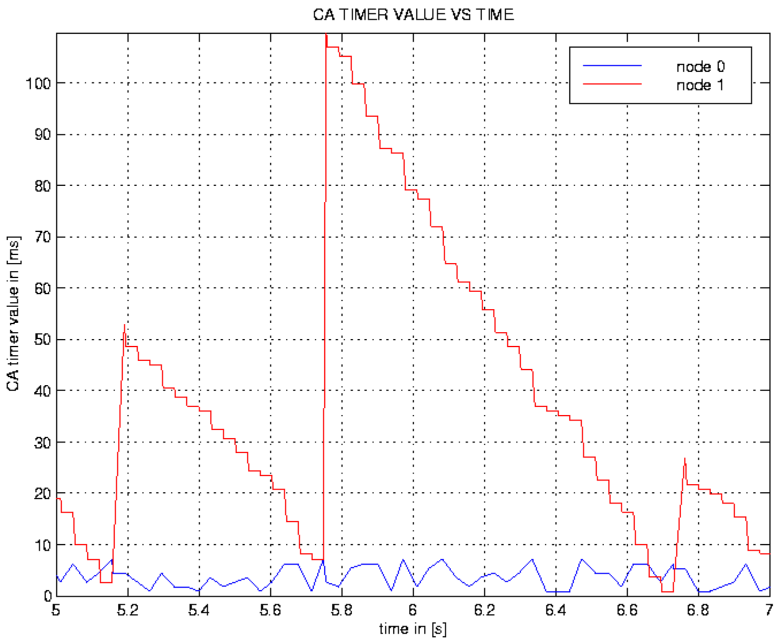

Figure 4.7: Test Model S2: Backoff timer

The previous figures have shown that setting the maximum window size to 128 is not a good choice. If the maximum window size is reduced to a lower value, long access delays should be avoided; however, the risk of collisions is increased.

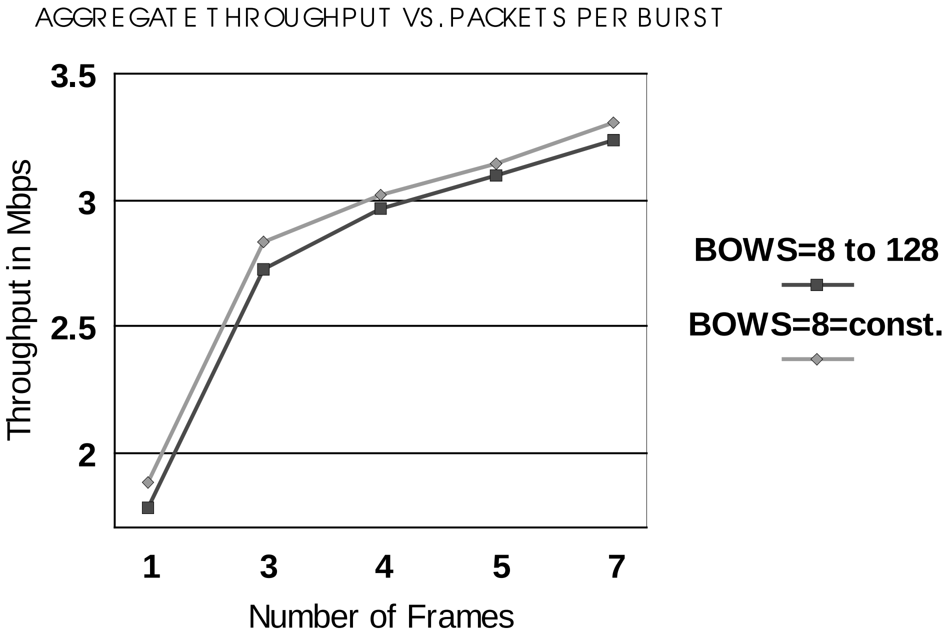

Figure 4.8: Test Model S2: Throughput vs. packets per burst

In a next experiment, the BOWS has been frozen to 8 (I_{base} = D_{base} = 1), Max_Retries has been set to 9, and the aggregate throughput has been measured as a function of the number of frames per burst selecting the RT/NACK scheme.
Fig. 4.8 indicates that the throughput is higher if the window size is not increased. Also, the access profiles shows that long access delays are now very rare.

4.6 Test Model S3

In this model, a unidirectional data transmission from station B to station A and one from station C to station D take place (see Fig. 4.9).

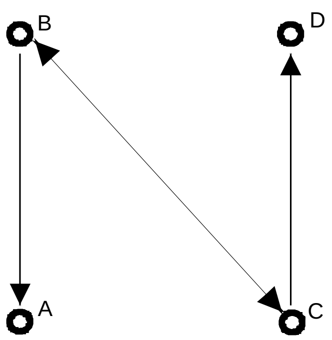

Figure 4.9: Test Model S3

The two transmitters can hear each other and their backoff slots are always synchronized. There is contention for the channel. When both station B and station C are in the CONTEND state, three different events can occur:

1. Station B first emits its RTS. It wins the reservation and station C has to remain silent until station B's reservation is over.
2. Station C first emits its RTS. It wins the reservation and station B has to remain silent until station C's reservation is over.
3. Station B and station C emit their RTS in the same slot. As the destinations A and D cannot perceive the RTS-RTS collision, an overlapping reservation occurs: both stations transmit their data packets simultaneously. As a consequence, no RTS retry is ever used by station B and station C.

It should be noted that, in this particular test model, overlapping reservations could occur without causing any frame loss, as each receiver can only hear the corresponding transmitter. The contention introduced by AIr MAC therefore causes a useless throughput reduction.
Some measurements have been done with $I_{base} = D_{base} = 2$, Max_Retries=4 and PPB=7.
For the EOB/NACK and RT/NACK protocol paths, system throughputs of, respectively, 3.71 and 3.83 Mbps have been measured.
Let's consider the RT/NACK protocol path. If the two transmission links were fully independent (i.e., if also station B and station C could not hear each other), the system throughput would be two times larger than the one obtained for the single transmitter of test model S1: 2 x 3.18 = 6.36 Mbps. On the other hand, if all stations could hear each other, the system throughput would be the same as the one of test model S2, i.e., 3.26 Mbps. The measured system throughput, 3.83 Mbps, situates itself between these two values. It is smaller than 6.36 due to the introduced channel contention (see events 1 and 2 mentioned above) and to the BOWS variation; it is larger than 3.26, due to the existence of overlapping reservations.
As expected, $FD_B = FD_C = 0\%$ and $FL_{BA} = FL_{CD} = 0\%$, in all cases.

4.7 Test Model S4

In this model, a unidirectional data transmission from station B to station A and one from station D to station C take place (see Fig. 4.10).

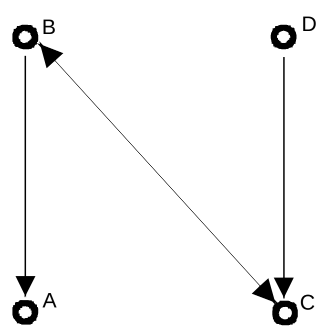

Figure 4.10: Test Model S4

The receiving station C and the transmitting station B can hear each other, but station D is unaware of stations A and B. Because of this unawareness two problems arise:

1. For all protocol paths, the collision avoidance slots of stations B and D cannot be kept synchronized

2. station D gets "locked out" (i.e., it has a high frame drop and a low throughput), if the default backoff algorithm is used ($I_{base} = D_{base} = 2$).

These two phenomena will be explained in the rest of this section.

4.7.1 Backoff Slot Asynchronism

In a four-station scenario with two data links B->A and D->C, the two transmitting stations B and D can keep their collision avoidance slots fully synchronized only in the following cases (where "to hear" means "to be able to receive data encoded at RR=16"):

1. station B and D can hear each other

2. station B can hear station D, and station D can hear station A

3. station B can hear station C, and station D can hear station B

4. station B can hear station C, and station D can hear station A

5. A combination of the conditions 1 to 4.

In other words, two transmitting stations can be fully synchronized only if both can hear at least one station of the other link (be it the transmitting station, the receiving station, or both). This condition is not fulfilled in this test model, because station D cannot hear neither station A nor station B.
In order to investigate the asynchronism, a logic analyzer/scope has been simultaneously connected with the TX_ACTIVE signals (see section 3.4) of stations B and D. In the following pictures, signal 1 corresponds to the TX_ACTIVE signal measured at station D, while signal 2 corresponds to the TX_ACTIVE signal measured at station B.

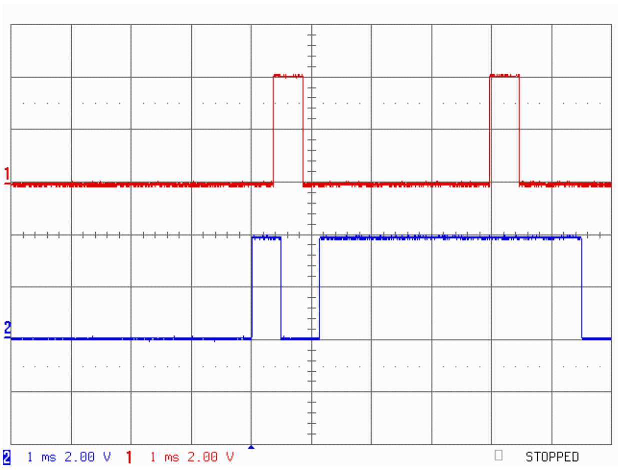

Figure 4.11: Test Model S4: TX_ACTIVE at station D (1) and B (2)

In Fig. 4.11, station B sends a RTS, receives a CTS, and transmits a DATA packet. Station A also sends a RTS, with a time shift of about 400 μs. The backoff slots of station D and B are clearly not synchronous. The RTS emitted by D collides with the one emitted by B, so no CTS is sent back to D. Since station D is not aware of the ongoing data transmission on the other link, D initializes another access request by transmitting a RTS after a random delay.

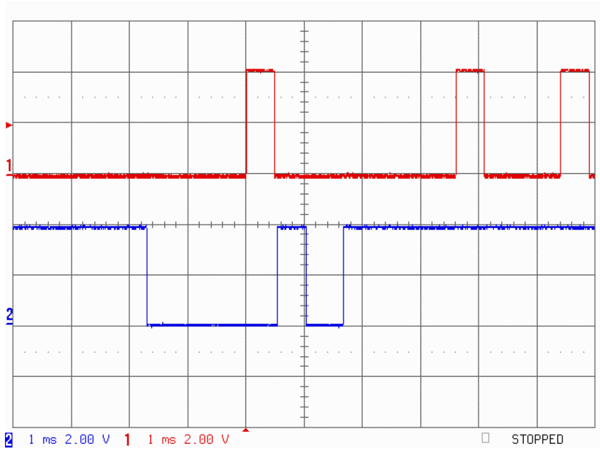

Figure 4.12: Test Model S4: TX_ACTIVE at station D (1) and B (2)

In Fig. 4.12, the time shift between the RTS emitted by station D and the one emitted by station B is about 500 μs. Therefore, the backoff slots of station D and B are asynchronous.
There are two undesirable effects due to the asynchronism that can be seen with the protocol measurements results:

- Some frames (up to 1.2% of $N_{TxSuccess}$) successfully sent by D to C are discarded by C. This happens when C responds with a CTS to a RTS emitted by D and, being in the WAIT FOR DATA state, receives a RTS from B before the first data frame has come from D.

- The frame loss of link B->A is always FL_{BA} = 0% (no station can interfere with station A). The one of link D->C can reach FL_{DC} = 4% with I_{base} = D_{base} = 2, Max_Retries=4 and PPB=1, and FL_{DC} = 8.5% with I_{base} = D_{base} = 1, Max_Retries=9 and PPB=7. This degradation occurs because sometimes B cannot hear a CTS from C because it already is transmitting an RTS.

If the stations B and D were synchronous, none of these effects should occur.

4.7.2 Station Lockout

Some measurements have been done for the EOB/NACK protocol path. If the default backoff algorithm is used (I_{base} = D_{base} = 2), station D gets locked out and the packet transfer on link B->A suffocates the one on link D->C, no matter how the parameters Max_Retries and PPB (Packets Per Burst) are set. While for Max_Retries=4 station B has a transmission throughput TP_B = 2.95 Mbps with PPB=7 and
TP_B = 1.39 Mbps with PPB=1, station D is severely penalized reaching only TP_D = 0.33 Mbps with PPB=7 and TP_D = 0.26 Mbps with PPB=1. The disadvantage suffered by station D is also outlined by its high frame drop. Station B's frame drop is always FD_B = 0%, while station D's frame drop ranges from FD_D = 4% with Max_Retries=4 and PPB=1 to FD_D = 48% with Max_Retries=4 and PPB=7.
The lockout effect can be explained as follows. Station B can always reserve the channel after the first RTS emission (no retries are needed), as station A is completely isolated from link D->C and cannot experience any RTS collisions. Station C can hear the RTS frames emitted by station B and must remain silent when a reservation is in progress. Consequently, station D can reserve the channel only if it emits its RTS while station B is in the CONTEND state waiting for the backoff timer to expire. Being unaware of

stations A and B, station D continuously sends RTS frames to station C while C is waiting as a non-participant for the reservation between A and B to terminate and cannot respond with a CTS. As a result, station D wastes RTS retries and its BOWS (and backoff timer) reaches quickly a large value, so that a lot of time is wasted between two consecutive RTS emissions.
Reducing PPB increases the mean time during which station B is waiting to emit a RTS in the CONTEND state. Therefore, it is more likely that station D can reserve the channel and its throughput should increase. However, the throughput improvement is little.

The only way to eliminate the lockout problem is to change the backoff algorithm, using $I_{base} = D_{base} = 1$. Max_Retries has to be larger than 4, otherwise station D drops too many frames (FD_D = 54% with Max_Retries=4 and PPB=7). The use of a constant backoff window size BOWS=8 eliminates (or, at least, attenuates significantly) the lockout problem:

- Max_Retries=5 to 9, PPB=7: in this scenario TP_B = 1.96 Mbps and TP_D = 1.76 Mbps. However, while station B do not drop any frames, station D's frame drop is FD_D = 21%.

- Max_Retries=9, PPB=1: in this scenario TP_B = 1.02 Mbps and TP_D = 1.05 Mbps. Station D's frame drop is reduced to FD_D = 0.65%. This is a fair scenario, but the price for this fairness is a diminution of the aggregate throughput.

4.8 Test Model S5

In this model, a unidirectional data transmission from station A to station B and one from station D to station C take place (see Fig. 4.13).

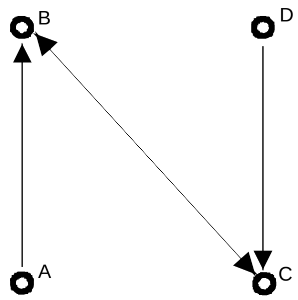

Figure 4.13: Test Model S5

The receiving stations B and C can hear each other, but both transmitting stations can only hear the correspondent receiver. As explained in the preceding section, in such a scenario the transmitters cannot be synchronous: for all protocol paths, the backoff slots of stations A and D cannot be kept synchronized.

4.8.1 Backoff Slot Asynchronism

In order to investigate the asynchronism, a logic analyzer/scope has been simultaneously connected with the TX_ACTIVE signals (see section 3.4) of stations B and D. In the following pictures, signal 1 corresponds to the TX_ACTIVE signal measured at station D, while signal 2 corresponds to the TX_ACTIVE signal measured at station B.

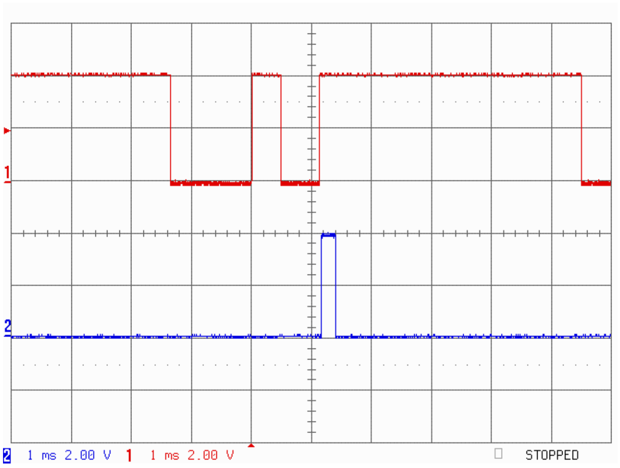

Figure 4.14: Test Model S5: TX_ACTIVE at station D (1) and B (2)

Fig. 4.14 again shows the signal TX_ACTIVE of stations B and D. Note that in this experiment station D transmits RTS and DATA frames but station B only CTS frames. In the figure, station D transmits a RTS frame, receives a CTS frame, and transmits a DATA packet. After receiving a RTS frame, station B answers with a CTS frame, and starts receiving a DATA packet. The emitted CTS frame also collides with the beginning of the DATA packet emitted by station D. The beginning of the RTS frame indicates the start of a backoff slot of station D, while the end of the CTS frame indicates about the end of a backoff slot of station B. Since the time difference between these two events is 1.2 ms, the access slots of the two stations are clearly asynchronous.

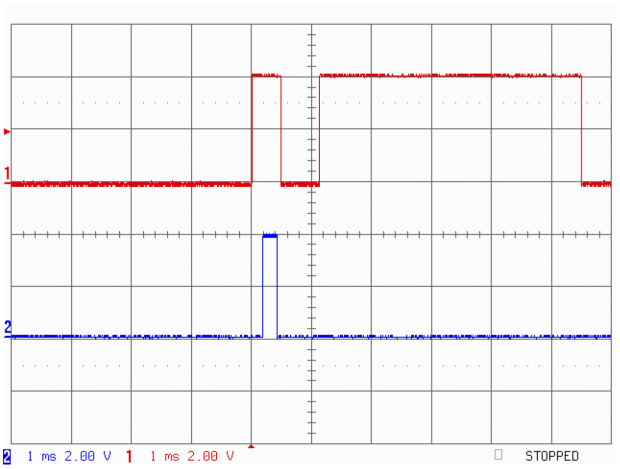

Figure 4.15: Test Model S5: TX_ACTIVE at station D (1) and B (2)

In Fig. 4.15, a collision between a RTS emitted by station D and a CTS emitted by station B is shown. This is again an indication for asynchronism.

There are two undesirable effects due to the asynchronism that can be seen with the protocol measurements:

- Some frames (up to 2.1% of $N_{TxSuccess}$) successfully sent by A (respectively D) to B (respectively C) are discarded by the receiving station. This happens when B (respectively C) responds with a CTS to a RTS emitted by A (respectively D) and, being in the WAIT FOR DATA state, receives a CTS from C (respectively B) before the first data frame has come.
- Frame losses can be observed on both links (particularly for the EOB/NACK protocol path), due to CTS and EOBC frames that collide with data packets. For example, if A sends a RTS and (at the same time) C sends a CTS, then there will be a RTS-CTS collision at B. B cannot decode the CTS and therefore is unaware of the reservation on link D->C. If A sends another RTS, B responds with a CTS which will perturb data arriving at C. Moreover, a reservation begins on link A->B that will be perturbed by the EOBC emitted by C.

Some measurements have been done for the EOB/NACK protocol path.

With $I_{base} = D_{base} = 2$:

1. Max_Retries=4, PPB=7: in this scenario $TP_A = TP_D = 1.64$ Mbps, $FD_A = FD_D = 6\%$ and $FL_{AB} = FL_{DC} = 0.85\%$.
2. Max_Retries=9, PPB=7: increasing the maximum number of retries the frame drop decreases $TP_A = TP_D = 1.65$ Mbps, $FD_A = FD_D = 1.5\%$, $FL_{AB} = FL_{DC} = 0.75\%$.
3. Max_Retries=9, PPB=1: decreasing the number of packets per burst further reduces the frame drop but the throughput decreases and the frame loss increases. $TP_A = TP_D = 0.85$ Mbps, $FD_A = FD_D = 0.1\%$, $FL_{AB} = FL_{DC} = 2.1\%$.

With $I_{base} = D_{base} = 1$:

1. Max_Retries=4, PPB=7: in this scenario $TP_A = TP_D = 2$ Mbps, $FD_A = FD_D = 39.9\%$ and $FL_{AB} = FL_{DC} = 2.85\%$.
2. Max_Retries=9, PPB=7: increasing the maximum number of retries the frame drop decreases. $TP_A = TP_D = 2$ Mbps, $FD_A = FD_D = 3.6\%$ and $FL_{AB} = FL_{DC} = 2\%$.
3. Max_Retries=9, PPB=1: $TP_A = TP_D = 1.04$ Mbps, $FD_A = FD_D = 1.4\%$ and $FL_{AB} = FL_{DC} = 4.5\%$.

4.9 Conclusions

The measurements done with test models S1 to S5 have investigated and highlighted the following issues:

1. **Utility of EOB/EOBC handshake and data acknowledgments:**

   The measurement results obtained with test model S1 shows that, even if the RT termination performs better than the EOB/EOBC termination in terms of throughput, the difference between the two is very little (30-90 kbps). The throughput lowering caused by the EOB/EOBC handshake overhead is counterbalanced by the advantages offered by this reservation termination. Firstly, premature reservation termination is possible (only if data acknowledgements are used) and can improve significantly the performance in the presence of interference. Secondly, the destination device can use the EOBC frame to inform the source device about RR* even if data acknowledgments are not used. Thirdly, backoff slot synchronization is possible upon reception of EOB or EOBC with non-matching RID in the READY state. The EOB/EOBC handshake has been adopted in the latest AIr MAC specification.

The use of data acknowledgements (ACK-DATA instead of DATA), introduces a large overhead that has a significant impact on the throughput performance (about 350 kbps throughput reduction). With ACK-DATA transmission it is possible to prematurely terminate a reservation if the source device does not correctly receive the ACK frame from the destination device. The ACK frames are also used by the destination device to propagate the RR* back to the source device, allowing it to change the repetition rate of the subsequent data frames. Moreover, the reception of an ACK frame allows stations that are blind to the source device but not to the destination device to dynamically join the ongoing reservation as a non-participant. The repetition rate adaptation and the capability of prematurely terminate a reservation can be very useful if the channel quality deteriorates during a reservation and in the presence of interference. ACK-DATA should be used if the performance improvement obtained by these two features compensates the overhead introduced by sending the ACK frames.

2. **Influence of backoff parameters on channel access delay and fairness:**

The results obtained with test models S2 and S4 show that the default backoff mechanism with I_{base} = D_{base} = 2 allows scenarios in which the channel access delay can reach very high values and a station can get locked out by other stations. Doubling the backoff window size after each collision actually reduces the risk of subsequent collisions, but the backoff window size and the related random backoff time can assume very large values making the channel access delay of the source devices enormous (see test model S2). In some cases, this disadvantage is experienced only by a subset of the source devices. In test model S4, a single station get locked out by the “explosion” of its backoff timer: its throughput is little and its frame drop high, if compared with the ones of the other source device. The severity of this lockout is common to all the protocol paths.
These problems can be avoided by freezing the backoff window size to its minimum value (I_{base} = D_{base} = 1) and by allowing more backoff retries (Max_Retries=9). The results show that the throughput increases significantly and that long channel access delays are very rare.

3. **Utility of backoff slot synchronization:**

The resynchronization of the backoff slots after each reservation is possible only at those non-participant devices that are within RR=16 range of the source device and/or of the destination device. At the end of the reservation, the source device, the destination device and all non-participant devices that are within the RR=16 range of source and/or destination device, go simultaneously from the IDLEWAIT state to the IDLE state and start the slot clock generator at the same time.
Test models S4 and S5 show that some undesirable effects can be caused by the asynchronism of the backoff slots (see sections 4.7.1 and 4.8.1). This is due to the fact that the size of the backoff slots and the whole contention mechanism have been designed with synchronism in mind.

4. **Utility of RTS putting to silence non-participant stations:**

Another question which arises is the following: is it really useful that a non-participant station, which hears an RTS with non-matching destination address, should remain silent during the subsequent reservation even if the station does not hear the related CTS transmitted by the destination device? This feature has been introduced in the AIr MAC specification because it can help dealing with the hidden terminal problem: a potential hidden interferer, which can hear the RTS from the source but cannot hear the CTS from the destination, has to remain silent during the ongoing reservation after receiving the RTS. However, test model S3 shows that there are cases in which this feature introduces an unnecessary throughput lowering. In test model S3, only the source devices can hear each other, so there is no interference at the destination devices even if the sources transmit data packets simultaneously. If one source device transmits an RTS (and there is no RTS-RTS collision), the other source device has to remain dumb during the subsequent reservation even if it could transmit its RTS and data packets without generating interference. This causes the throughput lowering.

A more general case is shown in Fig. 4.16. Two four-stations scenarios are considered. For each station, both the RR=16 and the interference range are depicted (see also Fig. 3.5).

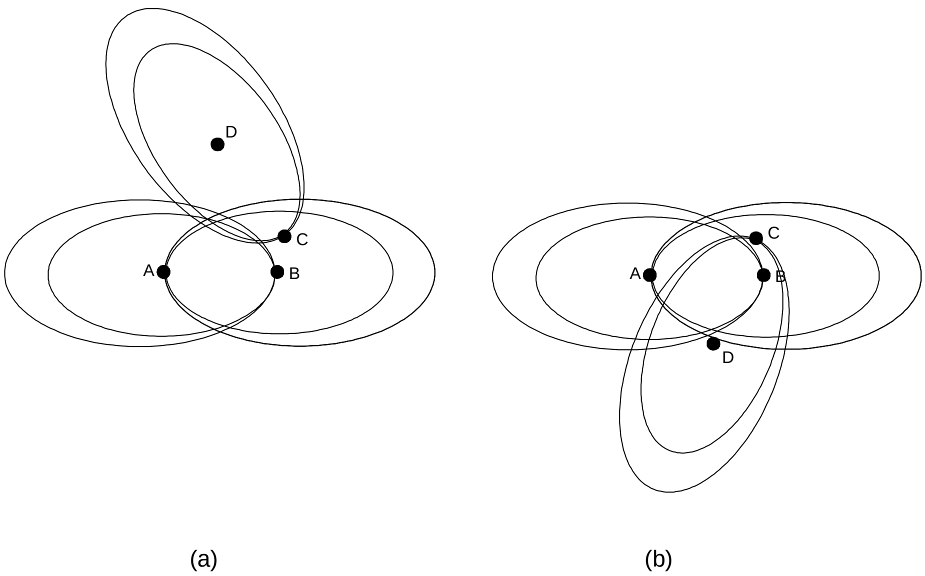

Figure 4.16: Utility of RTS

A is transmitting data to B, and C is transmitting data to D. In both scenarios (a) and (b), C is within the RTS range of A but is completely unaware of B and cannot receive any CTS frame from this station. Also, D is unaware of both A and B and cannot receive any AIr MAC frames from these stations. In scenario (a), B is outside the interference range of C: C could deliberately transmit to its destination D without interfering with B. However, the RTS frames coming from A obliges C to remain silent when a data transmission between A and B is going on. For C, this fact represents a useless waste of time and throughput lowering. In scenario (b), B is located within the interference range of C (and also within its RR=16 range). If C would not be able to receive any RTS from A, its transmissions to D would generate interference at B. Therefore, in this scenario the RTS form A to C are useful and necessary to avoid such an interference.

5. AIr MAC Performance: Interference Measurements

The goal of the interference measurements is to investigate how the AIr MAC performance (when the Reserved Data Transfer method is used) can be influenced by the following features:

- Transceiver parity
- RTS body format (repetition rate)
- Backoff slots synchronization mechanism

5.1 Test Model S6

The test model S6, that has been used for the first interference measurements, is illustrated in Fig. 5.1.

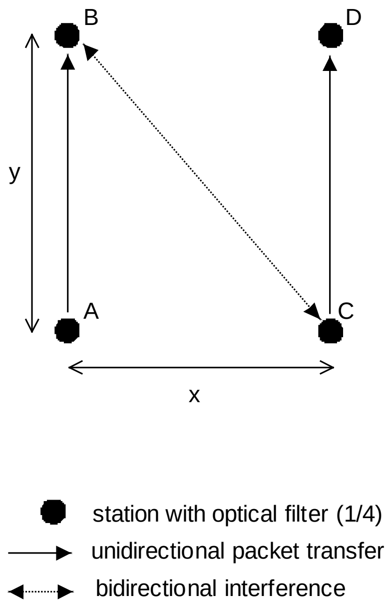

Figure 5.1: Test Model S6

Four stations are involved in this test model. An optical attenuation filter with factor 4.2 has been placed on each station's transceiver, so that the values of radiant intensity and of sensitivity were adjusted to accommodate the maximum interference ranges within the size of the test room (a 4.2-fold down-scaled range scenario). Y is the distance between stations A and B, and is always equal to the distance between stations C and D. X is the distance between stations A and C, and is always equal to the distance between stations B and D.
Station A is always unaware of stations C and D, no matter how the distance x is reduced, so the backoff slots of station A and C cannot be kept synchronized. Station D is always unaware of stations A and B, no matter how the distance x is reduced. During a measurement test, one unidirectional data transmission from station A to station B and one from station C to station D take place simultaneously. The stations

operate using the AIr reserved data transfer method, so that with ideal channel reciprocity between all stations no mutual interference should be present. However, in this scenario, station C is unaware of station A, so it is a potential hidden terminal and could act as an interferer, perturbing the data packets arriving at B from A. This is likely to occur if some of the CTSs emitted by B cannot be successfully received by C and C feels free to transmit data frames to D that overlap partly or totally with the data frames from A to B. So, the quality of the transmission on link A-B depends on the B-C link and, consequently, on the distance x. As x increases, the B-C distance also increases and we gradually move from a scenario in which B and C are within each other's RR=16 range to a scenario in which they are outside each other's interference range.
The following test parameters have been adopted:

Protocol path	RT/NACK
Number of Transmitted Frames	1000
Packet Payload Size	2048 Bytes
Packets Per Burst (PPB)	1
Backoff Retries (Max_Retries)	9
Minimum Backoff Window Size	8
BOWS Algorithm parameter I_{base}	1
BOWS Algorithm parameter D_{base}	1

- All timer values have been kept the same as the ones used for the measurements presented in section 4.2.
- RR values are fixed, i.e. no rate selection algorithm is being used. Data packets payloads are transmitted with fixed RR=1.
- The backoff slot synchronization is always active.
- The RTS body contains the destination and the source address (no CRC field is present), and is always sent at RR=16. In test model S8, some measurements have also been done with another RTS body format: in that case, the RTS body contains the destination address, the source address and a CRC field, and is always sent at RR=1.

The packet success rate of a link, C_p, is defined as follows:

$$C_p = 1 - FL$$

where FL is the frame loss of the considered link (see section 4.1 for the definition of the frame loss). The packet success rate of a particular link, say link A->B, will be denoted as C_p(A,B).
The distance y, between station A and station B, has been initially adjusted so that the packet success rate of the unperturbed link A-B was C_p(A,B) = 90% at 4 Mbps (RR=1). This value has been arbitrarily defined as the limit of useful link quality, corresponding to a bit error rate of approximately 10^{-5}. The obtained A-B distance was y=80 cm.
Then the packet success rate C_p(A,B) has been measured, at the end of an experimental test, as a function of the distance x. The distance x has been varied moving only the interfering link C-D as a whole and keeping the optical port axis of station C parallel to the port axis of station B. For each distance x, 1000 packets have been sent by station A to station B with the reserved data transfer method and the link C-D has been kept active, also in the reserved data transfer method, until the transmission on link A-B was over. The resulting packet success rate of link A-B was then recorded. Also, the number of CTS frames that reached station C from B during the test has been measured at each distance x.
The following behavior was expected:

1. If x is so small that C is still within the CTS range of B:
   C is able to successfully receive all the CTS frames emitted by B, thus avoiding any interference. C_p(A,B) should be 90%.

2. As x increases and C is gradually getting outside the CTS (RR=16) range of B:
   the number of CTS frames correctly received by C is expected to decrease gradually. Consequently, C_p(A,B) is expected to decrease due to the interference caused by C (B cannot capture all data packets in the presence of the interference frames from C). When x is so large that C gets outside the CTS range of B, then there is no medium access control anymore on link C-D. C emits its frames without caring if a reservation is proceeding on link A-B and if B is still in the interference range of C, C_p(A,B) will be smaller than 90%.

3. If x is further increased, C is already outside the CTS range of B and B is gradually getting outside the interference range of C:
   C deliberately interferes with B but the received signal power of the interference signal is gradually decreasing. As a consequence, C_p(A,B) is going back to 90%.

4. If x is so large that B is completely outside the interference range of C:
   the two links are fully independent and no more interference is present. So, C_p(A,B) is expected to be 90%.

Two measurement series have been done. For the first series, good parity was provided by the transceivers at both station B and station C. For the second series, an additional optical attenuation filter with factor 4.2 has been put on station C’s photodiode, so that a significant parity mismatch has been introduced at this station. The results are presented in the next two sections.

5.1.1 Approximate Parity

The results presented above assume approximate parity for both the interfering station (C) and the interfered station (B) transceivers. The used prototyped AIr transceivers have each 4 LEDs, driven with equal peak currents, and one photodiode. The positions and peak currents of the four LEDs at the transceiver used by station B (called T1) and at the transceiver used at station C (called T2) have been adjusted, until a good parity has been measured. Fig. 5.2 shows the radiant emission characteristics of transceivers T1 and T2 and the radiant reception characteristic of their photodiodes.

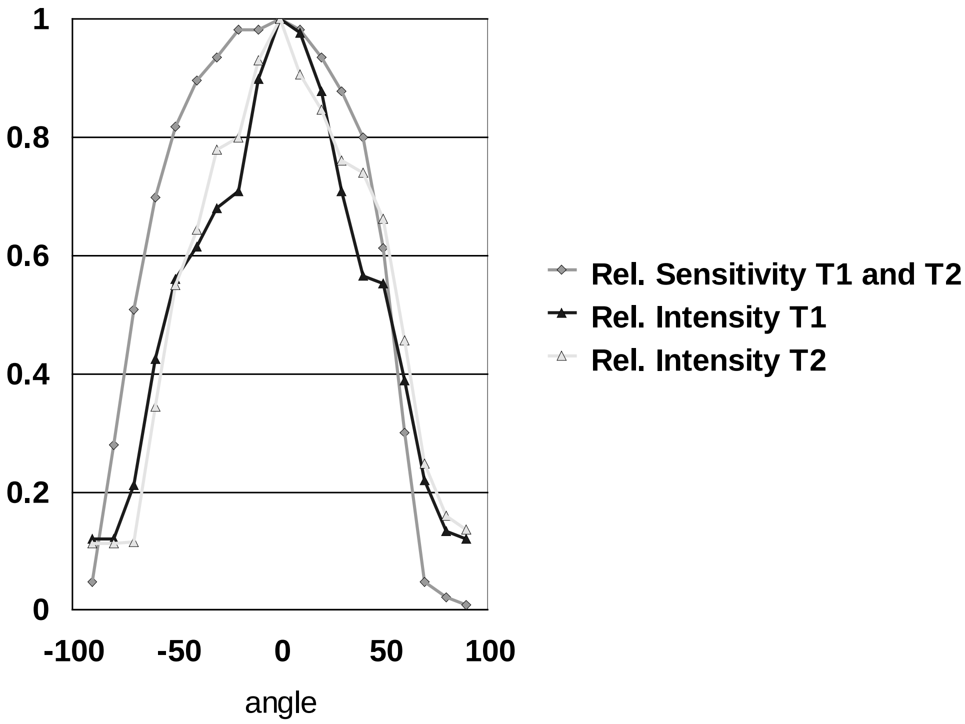

Figure 5.2: Characteristics of T1 and T2

There is no significant difference between the emission characteristics (relative intensity) of T1 and T2, and the curve of their relative sensitivity shows that they provide a good approximate parity.

Fig. 5.3 illustrates the measured curves. The blue curve shows the packet success rate $C_p(A,B)$ versus the distance x. The green curve shows the CTS success rate $C_{CTS}(C)$ (i.e., the number of CTS frames successfully received by C from B, divided by the total number of CTS frames emitted by B during the test) versus the distance x.

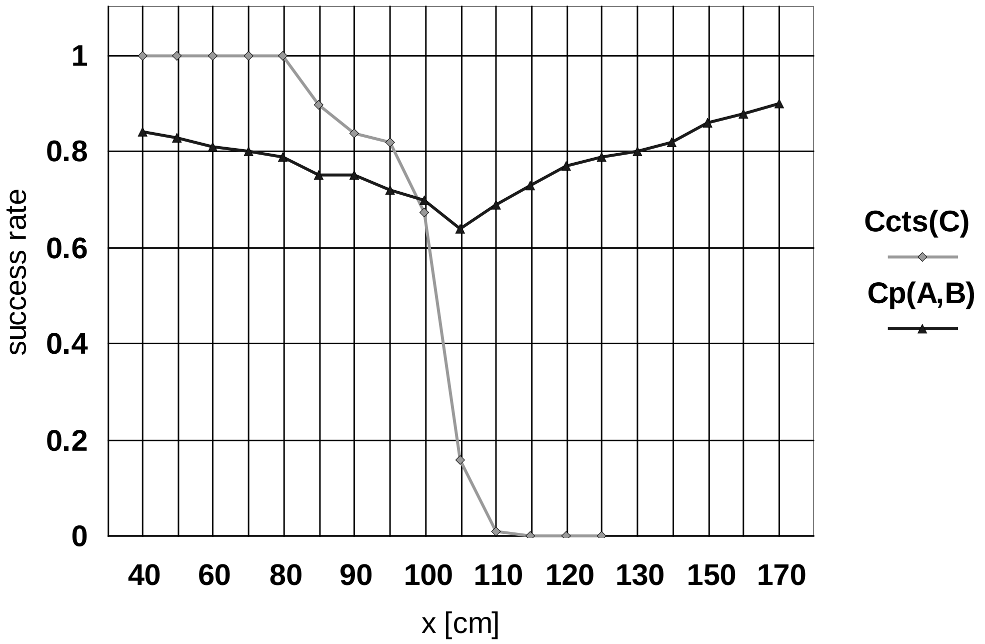

Figure 5.3: Test Model S6: curves for approximate parity at C

Three domains can be distinguished:

1. $x \leq 80cm$:
   C is located within the CTS range of B and can successfully receive all CTS frames emitted by B during the test. C_p(A,B) is expected to be 90%, but values between 79% and 84% have been measured. The cause of this packet success rate degradation has been investigated connecting the logic analyzer with the TX_ACTIVE signal (see section 3.4) of stations A and C. The following behavior has been observed: when these two stations launch an RTS frame in the same backoff slot, then two parallel simultaneous reservations (also called overlapping reservations) may exist. This is due to the following reason. Station D is unaware of link A-B, so it always responds with a CTS to all RTS frames from C. So, each RTS transmission from C is followed by the transmission of one data frame to D, which also interferes with B. If both A and C emit an RTS in the same slot, an RTS-RTS collision occurs at B. B should recognize this collision (the RTS from A cannot be successfully received) and should not respond with a CTS to A, avoiding an overlapping reservation. However, it has been observed that the recognition of the RTS collision does not always occur, due to the so-called *effect of capture*: if the ratio between the power received at B from A and the power received at B from C is greater than a given, physical layer dependent value, then B will be able to detect (to capture) the RTS from A even though it is colliding with the RTS from C. In this case, B responds with a CTS frame, an overlapping reservation occurs and, as the data frame is transmitted at RR=1 and is longer than the RTS frame (about 8.2 times longer), it is likely that B will not be able to successfully detect (capture) the entire data frame from A.
   The effect of capture gets stronger if the backoff slots of the transmitters are not synchronized (as occurs in this test model), because the overlap of the RTS frames during a collision is not total.

2. $80cm < x \leq 105cm$:
   C gradually gets outside the CTS range of B, as shown by the CTS success rate curve. B is still located within the interference range of C and, as expected, C_p(A,B) decreases as the distance x increases. This is due both to overlapping reservations and to the lack of the CTS frames at C from B (C is reaching the hidden terminal status). C_p(A,B) reaches the minimum value of 64% at x=105 cm.

3. $x > 105cm$:
   C is totally outside the CTS range of B. B gets gradually outside the interference range of C and, as expected, C_p(A,B) begin to increase until it reaches its maximum value 90% at x=170cm. At this distance, the links A-B and C-D are fully independent because B is located outside the interference range of C.

5.1.2 Parity Mismatch

An additional optical attenuation filter with factor 4.2 has been placed in front of station C's photodiode, to investigate the effect of this parity mismatch at C on the packet success rate C_p(A,B). The presence of the additional optical attenuation filter increases the required irradiance of C by a factor 4.2, so that a systematic system parity mismatch is introduced. The results intend to demonstrate that attention should be paid for achieving the smallest possible transceiver parity tolerances that economical constraints allow. Fig. 5.4 illustrates the measured curves. The blue curve shows the packet success rate C_p(A,B) versus the distance x. The green curve shows the CTS success rate C_{CTS}(C) at C versus the distance x.

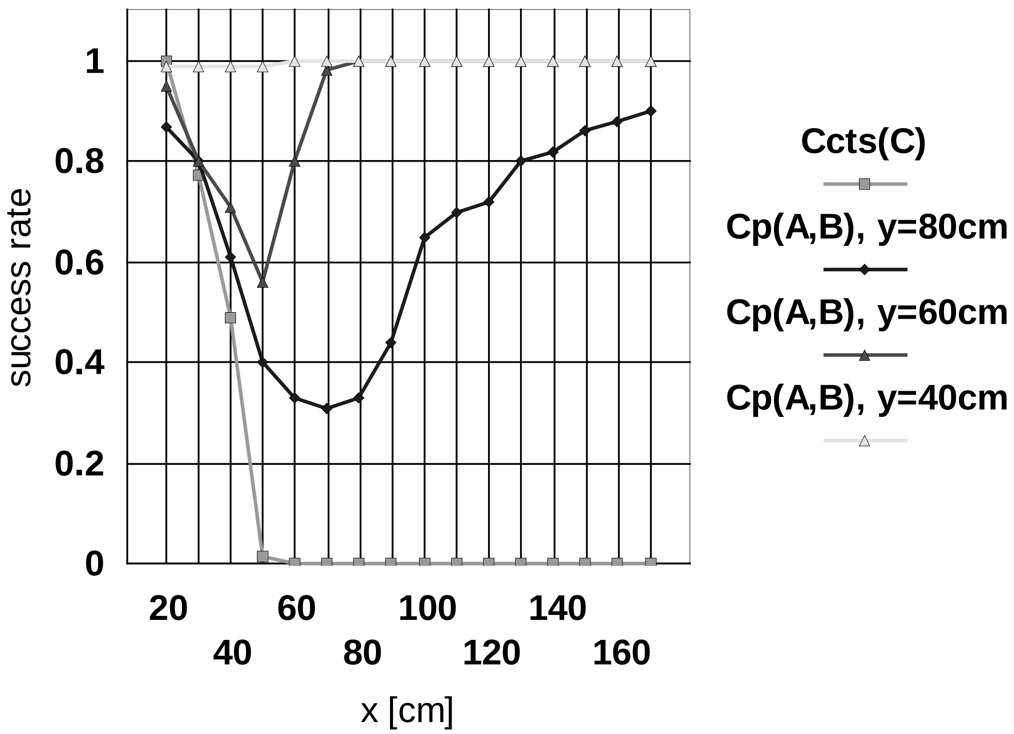

Figure 5.4: Test Model S6: curves for parity mismatch at C

It can be seen that the presence of the additional attenuation filter on station C's PD causes a reduction (2-fold down-scaling) of station B's CTS range, as it is perceived by C. With x=60cm, station C is already outside the CTS range of B. This fact, together with the overlapping reservations caused by the capture effect (see previous section), causes a dramatic lowering of C_p(A,B), which decreases faster and reaches its minimum of 31% at x=70cm. Then, as x is further increased and B gets outside the interference range of C, C_p(A,B) increases and can reach its maximum value 90%. For x > 130cm, the same C_p(A,B) values as in the scenario without additional attenuation filter at C have been measured (C is outside the interference range of B and the introduced parity mismatch has no effect).
The effect of the reduction of the distance y on the packet success rate has also been investigated. All measurements have been remade for y=60cm (about 75% of the initial A-B distance) and for y=40% (50% of the initial A-B distance). The result is also shown in Fig. 5.4. The red curve shows the packet success rate C_p(A,B) versus the distance x for y=60cm. The yellow curve shows the packet success rate C_p(A,B) versus the distance x for y=40cm. Even if the capture effect gets stronger, the packet success rate increases. For y=40cm, it remains about 100%. The reduction of the distance y increases the power received at B from A and the data frames coming at B from A can be well captured in the presence of interference from C. This leads to the idea to introduce *low-power RTSs*: reducing the optical peak power of the RTS frames transmitted by A obliges the source B to approach A in order to receive them correctly and respond with a CTS frame. The subsequent data frame is then transmitted with the full specified power. As a consequence, the subsequent packet success rate of the link A-B can be kept at sufficiently low values.

5.2 Test Model S7

The test model S7 is illustrated in Fig. 5.5.

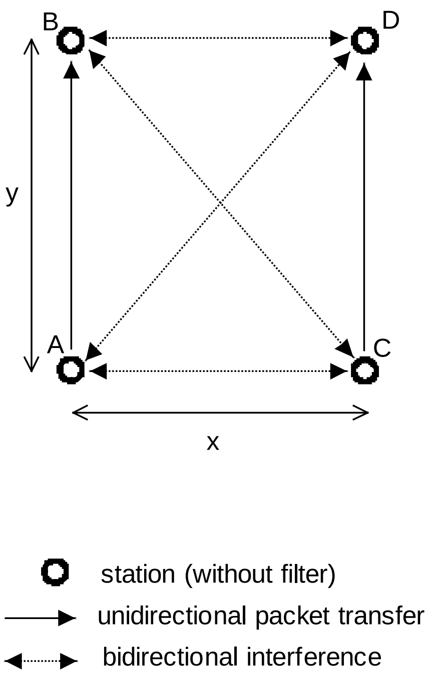

Figure 5.5: Test Model S7

This test model results from test model S6, after having removed the optical attenuation filters from all the transceivers. It follows that, depending on distance x, not only B and C can get within each other RR=1,

RR=16 or interference ranges: also A and C, B and D, and A and D can. In a scenario where A is within the RR=16 range of C and/or D, and C is within the RR=16 range of A and/or B, the backoff slots of the transmitters A and C can be kept synchronized.
The distance y (from A to B, the same as from C to D) has been set to 160 cm. At this distance, C_p(A,B) is 100%. As in scenario S6, the test has been executed at various x distances. Fig. 5.6 illustrates the measured curves. The blue curve shows the packet success rate C_p(A,B) versus the distance x. The green curve shows the CTS success rate at C versus the distance x.

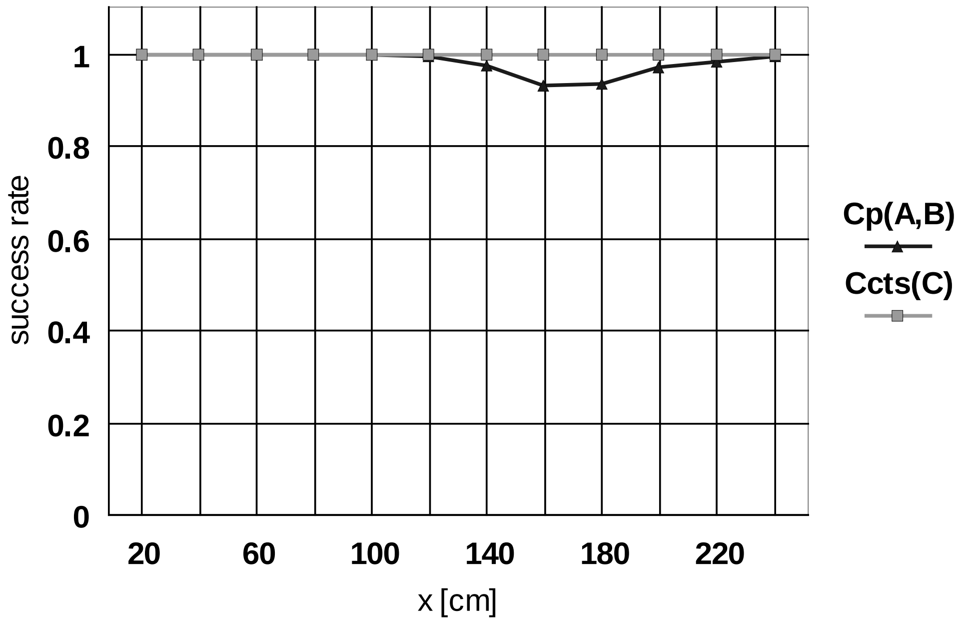

Figure 5.6: Test Model S7: curves for approximate parity at C

It can be seen that the CTS success rate is always 100%, i.e., C is within the CTS range of B for the distances x at which the measurements have been done ($x \leq 240cm$). As a consequence, it would be expected that C should not generate any interference on link A-B. However, it is not the case. The blue curve clearly shows that there is a degradation of C_p(A,B) for $130cm < x \leq 220cm$, with the minimum C_p(A,B) = 93% at x=160cm. This behavior is due to the capture effect (explained in the previous section) and to the consequent overlapping reservations. It has been observed, connecting the logic analyzer/scope with the TX_ACTIVE signal of stations A and C, that overlapping reservations occur for $x \geq 130cm$. After the transmission of a RTS frame in the same backoff slot by both A and C, both receivers B and D capture their RTS even in presence of the interference caused by the colliding RTS. B and D respond with an CTS and a fully overlapping transmission of the data frames from A to B and from C to D follows. These mutual interference causes the packet success degradation. As x get larger than 160cm, the strength of the interfering data frames at B and D get smaller and smaller, and it is more likely that these stations can capture also their data frames. The packet success rate C_p(A,B) begin to increase and reaches its maximum value 100% for $x > 220\,cm$.

5.3 Test Model S8

The test model S8 is illustrated in Fig. 5.7.

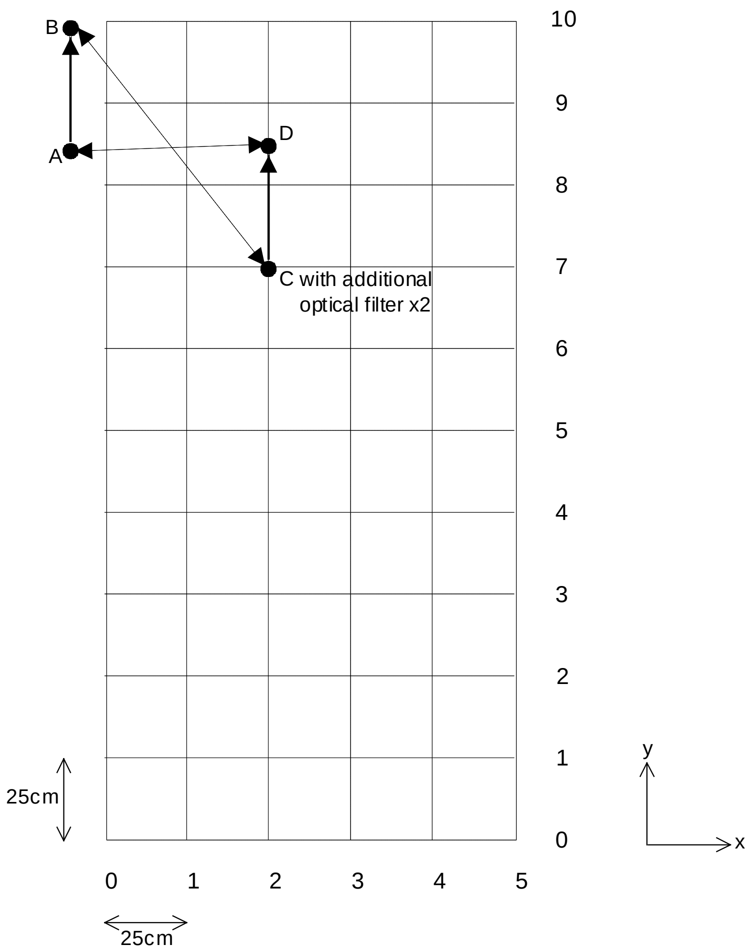

Figure 5.7: Test Model S8

This test model is similar to test model S6, but both links B-C and A-D can be present (depending on the position of the four stations) and the interfering link is now moved in a bidimensional 125cm by 275cm grid, with 25cm spacing. The optical attenuation filter with factor 4.2 has been kept on each station's transceiver. Additionally, a second optical attenuation filter with factor 1.7 has been placed on each transceiver, so that the values of radiant intensity and of sensitivity were further decreased. The resulting scenario is a 5.9-fold down-scaled range scenario. To introduce a parity mismatch at the interferer, an optical attenuation filter with factor 1.7 has been placed on C's photodiode. The same AIr MAC test parameters as the ones that had been adopted for test model S6 and S7 has been used in test model S8. Then, another measurement series has been done after modifying the RTS frame body structure, as explained below.

Like in test model S6, the distance between station A and station B, has been initially adjusted so that the packet success rate of the unperturbed link A-B was $C_p(A,B)$ = 90% at 4 Mbps (RR=1). The obtained A-B distance was 45cm. The location of A and B is near the upper, left-hand edge of the measurement grid. Stations A and B are motionless during the whole test: only the interfering station C will be moved about the grid, together with its destination station D. The optical port axis of C and D are kept parallel to the port axis of A and B for each measurement in the grid. Each time C is moved to a new grid location, D is placed at a distance for which the packet success rate of the unperturbed link C-D is $C_p(C,D)$ = 90% at 4 Mbps (RR=1). For each grid location, the following has to be done:

1. Place C in the grid location, with its port axis kept parallel and in reverse direction to the port axis of B.
2. Place D, with its port axis kept parallel and in reverse direction to the port axis of A, at the distance from C needed to obtain $C_p(C,D)$ = 90%. This distance is about the same as the A-B distance, i.e., 45cm.
3. Execute the test (1000 packets are sent from A to B; the packet transmission from C to D is stopped when the transmission from A to B is over).

At the end of the test, the following data are recorded, for each grid location:

- $C_p(A,B)$ in presence of interference from link C-D
- The number of CTS frames that reached station C from B, and the resulting CTS success rate value $C_{CTS}(C)$
- The number of RTS frames that reached station D from A, and the resulting RTS success rate value $C_{RTS}(D)$

To visualize the results, the data were processed to show interpolated contour lines with constant $C_p(A,B)$, $C_{CTS}(C)$ and $C_{RTS}(D)$.

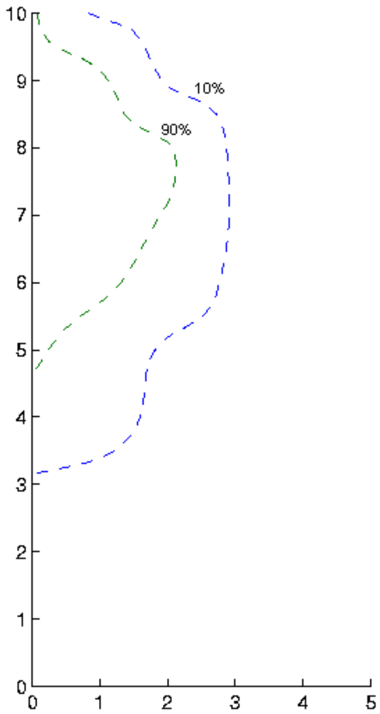

Figure 5.8: Test Model S8: Regions of $C_{CTS}(C)$

Fig. 5.8 shows the regions of CTS success rate at C, $C_{CTS}(C)$. The green line is the interpolated contour line at which $C_{CTS}(C)$ = 90%. The blue line is the interpolated contour line at which $C_{CTS}(C)$ = 10%. At the grid locations within the region delimited by the green line, $C_{CTS}(C) \geq 90\%$: the interfering station C can correctly receive more than 90% of the CTS frames transmitted by B, and is therefore expected to defer any pending transmission during the announced reservation period which could cause packet loss of link

A-B. The region between the blue line and the green line designates a transition zone with decreasing CTS reception probability: $C_{CTS}(C)$ goes from 90% to 10% and C gradually reaches the hidden status. In the region outside the blue line, $C_{CTS}(C)$ < 10% and C looses its CTS connectivity with B.

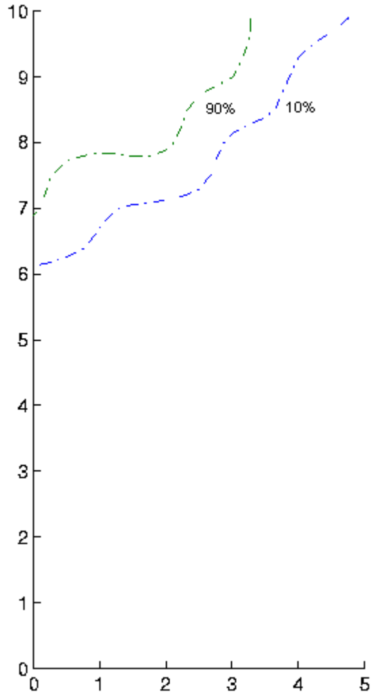

Figure 5.9: Test Model S8: Regions of $C_{RTS}(D)$

Fig. 5.9 shows the regions of RTS success rate at D, $C_{RTS}(D)$. The green line is the interpolated contour line at which $C_{RTS}(D)$ = 90%. The blue line is the interpolated contour line at which $C_{RTS}(D)$ = 10%. When C is placed at the grid locations within the region delimited by the green line, $C_{RTS}(D) \geq 90\%$: station D can correctly receive more than 90% of the RTS frames transmitted by A. With the RTS frames, A notify D that it is going to reserve the channel and that D should avoid responding with a CTS to a RTS from C during the reservation. This information is particularly useful in avoiding overlapping reservations when C has lost CTS connectivity with B. The region between the blue line and the green line designates a transition zone with decreasing RTS reception probability: $C_{RTS}(D)$ goes from 90% to 10% and D gradually get out the RR=16 range of A. In the region outside the green line, $C_{RTS}(D)$ < 10% and D looses its RTS connectivity with A.
Assuming that the link A-D is approximately symmetric (transceivers at A and D provide approximate parity), we can say that if D has RTS connectivity with A, then A has CTS connectivity with D. Consequently, if $C_{RTS}(D)$ = 100%, A is in the CTS range of D and can keep its backoff slots synchronized with the backoff slots of station C. The backoff slots of both A and C can be kept fully synchronized only if $C_{RTS}(D) = C_{CTS}(C)$ = 100%. This occurs when C is placed at the grid locations (0,7) to (0,10), (1,8), (1,9) and (2,8). If $C_{CTS}(C)$ = 100% but $C_{RTS}(D) \simeq 0\%$, C is able to remain synchronized with A but A cannot keep the synchronization with C. This occurs when C is placed at the grid locations (0,5), (0,6), (1,6) and (1,7). For all other grid locations, both A and C are unable to remain synchronized with each other.

The packet success rate $C_p(A,B)$ has been measured, for C at each grid location, using the usual RTS body format: the RTS body contains the destination address, the source address an is always sent at RR=16. Then, a second measurement series has been done with another RTS body format: the RTS body

now contains the destination and source addresses along with a CRC, and is always sent at the same repetition rate as the data, which in our tests is always RR=1. It is expected that, with the RTS body transmitted at RR=1, RTS-RTS collisions are more likely to be catastrophic and thus recognized by the sources at which the collisions take place. As a consequence, there should be less overlapping reservations due to the capture effect and the packet success ratio should increase.
The regions of packet success rate C_p(A,B) with the RTS body transmitted at RR=16 are shown in Fig. 5.10

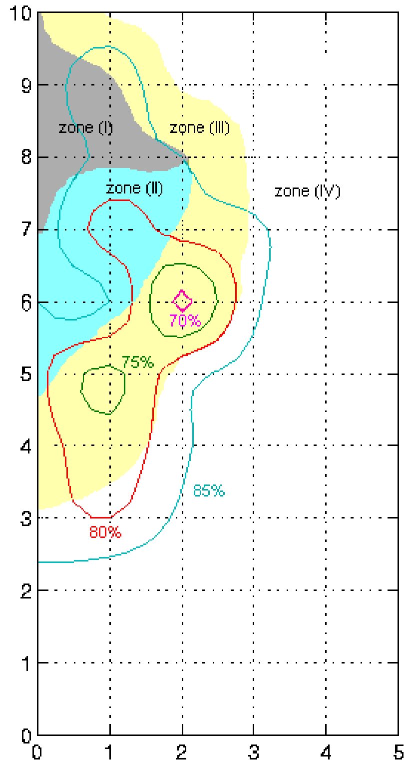

Figure 5.10: Test Model S8: Regions of C_p(A,B) with RTS body at RR=16

The blue line is the interpolated contour line at which C_p(A,B) = 85%; the red line is the interpolated contour line at which C_p(A,B) = 80%; the green line is the interpolated contour line at which C_p(A,B) = 75%. For C placed at the grid location (2,6), the packet success rate reaches its minimum C_p(A,B) = 70%.

The obtained contour plot can be divided in four zones:

1. Zone (I): C_{RTS}(D) = C_{CTS}(C) = 100%
   In this zone, the backoff slots of A and B are fully synchronized. C can receive all CTS frames from B and it is expected that C_p(A,B) = 90%. However, this happens only when C is placed at the grid locations (0,7) to (0,10), for which no overlapping reservations have been observed. When C is placed at (1,8), (1,9) or (2,8), some overlapping reservations occur and C_p(A,B) decreases, ranging from 82% at (1,9) to 85% at (2,8). These overlapping reservations are due to the capture effect at B and D. The signal power at B (D) from A (C) is high enough to allow B (D) to capture the RTS from A (C) even if it collides (total overlap) with an RTS from C (A). A and C then transmit one data frame simultaneously, interfering with each other. The data frames are transmitted with RR=1 and are so long (about 8.2 times longer than an RTS frame), that it is very unlikely that they can be entirely captured by the sources B and D in the presence of interference.

2. Zone (II): $C_{CTS}(C)$ = 100% and $C_{RTS}(D) \simeq 0\%$
   In this zone, the backoff slots of A and B cannot be kept synchronized, because A has lost its CTS connectivity with D. C can receive all CTS frames from B and it is expected that $C_p(A,B)$ = 90%. However, some overlapping reservations occur and $C_p(A,B)$ decreases, ranging from 78% at (1,7) to 85% at (0,6) and (1,6). These overlapping reservations are due to the capture effect at B. Station D is unaware of link A-B, so it always responds with a CTS to all RTS frames from C. Each RTS transmission from C is followed by the transmission of one data frame to D, which also interferes with B. If an RTS-RTS collision occurs at B. the signal power at B from A is high enough to allow B to capture the RTS from A even though it is colliding with the RTS from C. In this case, B responds with a CTS frame, an overlapping reservation occurs and, as the data frame is transmitted at RR=1 and is longer than the RTS frame, it is likely that B will not be able to successfully detect (capture) the entire data frame from A.
   The effect of capture is emphasized by the asynchronism of the backoff slots of the transmitters, because the overlap of the RTS frames during a collision is not total.

3. Zone (III): $C_{CTS}(C)$ decreases from 100% to 0%
   This is a transition zone with decreasing CTS reception probability. $C_p(A,B)$ is expected to decrease significantly due to the failure of B to capture all data frames in the presence of interference frames from C. This interference is due both to the persisting overlapping reservations (B captures an RTS from A even if it is colliding with an RTS from C) and to the lack of CTS frames from B at C (C is gradually reaching the hidden terminal status).
   In these zone the lowest $C_p(A,B)$ values have been measured. $C_p(A,B)$ reaches the minimum values of 70% at (2,6) and of 74% at (1,5).

4. Zone (IV): $C_{CTS}(C)$ = 0%
   C has no CTS connectivity with B. The interference from C continues, but B gets gradually outside the interference range of C and $C_p(A,B)$ begin to increase as the distance from the link A-B increases, until it reaches its maximum value 90%

$C_p(A,B)$ has then been measured with the RTS body transmitted at RR=1 The regions of packet success rate are shown in Fig. 5.11.

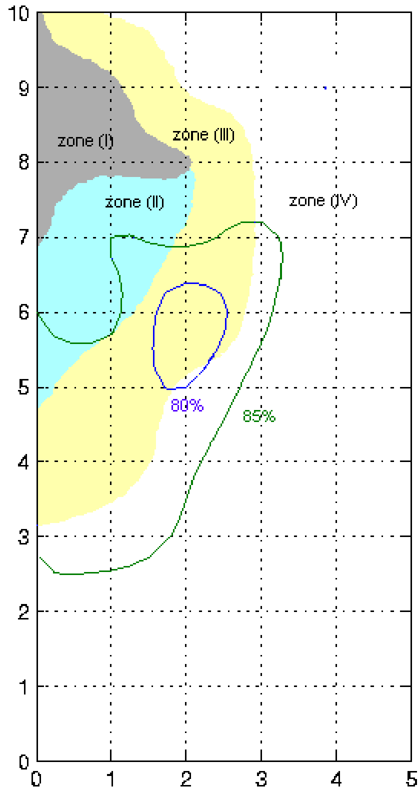

Figure 5.11: Test Model S8: Regions of C_p(A,B) with RTS body at RR=1

C_p(A,B) is now larger than 80% for all locations in the bidimensional grid, except at (2,6) where the new minimum 77% has been measured. The range of the region with C_p(A,B) < 85% (now within the green curve) is significantly reduced, because C_p(A,B) increased to 90% when C was at (1,8), (1,9) or (2,8) and thus C_p(A,B) has the maximum value 90% at all locations within zone (I). However, only a little improvement of C_p(A,B) has been measured in zones (II) and (IV).
The general amelioration of the packet success rate can be explained as follows.
The rate reduction of the RTS body, from RR=16 to RR=1, makes the RTS frame body less robust and more sensible to collisions. When A and C transmit an RTS simultaneously (in the same backoff slot), a RTS-RTS collision occurs at B. If the RTS bodies have been transmitted without repetition-encoding them, it is more likely that B will detect an error in the received RTS body (i.e., a CRC error, or simply an unmatching station address) and will not be able to capture its RTS, recognizing that a collision has occurred. Then, B will not respond with a CTS to A and no overlapping reservation takes place. So, the probability of a catastrophic RTS-RTS collision at B increases if the RTS body is transmitted at RR=1, and the result is a reduction in the number of overlapping reservations.
It has been observed, with the logic analyzer/scope, that when C is at the grid locations (1,8), (1,9) or (2,8), the number of RTS-RTS collisions recognized by B during the test with the RTS body sent at RR=1 is twice as much the number with the RTS body sent at RR=16. This effect declines as the distance between B and C increases, because the signal power received at B from C decreases allowing B to capture its RTS frame.

5.4 Conclusions

The measurements results, collected with test models S6, S7 and S8, have highlighted how some physical layer features have an impact on the AIr MAC performance. The following issues have been investigated:

- Capture of RTS frames during RTS-RTS collisions:

  The RTS capture during RTS-RTS collisions and the consequent interference due to overlapping reservations have been observed in test models S6, S7 and S8. If the signal-to-interference ratio SIR at an AIr station is higher than a given, AIr IrPHY-dependent value $SIR_{100\%capture}$, then this station is always able to capture its RTS frames even if they are colliding with an RTS (or another frame) from another station. If $SIR_{0\%capture} < SIR < SIR_{100\%capture}$, only a percentage of the RTS frames can be captured in the presence of interference. If $SIR < SIR_{0\%capture}$, the station is no more able to capture its RTSs in the presence of interference and the RTS-RTS collision are always catastrophic. The RTS capture during an RTS-RTS collision can cause scenarios where overlapping reservations occur and some stations mutually interfere. Fig. 5.12 shows a simple example.

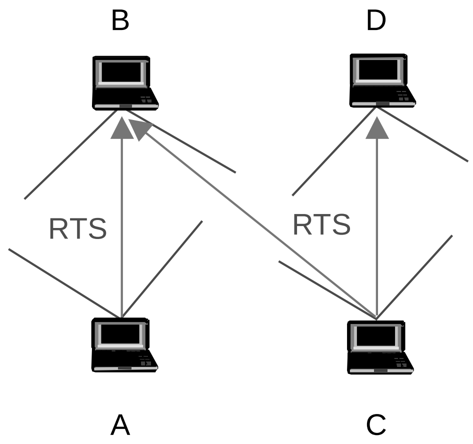

Figure 5.12: RTS-RTS collision

The red lines show the viewing angles of the stations. In this scenario station A sends an RTS to B and, at the same time (in the same backoff slot), C sends an RTS to D. Station D is outside the communication range of A, so no RTS-RTS collision occurs at D. After receiving its RTS, D responds with a CTS. On the other side, B is within the RTS range of C, so an RTS-RTS collision occurs at B. If B is able to capture its RTS even if it is colliding with the RTS from C, the RTS-RTS collision is not catastrophic and B responds with a CTS to A. Both A and C begin the data transmission simultaneously, after having received their CTS frames: this is a so-called overlapping reservation. The data frames are longer than the RTS frames (more than 8 times longer in the configuration of scenarios S6,S7 and S8). Consequently, the likelihood that station B will be able to capture its data frame in the presence of the interfering data frame from C is lower than the likelihood of RTS capture. The overlapping reservation will often result in a frame loss at B.
A means to mitigate the RTS capture effect is to reduce the repetition rate (RR) of the RTS body. This has been shown with the measurements of test model S8. If the RTS body is sent with RR=1 instead of sending it with RR=16, it will be less robust and more sensible to collisions. In other words, decreasing the repetition rate of the RTS body increases both $SIR_{100\%capture}$ and $SIR_{0\%capture}$. As a consequence, there will be more catastrophic collisions, less overlapping reservations and thus less mutual interference between the stations.

- Synchronization of backoff slots:

  Test model S8 has shown that the RTS capture effect is emphasized when no backoff slot synchronization can be provided. When two transmitters decide to emit an RTS in the same slot and their backoff slots are not in synch, then the transmission do not begin at the same time and the RTS-RTS collision do not occur with a full overlap. As a consequence, the interference will not last during the whole reception and the capture is facilitated.

- Transceiver parity:

  Test model S6 has clearly shown that transceiver disparity at the interfering station significantly emphasizes the hidden terminal problem and the consequently generated interference. It is therefore recommended that all AIr stations provide at least approximate transceiver parity.

6. Conclusions

After an overview of the IrDA standards and of the AIr system, this report has presented a series of preliminary measurements (test models S1 to S5) and a series of interference measurements (test models S6 to S8).
The goal of the preliminary measurements was to investigate how some protocol options (reservation termination path, data acknowledgement, backoff algorithm) can have an influence on the AIr MAC performance, evaluated in terms of system throughput and frame drop/loss at the AIr stations involved in the tests. The main results obtained with the preliminary measurements are resumed in the following table:

Area of Investigation	**Test Models**	**Related Problems**	**Results/ Recommendations**
Utility of EOB/EOBC handshake	S1, S2	EOB/EOBC introduces overhead	EOB/EOBC should be used. The overhead is small and is compensated by the offered advantages (early termination, RR* feedback, BO slots synchronization)
Utility of data acknowledgements	S1, S2	ACK-DATA introduces overhead	ACK-DATA should be used only when necessary. The overhead is large and is not always compensated by the offered advantages (early termination, RR* feedback)
Channel access fairness of backoff algorithm	S2, S4	Default BOWS adjustment algorithm ($I_{base} = D_{base =}$ 2) introduces channel access latency and station lockout	Backoff algorithm with $I_{base} = D_{base =}$ 1 (constant BOWS) should be used, with high max. BO retries number (e.g., 9)
Utility of backoff slots synchronization	S4, S5	Unexpected protocol paths caused by BO slots asynchronism	The synchronization mechanism should be used. However, it cannot work properly in some scenarios
Utility of RTS putting to silence non-participants	S3	Useful to deal with the hidden terminal problem, but can introduce unnecessary throughput lowering	CTS should always put to silence non-participants. RTS in some scenarios only

The goal of the interference measurements was to investigate how a transceiver parity mismatch at a potential interfering device can cause a packet success rate degradation, and have been inspired by the study presented in [8]. During the measurements, other problems (the effect of RTS capture during RTS-RTS collisions, the effect of backoff slot asynchronism) have been individuated and investigated. The main results obtained with the interference measurements are resumed in the following table:

Area of Investigation	Test Models	Related Problems	Results/ Recommendations
Capture of RTS frames during RTS-RTS collisions	S6, S7, S8	Causes overlapping reservation and consequent frame losses	Mitigate the RTS capture transmitting RTS body at RR=1 or utilizing low-power RTS
Physical layer disparity	S6, S8	Introduces interference due to the hidden terminal problem	Transceiver parity should be provided by all AIr devices
Utility of backoff slots synchronization	S8	Emphasizes the RTS capture effect	The problem cannot be avoided in some scenarios

More detailed explanations of the issues resumed in this two tables can be found in chapters 4 and 5.

This diploma thesis work has gone through several phases.
At the beginning, four prototype AIr stations have been set up to be used for the experimental performance measurements. The installation and configuration of the prototype AIr hardware and software components were particularly time-consuming. The available documentation was poor and a frequent information feedback has been demanded to member of the Infrared Wireless Communications Group at IMD Toronto. Moreover, several technical problems (incompatibility between some IBM ThinkPads and some docking stations, hardware conflicts) have slowed down this initial phase of the work.
During a second phase, some useful signals have been found at the available prototype interfaces. Connecting these interfaces with a logic analyzer/scope during the execution of the measurement tests has allowed to get some more insight in the operation of the devices and to understand better the measurement results. A logical signal that goes to the high voltage state when the station is transmitting and a five-bit AIr MAC FSM state vector, have been particularly useful.
During a third phase all the experimental performance evaluation measurements have been done. Firstly, the investigation areas have been defined. Secondly, some sensible test models have been designed. Thirdly, the actual measurements have been done and the results have been analyzed.
During a last phase the final report has been written. It summarizes the results of the diploma thesis work and tries to provide some recommendations for optimum AIr MAC usage.

Appendix A. Brampton Board (AIr Prototype) Tutorial

The following files are needed:

- For the FPGA update: **eeprom.exe, upload.exe, bramp50.102**
- For the diagnostic and test vehicle: **test7.cfg, app.cfg, menu.exe, menu.map**
- For the events tracer: **evtrace.exe, egavga.bgi**

ThinkPad and Docking Station Configuration

As configured in app.cfg, the Brampton board ISA interface uses IRQ 10 and DMA channels 0 and 3 (one for receiving and one for transmitting). They have to be reserved using the ThinkPad Features Setup and all existing hardware conflicts have to be eliminated.
Also, not all ThinkPad-Dock combinations do work:

	Dock1	**Dock2**	**Dock3**
ThinkPad760ED	YES	YES	YES
ThinkPad750C	NO	YES	NO
ThinkPad750Cs	YES	YES	NO

1. Brampton Board Hardware Configuration

First of all the FPGA has to be updated using the files eeprom.exe, upload.exe and bramp50.102.
Some other hardware configuration parameters can then be set in the configuration file app.cfg and running menu.exe.
Remember that you can run all files only under pure DOS.
Also, never hot-plug the transceiver and always use the ISA connector.

Brampton FPGA Update

In order to update the Brampton FPGA you have to go through the following points:

1. Power off ThinkPad and docking station
2. Ensure the transceiver is not plugged in
3. Set the 4-switch DIPs to programming operation, as below:

	1	**2**	**3**	**4**
SW1	ON	ON	OFF	OFF
SW2	OFF	OFF	OFF	OFF
SW3	OFF	OFF	OFF	OFF

4. Power up Think Pad and docking station
5. Run UPLOAD BRAMP50.102
6. When complete, power off ThinkPad and docking station

7. Reset the 4-switch DIPs to normal operation, as below:

	1	2	3	4
SW1	ON	ON	OFF	OFF
SW2	ON	OFF	OFF	OFF
SW3	OFF	OFF	ON	ON

Brampton Transceiver Connectors

There are two transceiver connectors on the board, one connector to the 15-pin ZRL transceiver, and the other through the standard 8-pin single ended IR Xcvr dongle interface. Two jumper changes are required to switch between which connector is used:

- **ZRL Configuration**
  Jumper labeled “DataIn”: place jumper between center pin and pin marked ‘I’
  Jumper labeled “IR RX”: open all pins

- **Ladybug Configuration**
  Jumper labeled “DataIn”: place jumper between center pin and pin marked ‘P’
  Jumper labeled “IR RX”: place jumper between center pin and pin marked ‘I’

The hardware settings in the configuration file app.cfg must always reflect the jumper settings. There is a switch in the app.cfg, called “ZUR=”. If you are using a ZRL transceiver you have to set “ZUR=1”, uncomment the step table values for ZRL and comment out the step table values for IBM Firefly. There is no step table available yet for Ladybug (if you are using a Ladybug you have to set the option within the menu.exe program itself, hitting two times ‘F’ from the main menu as explained below).

2. Brampton Board Driver (Test Vehicle)

The executable file menu.exe drives the Brampton Board through diagnostic tests and measurement studies.
It uses two configuration files: test7.cfg and app.cfg.
After running menu.exe you get the list of test configuration parameters that will be used later for the measurement study (using the Mac: Basic Reservation Menu). So you can hit ‘n’ to go to the main menu.
You can always use ‘n’ or ‘p’ to go back to the previous menu.

2.1 The Main Menu

In the main menu the values displayed at the top (Diagnostics Settings, white) reflect the hardware and timer values set in the app.cfg file. All values in app.cfg are in HEX:

1. TIMER VALUES Section: the five backoff timer values (BCK, corresponding to the first try plus four retries, all in 900 usec’s) will be used if you choose to adopt a fixed backoff policy. The other timer values represent the mean turn delay (MIN), the response timer (RSP), the WFDATAP timer (WDP) and the WFDATANP timer (WNP) (all in 4 usec’s). You also find the reservation timer (RSV, in 256 usec’s), the 2T2 timer (2T2, in 1 msec’s) and the maxburst timer (MBT, in 256 usec’s).

2. FLAGS section: These control the printing of message information (enable, disable, short/full), <u>rate select with RR* (fixed or adaptive)</u>, transceiver type (ZRL/Firefly/Ladybug) and the passing up to the test layer of CRC errors.

Below, you find several options:

0 Check Chip Type
1 Register Test Menu
2 VR Sniffer Menu
3 Packet Banger Menu
4 Send Frame Menu
5 Simple Backoff
6 Mac: Basic Reservation Menu
7 SIR Menu
8 DOS Demo
9 EV Trace
E To Exit Diagnostic
F Set Hardware Configuration

Selecting option F (Set Hardware Configuration), then F again (Misc. Options), you can set the FLAGS section values.

2.2 Measurement Study with Mac: Basic Reservation Menu

Selecting option 6 in the main menu (Mac: Basic Reservation Menu) it is possible to do measurement studies with the reservation method.
Two options are then displayed:

0 Set Tx Packet Parameters
1 Transmit RDATA in Auto Mode

Hitting 0 you can set all the diagnostics settings that will be an effect on the Mac: Basic Reservation tests. The diagnostics parameters will be presented in section 2.3. Here only the subset of diagnostics settings that will have an effect on the Mac: Basic Reservation test is presented:

- CHAIN: chaining frames
- TERM: use of EOB or of RT termination
- RT: calculate the reservation time or take it from test7.cfg
- OSR: slot reset on OLR
- RTS: Auto : automatic handling of RTS ABORT and RTS2 ABORT
- SLR: slot resync
- BOALG: the used backoff algorithm
- RSALG: the rate select method

Hitting 1 you can begin the test. The following table contains the test configuration parameters that can be set before launching a test and also appears when first executing menu.exe. All values, in decimal, are stored in the configuration file test7.cfg. They override the diagnostics values in case of redundancy.

PARAMETER NAME	MIN/MAX VALUE	EXPLANATION
RxNumberOfPkts	0/32767	Number of packets this station expects to receive (recommended: 20006)
TxNumberOfPkts	0/32767	Number of packets to be transmitted (also called NTX) (recommended: 20006)
NumOfPktsPerRsv	0/7	Number of packets sent per reservation (also called PPB)
AveragePDUSize	0/2048	Average packet (frame) size (also called PS)
PDUSizeFluct	0/1	0: the packet size fluctuates around the average size 1: the packet size is always equal to the average size
ShowStatus	0/10	This is NOT USED. Setting it to 3 will show backoff calculation values
DataAck	0/1	0: ACKs are not used (DATA frames) 1: ACKs are used (ADATA frames)
TxTrigger	0/1	0: this station is a slave 1: this station is the master which triggers the test beginning
RRValue	1/16	Start value of RR for repetition coding the data
RateSelectProc	0/1	is NOT USED
ReservationTime	1/4095	Reservation time (RT) in 3.6 msec's This value is not taken into account if RT is computed for each reservation (see RT: calculate in the packet settings)
MinTurnDelay	0/255	Value of MTD in 4 usec's (overrides the value contained in app.cfg)
AdjustOtherTimers	0/1	1: all timer values (RESPONSE, WFDATAP, WFDATANP) are adjusted according to MTD (they are not read from app.cfg)
DeltaDelay	0/65535	Time delay to use before the packet is sent out, in 1 usec's (overrides the DELAY value in the packet settings)
SrcStationNumber	0/10	Address of this station
DstStationNumber	0/10	Address of the destination station
BORetries	0/9	Maximum number of retries, when attempting to make a reservation, before giving up and discarding the transmit packets (recommended: 4)
BOAlgoIbase	1/128	Exponential base for increasing the backoff window size (recommended: 2)
BOAlgoDbase	1/128	Exponential base for decreasing the backoff window size (recommended: 2)
BOAlgoDslope	0/1	Linear slope for increasing the backoff window size (recommended: 0)
BOAlgoDdelta	0/128	Linear slope for increasing the backoff window size (recommended: 0)
PerfUpdateRes	1/32767	Number of packets to receive or transmit before updating test outcomes

2.3 Diagnostic Tests with Send Frame and VR Sniffer

With the main menu options 2 (VR Sniffer) and 4 (Send Frame) you can do some diagnostic tests exchanging single packets or packet bursts between stations. After selecting option 2 or 4, a common display (Packet Menu) appears at the top with the following parameters:

TXMODE: the transmission mode (AIr or IrDA)
SPEED: the speed, depending on TXMODE
MOD: the modulation (AIr does not use this)
DELAY: the delay in us before a packet is sent out
RTS: the RTS frame format; default is DA, SA, no CRC
RHCRC: Turn on/off the RH CRC generator
OSR: slot reset on OLR
CRC: Turn on/off the CRC generator
TERM: EOB or RT termination
RTS: automatic handling of RTS ABORT and RTS2 ABORTs
CHAIN: chaining AIr frames (note that you cannot chain UDATA frames)
RT: calculate the RT or use the RT set in app.cfg (or in test7.cfg if using the
    MAC: Basic Reservation Menu)
RTPAD: add RT padding
SLR: slot resync
RID: RID for Air
RT: reservation time
FRTYPE: type of frame to send (DATA, ADATA, UDATA)
RETRY COUNT: number of RTS retries
L: PPM rate
RR: Repetition rate
DA/SA: source and destination address
SIZE: size of data portion to send
BOALG: the used backoff algorithm (fixed, random, Olsen, Murray)
RSALG: the used rate select method (table lookup or Rupam adaptive)
DATA: data to send (ASCII only)

Auto Mode (AIr) supports only DATA, ADATA and UDATA. When sending DATA or ADATA in Auto Mode, the peer station has to exchange RTS/CLS packets according to the AIr MAC protocol.
Manual Mode can send all frame types and does not imply any RTS/CTS exchange.

Under the Send Frame Menu you can send single or multiple frames in manual or auto mode. In loopback mode a station receives the same frames it sends out.
Under the VR Sniffer Menu you can receive the frames sent by another station under Send Frame Menu. In manual mode the sniffer displays the frames it sees but does not respond to AIr frames. In auto mode the sniffer responds to AIr frames directed to it from other stations (if SA=DA) and shows DATA and EOBC frames. Using the “Detect DCD Test” allows to monitor the carrier sense (CS) flag. The information displayed varies depending on the options FLAGS set. A summary of the sniff is displayed at the end.

2.4 Tracing States and Events with EVTRACE

With the main menu option 9 EV Trace you can trace events and state transitions while running a test with Mac: Basic Reservation. EVTRACE software is modeled after a transitional based logic analyzer. The input vector is 3 bit state and 5 bit event from the AIr MAC FSM. When a change occurs on the input vector the hardware logs the state-event couple (8 bit) and the time (24 bit, 125 ns resolution) from the last state-event couple. You can log up to 64 state-events couples (a limitation of the 4020’s size) which is roughly 1 and ½ reservations with 7 packets per burst.

The numbers of states and of events are listed in the AIr MAC VHDL document. The states are: 0-IDLE, 1-RECEIVE, 2-CONTEND, 3-WFCTS, 4-TRANSMIT, 5-IDLEWAIT, 6-WFDATA (and 7-UDATAwait).

Hitting 9 you get the EVTRACE menu and hitting ‘↑’ or ‘↓’ you can go through the parameters. All parameters, except the first (1-Status), can be set with ‘→’ and ‘←’.
1-Status indicates if the tracer is running or stopped. You can start/stop hitting ‘F1’. In running mode the analyzer is waiting for a trigger or filling post-trigger memory after a trigger occurred. When stopped, the analyzer is not running, data can be viewed and new parameters entered.

The following parameters are used for setup, while the analyzer is stopped:

- Trigger point(s) (<: Find ‘A’ OR ‘B’ OR ‘C’>), or sequential mode (<1: Find ‘A’ 2: Trigger on ‘B’ 3: Elsif ‘C’ goto 1:>). With the trigger point(s) you can set up to three triggers (A OR B OR C). The sequential mode allows you to wait for trigger event A, then stop the analyzer on B, else reset the sequence if C is found, at which time the analyzer waits for A to occur as before. To set the trigger(s) A= Event[] State[], B= Event[] State[] and C= Event[] State[], press ENTER and enter a decimal number or ‘xx’ for don’t care. Each trigger can be STATE and EVENT or just STATE or EVENT only.

- Post Trigger %: set the maximum amount of memory to be used after the trigger has occurred. Example: if set to 0%, when the trigger occurs, the analyzer stops and up to 64 events up to the trigger are visible, but no events after the trigger will be captured; if set to 50%, up to 32 pre-trigger events and up to 32 post-trigger events are viewable.

Once these parameters are entered they are not stored in the analyzer parameter memory until F1 is pressed which also starts the analyzer. Once entered they remain until changed or until power is removed from Brampton (software resets have no effect). Pressing F1 again stops the analyzer whether or not its memory is full (note that EVTRACE will stop the analyzer automatically if a trigger is acquired and its memory is filled).
Once the trigger and post trigger % are set, then press F1 before exiting to save the settings and start the analyzer (press two times F1 to stop the analyzer if you do not want it running this time). You can then exit EVTRACE pressing ‘esc’ and run any diagnostic program. During a test, you can start and stop the analyzer by pressing F1 and F12 respectively. So once you are ready press F1 during a test and then at the end of the test enter EVTRACE to see if a trigger occurred.
When the analyzer is stopped EVTRACE downloads the memory contents and formats it for waveform display. The waveform display settings are listed below:

- Delay: the center point of the screen is the time from the trigger point. You can look before or after the trigger
- Time/Div: basically a zoom in and out in time
- Trig to X: a colored marker to measure the time between the marker and the trigger point
- Trig to 0: a second colored marker. X to O shows the time between the two markers

The waveform display shows the value the states and events. Vertical lines depict the boundaries on individual events and the end of acquired data. A concentration of more than one vertical line makes it turn blue, which means you need to zoom in to see the hidden events. The lack of a vertical line on the left or right side of the screen means that the data are off the edge of the screen.
The trigger line is red bar and the center of the screen is blue bar. The X marker is green and the O marker is yellow.

References

[1] IrDA official web site, http://www.irda.org/.

[2] I. Millar, M. Beale, B.J. Donoghue, K.W. Lindstrom, S. Williams, “The IrDA Standards for High-Speed Infrared Communications”, *The Hewlett-Packard Journal*, February 1998.

[3] J. Taglione, J. Molson, “IR provides high speed wireless connectivity for portable devices”, *Wireless Systems Design*, May 1997.

[4] I. Millar, M. Beale, S. Williams, F. Gfeller, P. Kermani, M. Naghshineh, “Technical Overview of an Advanced IrDA Link Access Protocol and Physical Layer Draft Version 0.1”, April 1997.

[5] F. Gfeller, W. Hirt, “Advanced Infrared (AIr) IrPHY Draft Physical Layer Specification. Version 0.4”, January 1998.

[6] C. M. Olsen, D. McKay, “Super FSM for Advanced Evaluation of AIr IrMAC Protocol Options and Performance Version 0.5”, November 1997.

[7] I. Millar, M. Naghshineh, M. Olsen, “Advanced Infrared (AIr) IrMAC Draft Protocol Specification Version 0.3), October 1997.

[8] F. Gfeller, W. Hirt and P. Chevillat, “Robust Wireless Infrared System with Channel Reciprocity”, to be published in *IEEE Communications Magazine*, December 1998.

Glossary

ACK	Acknowledgement frame
ACK-DATA	Acknowledged Data frame
AIr	Advanced Infrared
BER	Bit Error Rate/Ratio
BL	Body Length
BO	Back-Off
BOWS	Back-Off Window Size
CA	Collision Avoidance
CAS	Collision Avoidance Slot (= BO slot)
CCITT	International Telegraph and Telephone Consultative Committee (now ITU-T)
Ccts	CTS success rate
Cp	Packet success rate
Crts	RTS success rate
CRC	Cyclic Redundancy Check
CSMA	Carrier Sense Multiple Access
CSMA/CA	CSMA with Collision Avoidance
CSMA/CD	CSMA with Collision Detection
CTS	Clear-To-Send
DA	Destination Address
DATA	Data frame
DBP	Data Bit Pair
DD	Data symbol
DMA	Direct Memory Access
E	Irradiance
EOB	End-Of-Burst
EOBC	End-Of-Burst-Confirmed
FCS	Frame Check Sequence
FD	Frame Drop
FL	Frame Loss
FPGA	Field Programmable Gate Array
FSM	Finite State Machine
HDLC	High-Level Data Link Control
I	Radiant Intensity
ICR	Infrared Channel Reciprocity
IEEE	Institute of Electrical and Electronics Engineers
IMD	IBM Microelectronics Division
IrCOMM	IrDA protocol for emulation of serial and parallel communications
IrDA	Infrared Data Association
IrLAN	IrDA LAN access protocol
IrLAP	IrDA Link Access Protocol
IrLMP	IrDA Link Management Protocol
IrMAC	Advanced Infrared Medium Access Control
IrOBEX	IrDA Object Exchange Protocol
IrPHY	Infrared Physical layer
IR	Infrared
IRED	Infrared Emitting Diode
IRQ	Interrupt Request
ITU-T	International Telecommunications Union – Telecommunications (Sector)
kbps	Kilobits per second
LAN	Local Area Network
LM-IAS	Link Management Information Access Service

LM-MUX	Link Management Multiplexer
LED	Light Emitting Diode
MAC	Medium Access Control
Max_Retries	Maximum Number of Backoff Retries
MTD	Min Turn Delay (turnaround time)
Mbps	Megabits per second
MB	Main Body
MGMT	Management
NACK	No data acknowledgement
PA	Preamble
PD	Photodiode
PDA	Portable Digital Assistant
PLL	Phase Locked Loop
PPB	Packets Per Burst
PPM	Pulse Position Modulation
PO	Optical Parity
PS	System Parity
PT	Transceiver Parity
RF	Radio Frequencies
RH	Robust Header
RID	Reservation Identifier
RR	Rate Reduction (also: Repetition Rate)
RT	Reservation Termination
RTS	Request-To-Send
S	Sensitivity
SA	Source Address
SDL	Specification Description Language
SIR	Signal to Interference Ratio
SNR	Signal to Noise Ratio
SSN	Source Station Number
SYNC	Synchronization
TAT	Turnaround Time
TinyTP	IrDA Tiny Transport Protocol
TP	Throughput
TSN	Target Station Number
UART	Universal Asynchronous Receiver-Transmitter
TT	Transmission Time
UDATA	Unreserved Data frame
VR	Variable Rate
WFCTS	Wait For CTS
W_max	Maximum backoff window size
W_min	Minimum backoff window size
ZRL	Zurich Research Laboratory
4-PPM	Four-slot Pulse Position Modulation

Acknowledgments

I would like to thank the following people, for their support and contributions during the evolution of my industrial thesis:

Dr Pierre Chevillat, manager of the Wireless Communication Systems Group at IBM ZRL.
Dr Fritz Gfeller, my supervisor at IBM ZRL.
Dr Wolfgang Schott, who also supervised my work and allowed me to utilize some results obtained with its AIr MAC simulation model.
All the members of the Infrared Communication Systems Group at IBM ZRL, particularly Mr Beat Weiss.
The following members of the IBM Infrared Wireless Communications Group in Toronto, Canada:
Mr Danny McKay, Mr Mike Chang, Mr Patrick Kam and Mr Brian Varga.

Zurich, September 30, 1998
Dewis Bernasconi