CHAPTER 5

Advanced Infrared Medium Access Control Layer

The IrDA 1.x protocol has been proven very popular and millions of devices are equipped with an IrDA 1.x infrared port. However, IrDA 1.x specifications are addressing the 'point and shoot' user model; only one pair of devices can communicate in the same infrared space and the link range is limited. The significant increase on the number of mobile devices on the market today and recent advances in infrared technology have led to the decision to address the communication requirements of a pool of users.

IrDA proposed the Advanced Infrared (AIr) protocol specifications for indoor, high-speed, low cost and multipoint wireless communications. This chapter describes the AIr protocol and evaluates AIr performance by using simulation results. The outline of this chapter is as follows. Section 5.1 introduces the AIr protocol architecture and section 5.2 presents the frame formats supported by the AIr MAC layer. Section 5.3 discusses the AIr MAC transfer schemes including the Reserved and Unreserved transfer modes of the protocol. Section 5.4 presents the collision avoidance procedures of the AIr protocol and section 5.5 explains the interface between the MAC and PHY layers. Section 5.6 presents an OPNET simulator developed for the AIr protocol. Finally, sections 5.7 and 5.8 evaluate, using simulation results, the performance of the Unreserved and of the Reserved transfer modes respectively.

5.1 Architecture overview

The primary goal in developing AIr specifications was to introduce multipoint connectivity and, at the same time, preserve the investment in upper layer applications by making certain that existing IrDA 1.x applications will be able to utilize the proposed extensions in lower layers [101]. Fig. 5.1 presents the OSI standards level overview of the AIr protocol architecture. A new physical layer, AIr PHY [46], is introduced and the IrDA 1.x IrLAP layer [49] is split into three sub-layers, the AIr Medium Access Control (MAC) [45], the AIr Link Manager (LM) [42] and the AIr Link Control (LC) [43] sub-layers. IrDA 1.x and AIr LC procedures for establishing device-to-device connections is transparent to network layer entities.

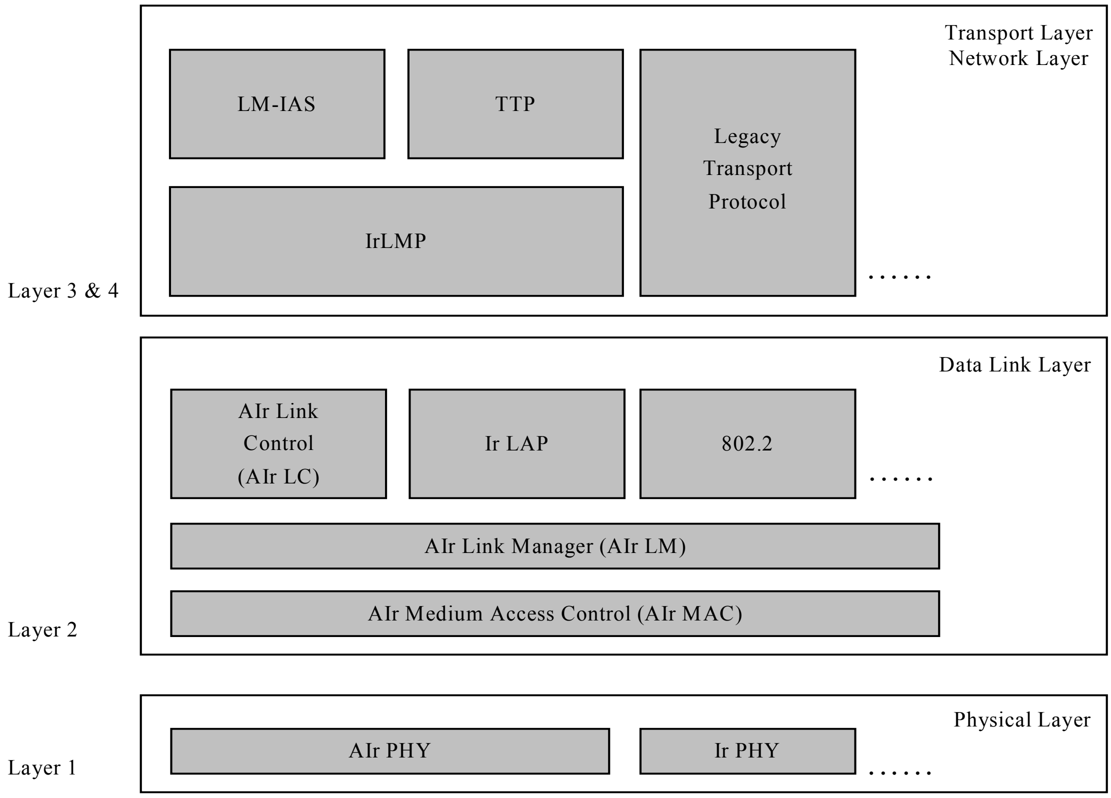

Figure 5.1 AIr architecture overview

AIr PHY supports wide-angle ports operating at ±60 to ±75 degrees, compared to narrow-angle ±15 to ±30 degrees for the IrDA 1.x, in order to achieve multipoint connectivity. AIr utilizes one common modulation format for all supported data rates. This format is defined as four slot Pulse Position Modulation with Variable Repetition Rate (RR) encoding (4PPM/VR). The base rate is 4 Mbit/s. Lower data rates (up to 256Kbit/s) may be utilized by repeating the transmitted symbols RR times at the base rate. The introduced redundancy improves link signal to noise ratio (SNR) at the expense of a lower data rate. It is expected that halving the data rate results in a 19% range increase at each reduction step [101]. Variable rate coding is introduced in [30] and physical layer issues for achieving the required channel symmetry are discussed in [29][14]. AIr physical characteristics with experimental results can be found in [28]. Standard range (S-class) AIr transceivers are expected to provide a transmission distance from 1m to 2.5m at 4Mbit/s. At 256 Kbit/s, a range from at least 2m to at least 5m is achieved. Long-range (L-class) AIr transceivers accomplish a transmission range from 2.5m to 6m at 4 Mbps and a range of at least 5m to at least 12m at 256 Kbit/s [46]. An AIr prototype port on a US quarter dollar presented by IBM is shown in Fig. 5.2.

Figure 5.2 AIr prototype port on a US quarter dollar presented by IBM

The AIr MAC sub-layer allows upper layers to cope with the relaxing of restrictions on the angle and range of AIr PHY ports. AIr MAC is responsible for coordinating the access to the infrared medium among AIr and IrDA 1.x devices. AIr MAC supports reservation based media access control, reliable and unreliable data transfer, data sequencing and data rate adaptation. AIr MAC coordinates medium access by employing Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) techniques [45]. AIr LM is a 'thin' layer that allows multiplexing of multiple different client protocols. It also provides dynamic addressing, station grouping and priority and non-priority data channels [42]. Dynamic addressing is used to cope with MAC address conflicts and station grouping is utilized to enable multicast transmissions. AIr LC supports connections to multiple devices. AIr LC is a derivative of the widely used HDLC protocol operating at the Asynchronous Balanced Mode of the protocol. AIr LC does not assign primary and secondary roles to communicating devices. It supports error detection and recovery services, address conflict resolution procedures and guaranteed data delivery services.

5.2 AIr MAC frame formats

The following Repetition Rate (RR) values are supported by AIr PHY and MAC layers:

Repetition Rate	Data rate
RR=1	4 Mbit/s
RR=2	2 Mbit/s
RR=4	1 Mbit/s
RR=8	512 Kbit/s
RR=16	256 Kbit/s

AIr MAC utilizes 12 frame types in total. Two general classes are defined; the reservation control frames and the data transfer frames. The following reservation control frames are used to contend, initiate and terminate reservations:

Type	Description
RTS	Request To Send
CTS	Clear To Send
SOD	Start Of Data
EOB	End Of Burst
EOBC	End Of Burst Confirmed

The following data transfer frames are use to transfer payload data:

Type	Description
DATA	Reserved data frame
ADATA	Reserved data frame with acknowledgment
UDATA	Unreserved data frame
SDATA	Reserved data frame with sequencing
ACK	Acknowledgment frame
SPOLL	Sequenced poll frame
SACK	Sequenced acknowledgment

The AIr MAC frame format definitions are shown in Fig. 5.3. Three major elements are defined:

a) The *initialisation* element consists of the Preamble (PA) and the synchronisation (SYNC) fields. These two fields allow the receiver to detect an incoming frame. The PA field consists of a predefined sequence that helps the Phase Locked Loop (PLL) to lock on to and detect the start of an incoming transmission. PA provides an early form of carrier sense and lasts for 64 µsec (TT_{PA}=64 µsec). The SYNC field qualifies the carrier detection and allows exact identification of the beginning of the robust header element. SYNC lasts for 40 µsec (TT_{SYNC}=40 µsec).

a) The *robust header* (RH) element contains the essential information required by the PHY and MAC layers to co-ordinate medium access. It is always transmitted using maximum RR=16 to allow reception of this essential information by all stations in range. Thus, all stations capable of interfering with the current transmission refrain from transmitting. RH contains 32 bits and lasts for 128 µsec (TT_{RH}=128 µsec).

c) The *main body* (MBR) element, which includes the Cyclic Redundancy Check (CRC). MBR contains non-essential MAC protocol information and is transmitted using variable RR. MBR contains payload data and has a variable length.

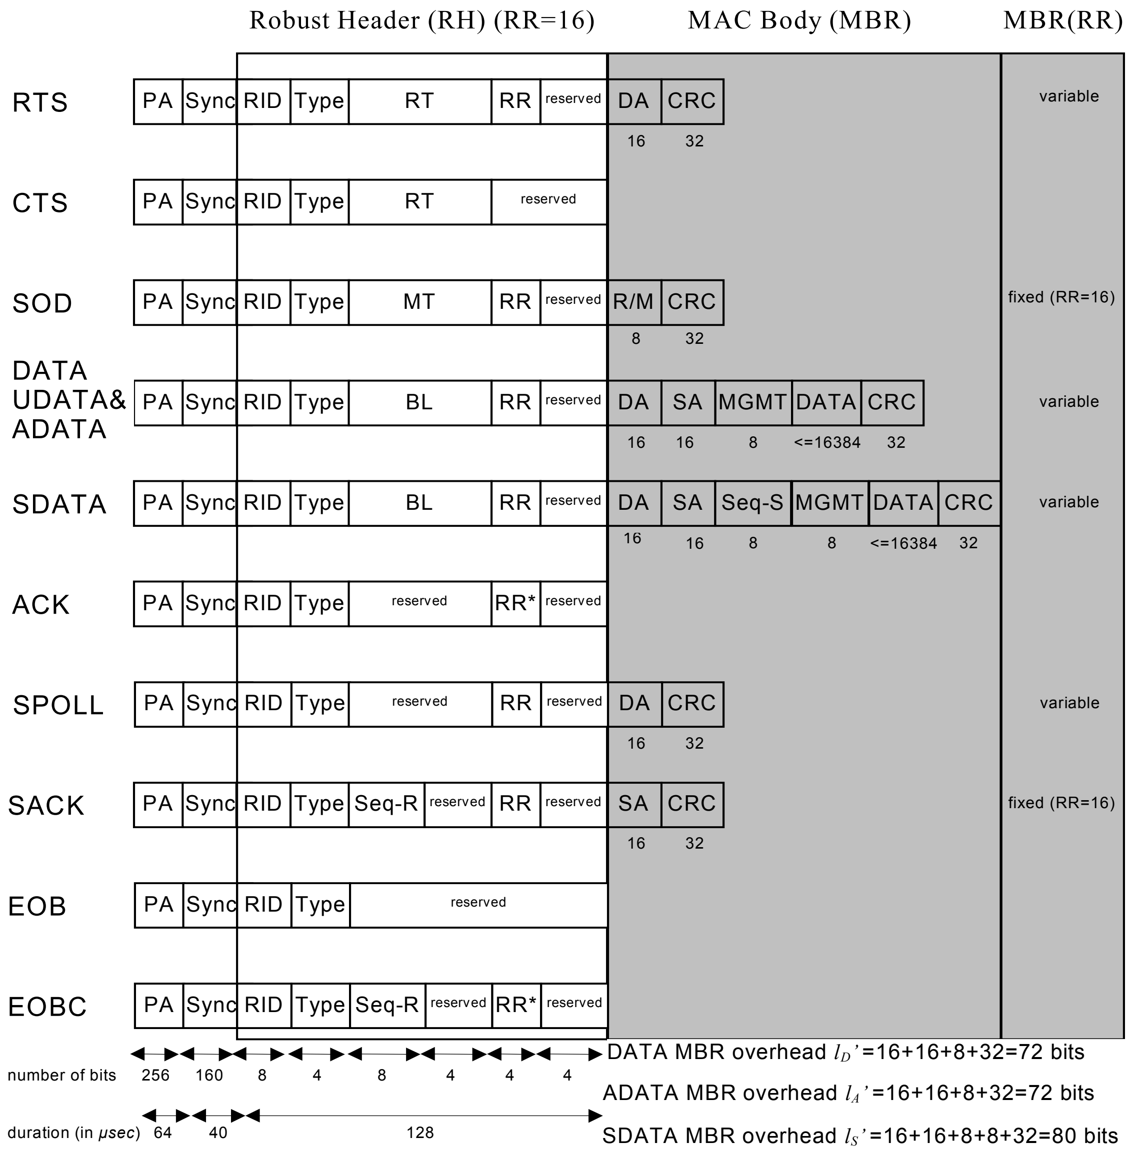

Figure 5.3 AIr MAC frame definitions

PA, SYNC and RH fields are present in all AIr MAC frame types. However, MBR is not present in some frame types. In this case, the RH field is not protected by a CRC because it is transmitted using maximum RR=16. The transmission time of frames with no MBR field is $T_{RH}=TT_{PA}+TT_{SYNC}+TT_{RH}=232$ μsec.

5.2.1 Robust Header (RH) fields

The RH fields contain all the PHY and MAC information needed by all stations in order to operate efficiently in accessing the infrared medium.

- Reservation Identifier (RID) field: This field specifies the Reservation Identifier of a reservation attempt or an on-going reservation. A station randomly selects a

RID value every time it attempts to reserve the infrared medium. RID associates AIr MAC frames with a specific reservation and acts as a shorthand notation for the source and destination addresses for the entire reservation duration.

- TYPE field: This field defines the type of frame. The frame type specifies the frame structure, which is used by all receiving stations to extract the information carried in the frame fields. It is also used to determine the suitable response (if any). There is no difference in the frame structure of DATA, UDATA and ADATA frames but different frame types are assigned to these frames because they require different actions from the receiver.
- Reservation Time (RT) field: This field specifies the length of the time that a reservation will last. It is present in both RTS and CTS frame types. Although EOB and/or EOBC frames normally inform all stations in range that current reservation is over, the RT field ensures that a reservation will terminate even in cases where overlapping reservations and collisions will prevent a normal reservation termination.
- Block Length (BL) field: The BL field defines the length of the frame in bytes.
- Modulation Time (MT) field: This field specifies the time duration a modulation will be in use. It is present only in SOD frames that support IrDA 1.x communications.
- Repetition Rate (RR) field: This field specifies the Repetition Rate used in the MBR field of the frame. It is present only in frames that have an MBR field (i.e. RTS, SOD, DATA, UDATA, ADATA, SDATA, SPOLL and SACK frame types).
- Sequenced Receive (SEQ-R) field: This field carries MAC sequencing information. It is present only in SACK and EOBC frame types and carries the next sequence number expected by the receiving station.
- Recommended Rate (RR*) field: This field is similar to RR field except that it specifies recommendations for the RR in the reverse direction based on the receiver's evaluation of the link quality. The RR* field is present only in ACK and EOBC frames.

5.2.2 Main Body (MBR) fields

The MBR field contains upper layer and non-essential AIr MAC protocol information. When present, it contains a CRC field that protects both RH and MBR fields.

- Rate/Modulation (R/M) field: This field specifies the rate and modulation format of the IrDA 1.x transmission
- Source and Destination Address (SA & DA) fields: These fields specify the Source Address (SA) and Destination Address (DA) of the communicating stations. When appropriate, the DA address can be a multicast or the broadcast (i.e. all stations) address.
- MGMT field: This field carries upper layer control information in frames with payload data.
- Sequenced Send (SEQ-S) field: This field carries the MAC generated sequence numbers on SDATA frames. The first SDATA frame in a reservation always has sequence number 0.
- DATA field: This is a variable length field that carries information data provided by the upper layers. The maximum field length is 16Kbits.
- CRC field: This field carries the CRC of the frame.

5.3 AIr MAC transfer modes

When the transmitter is sending a frame, it 'blinds' its own receiver such that it can not receive remote infrared pulses. The transmitter's receive circuitry needs a minimum Turn Around Time (TAT) to recover. As a result, when a frame reception is completed, the receiving station must always wait a TAT delay before transmitting its response to ensure that the intended station's receive circuitry will have successfully recovered. According to AIr MAC specification, TAT duration is 200μsec. This duration is significantly higher than the TAT period of similar WLAN protocols, such as the IEEE 802.11 [39][95].

AIr MAC supports medium reservation utilizing the Request To Send (RTS) / Clear To Send (CTS) frame exchange. The transmitting station reserves the medium for the duration contained in the Reservation Time (RT) field of the RTS frame it transmits. After a TAT delay, the receiving station echoes the reservation period in the RT field of the responding CTS frame. Thus, even stations being able to hear only the RTS or only the CTS frame refrain from transmitting for the entire reservation period. The RTS/CTS scheme is used to address the hidden station (a station not being able to hear the transmitter or the receiver) problem [31][60][61] at the expense of the time required for transmitting the RTS and CTS frames. When the CTS frame is successfully received, the transmitter waits a TAT delay and initiates a window frame transmission. After the last data frame is transmitted and before the reservation time expires, the transmitter can request termination of current reservation by transmitting an End Of Burst (EOB) frame. The receiver waits a TAT period and responds with an End Of Burst Confirm (EOBC) frame confirming termination of current reservation. As with RTS/CTS exchange, a station hearing only the EOB or the EOBC frame realizes that the current reservation is over and that it is able to contend for the medium again. AIr MAC implements two timers (EXIT1 and EXIT2) to ensure that stations hearing only the EOB or the EOBC frame will contend for medium access at exactly the correct time.

AIr MAC provides reserved and unreserved transfer modes. AIr transfer modes are shown in Fig. 5.4.

5.3.1 Unreserved transfer mode

Unreserved transfer mode (Fig. 5.4(a)) does not utilize the RTS/CTS exchange and transmits only one UDATA frame to a multicast or broadcast (i.e. all devices) address. Unreserved transfer mode is unreliable because no acknowledgment is received. Protocol specifications suggest that UDATA frames should be transmitted at maximum RR = 16 to ensure maximum coverage.

5.3.2 Reserved transfer modes

5.3.2.1 Reserved transfer mode with DATA frame

Reserved transfer mode with no acknowledgment (Fig. 5.4(b)) reserves the medium using the RTS/CTS frame exchange, transmits a window of DATA frames in a successful reservation, terminates the reservation using the EOB/EOBC frame exchange but it is still unreliable as no acknowledgment is received. When an unreliable data transfer mode is used, (Fig. 5.4(a) & (b)), a MAC successful transmission indication to LM layer means that the frames are sent and not that the frames are correctly received.

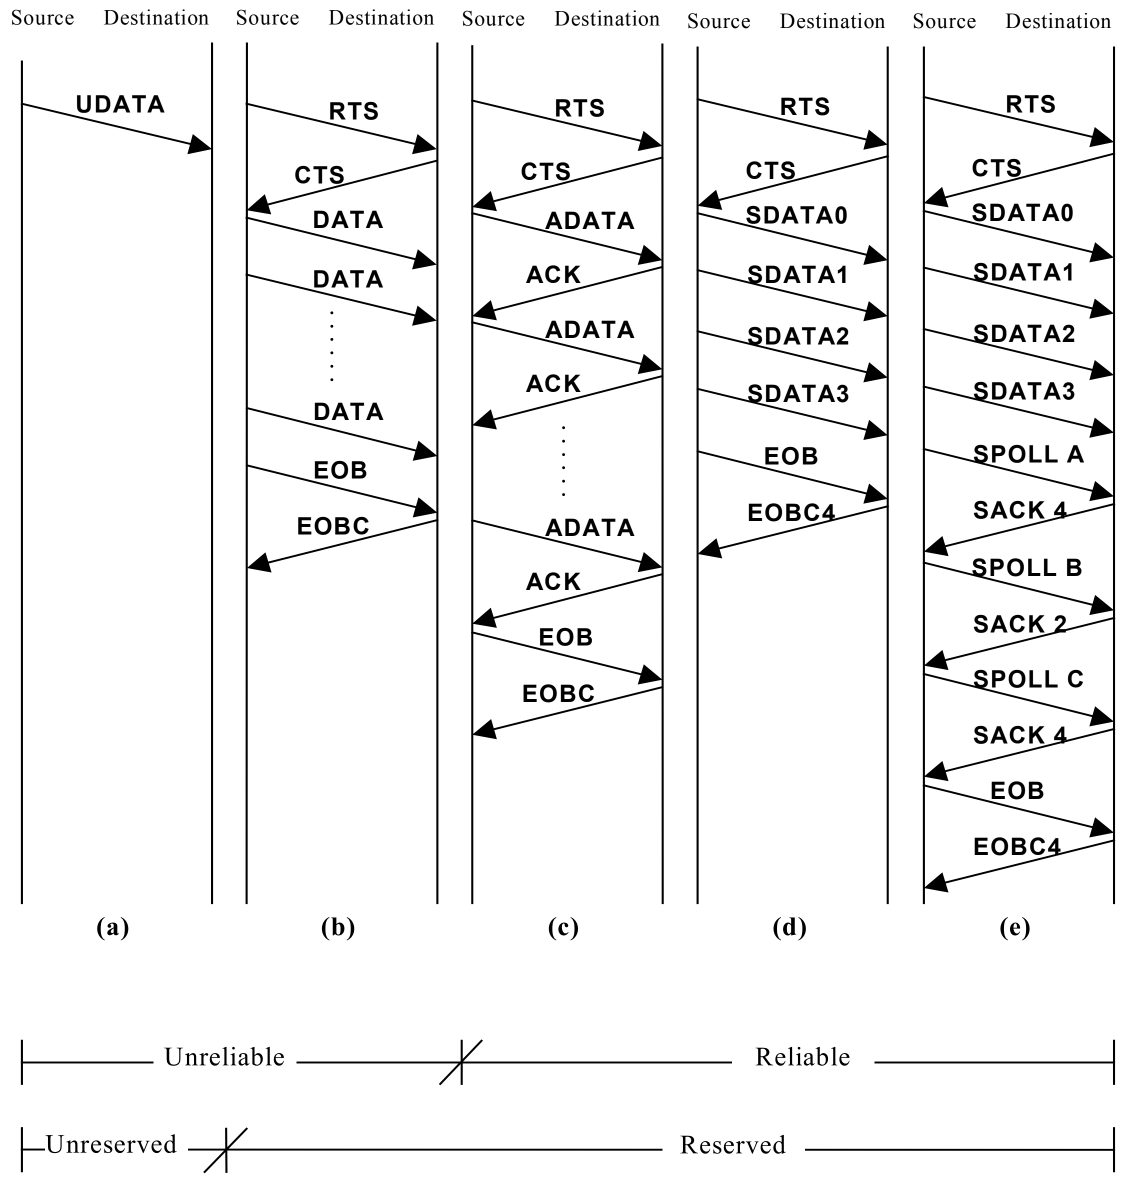

(a) Unreserved transfer mode with UDATA frame
(b) Reserved transfer mode with DATA frame (no acknowledgment)
(c) Reserved transfer mode with Acknowledgment (ADATA frame)
(d) Reserved transfer mode with Sequenced Data (SDATA frame)
(e) Reserved transfer mode with reliable multicast

Figure 5.4 AIr transfer modes

In this case, the LC layer implements a retransmission scheme to handle frame errors.

5.3.2.2 Reserved transfer mode with acknowledgment

Reserved transfer mode with frame acknowledgment (Fig. 5.4(c)) utilizes the RTS/CTS reservation scheme, ADATA frames to carry payload data and transmits a window of frames in a successful reservation. Successful ADATA frame reception is based on an immediate ACK (acknowledgement) frame transmitted by the receiver. Finally, the reservation is terminated by using the EOB/EOBC control frame exchange.

ACK frames contain RR recommendations based on the receiver's evaluation for the link quality. The transmitter's LM layer specifies when its MAC layer can change the RR it implements during a reservation based on the recommended rates provided by the receiver. Only this transfer mode can quickly adapt the RR to meet varying channel conditions because it includes multiple RR recommendations in ACK frames during a successful reservation.

5.3.2.3 Reserved transfer mode with sequenced data

Reserved transfer mode with sequenced data (Fig. 5.4(d)) uses the RTS/CTS frame exchange and sequences transmitted data using SDATA frames. An SDATA frame contains its sequence number in the Seq-S field. This mode terminates a reservation by using the EOB/EOBC frame exchange but, in this case, the EOBC frame contains the next frame sequence number expected by the receiver in the Seq-R field. Thus, the transmitter is informed of the correctly in-sequence received frames when the reservation terminates.

5.3.2.4 Reserved transfer mode with reliable multicast

Reserved transfer mode with reliable multicast (Fig. 5.4(e)) is similar to reserved transfer mode with sequenced data but transmits data to a multicast address. As a result, each target station must be separately polled at the end of the SDATA frame transmissions. An 'SPOLL *n*' frame is used to poll station *n*, which responds with an SACK frame indicating in the Seq-R field the next frame sequence number expected. When a reliable data transfer mode is used, (Fig. 5.4(c), (d) & (e)), a MAC successful transmission indication to LM layer means that the frames are correctly received.

5.4 Collision avoidance procedures

AIr MAC employs Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) techniques to minimize collision probability. A station wishing to transmit and regardless of the transfer mode it employs, it first invokes the Collision Avoidance (CA) procedures in an effort to minimize collisions with other stations. Considering the AIr transfer modes presented in the previous section, a contending station always

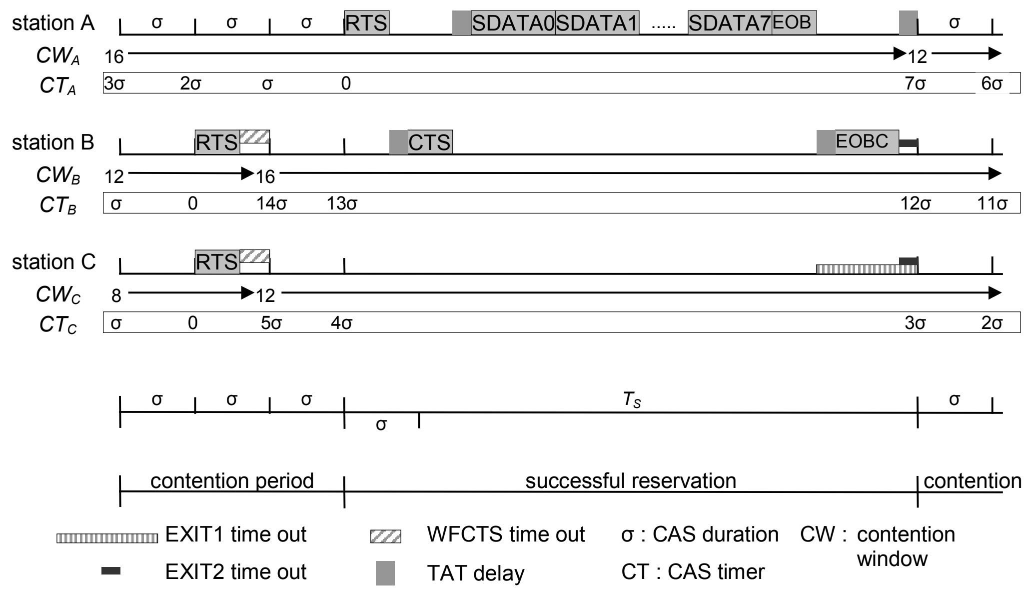

Figure 5.5 Operation of Collision Avoidance procedures

invokes the CA procedures before an RTS and a UDATA frame transmission. If the medium is busy, the contending station first waits until the medium is idle and the next contention period starts. Contention periods start a TAT period after the UDATA or the EOBC frame that terminates a reservation. The TAT delay is required because the station that transmitted the UDATA or EOBC frame should be able to receive the next frame (Fig. 5.5). Contention periods end when a non-colliding RTS or UDATA frame is transmitted. AIr MAC specification ensures that even stations hidden from the transmitter or the receiver are synchronized in contending for the medium at the beginning of the next contention period after a successful RTS/CTS medium reservation. Synchronization is accomplished by means of the EXIT1 and EXIT2 timers. EXIT1 is started when the EOB frame is received. EXIT1 expires at the beginning of the next contention period synchronizing stations not receiving the EOBC frame. EXIT2 is started when the EOBC frame is received and, as it accounts for the TAT delay, it expires when the next contention period starts (Fig. 5.5). Thus, stations hearing only the EOB or the EOBC frame will contend for medium access at exactly the correct time [45].

The contention period is slotted and a station is allowed to transmit only at the beginning of a Collision Avoidance Slot (CAS). The CAS duration (σ) is defined as being greater than the transmission time of the RTS frame plus the TAT delay plus the time need to detect the PA and SYNC fields of the responding CTS frame ($\sigma > T_{RTS}$+TAT+TT_{PA}+TT_{SYNC}). AIr MAC defines that σ=800 μsec. Such a long CAS duration aims to avoid collisions caused by stations hidden from the transmitter that are not able to hear the RTS frame transmission. These stations will detect the beginning of the CTS frame during the CAS duration, realize that another competing station has selected this CAS and refrain from transmitting [95].

The CAS duration AIr protocol employs is significantly longer than the CAS duration similar CSMA/CA protocols define. These protocols define that the CAS duration should only be long enough for a station to detect the beginning of an incoming RTS frame. For example, the IEEE 802.11 specification [39] defines that the CAS duration accounts for the propagation delay, for the time needed to switch from receiving to transmitting state and for the time needed for the physical layer to signal the channel state to the MAC layer. AIr MAC defines a longer CAS duration that provides a much better hidden station approach at the expense of longer contention periods caused by the longer CAS duration.

As discussed earlier, when the medium is busy, a competing station for medium access first waits for the transmitting station to finish and for the beginning of the next contention period. It then selects a random number of CASs to wait before transmitting and assigns the deferral period to the CAS timer (CT). This integer random number is uniformly chosen in the range (0, *CW*-1), where *CW* is the current Contention Window (CW) size. CW size represents the range the random number is picked from and competing stations may utilize different *CW* values during the contention period. If during the station's deferral period another transmission is observed, the station freezes the CAS timer and restarts it again when the on-going transmission is finished and the next contention period is started. When the CAS timer reaches zero, the station transmits.

If the station employs a Reserved transfer mode scheme, it transmits the RTS frame and starts the Wait For CTS (WFCTS) timer. If another (or more) station has selected the same CAS, as in the second CAS in Fig. 5.5, it transmits its RTS frame at the same time. The resulting collision is determined by the transmitting stations by the WFCTS timer expiration. The colliding stations select a new CAS and continue contending for medium access. To synchronise the colliding stations with the remaining stations, the

WFCTS timer should expire at the end of the current CAS.

The *CW* parameter is of great importance because in the standard a long CAS duration is defined. If a small *CW* value for the specific network size and load is used, a low utilization figure will be reached caused by the increased number of collisions. If a large *CW* value is used, the increased number of empty CAS will result in low medium utilization. A station can only estimate the appropriate *CW* value it should implement based on the experienced successful reservations and collisions. AIr specifications define that the LM layer selects the *CW* value to be used in every reservation attempt and pass it down to the MAC layer [42].

The AIr LM specification defines guidelines for adjusting *CW*. It presents a linear algorithm that increments and decrements *CW* after a collision and a successful reservation respectively. The transmitter 'remembers' the *CW* value used in the previous reservation attempt. If this attempt was successful, *CW* is decreased by 4 (see station A in Fig. 5.5); if it resulted in a collision, *CW* is increased by 4 (stations B and C in Fig. 5.5). A minimum *CW* value of 8 and a maximum *CW* value of 256 are also defined [42].

Fig. 5.5 illustrates the station behavior for a LAN with three contending stations employing the Reserved transfer mode with sequenced data. It is assumed that at the beginning of the first contention period, station A has a current *CW* value of 16 (CW_A=16) and defers transmission for three CAS. The CT of station A has a value of 3σ (CT_A=3σ), where σ is the CAS duration. We also assume that Station B has CW_B=12, CT_B=σ and station C has CW_C=8, CT_C=σ. As all stations defer transmission, the first CAS is empty. All stations decrease their CTs for the entire CAS duration. As a result, at the beginning of the second CAS, CT_A=2σ, CT_B=0 and CT_C=0. The deferral period of stations B and C has expired and these stations transmit an RTS frame at the beginning of the second CAS resulting in a collision. As no CTS frames are generated, stations B and C realize the collision by the expiration of their WFCTS timers. As explained earlier, the WFCTS timer expiration occurs at the beginning of the next CAS synchronizing stations B and C with station A. Stations B and C increase their *CW* values (CW_B=16 and CW_C =12) and select new deferral periods. Assuming station B selects 14 and station C selects 5, CT_B=14σ and CT_C=5σ at the beginning of the third CAS. At this point CT_A=σ as station A was decreasing its CAS timer for the entire duration of the second CAS because it did not receive any valid frame. As all stations

PSAP Primitive	abbreviation	type	explanation
AIr PHY Carrier Sense Detect	CSD	indication	Preamble (PA) portion of an incoming frame detected
AIr PHY Carrier Sense Confirm	CSC	indication	Synchronization (SYNC) portion of an incoming frame detected
AIr PHY Robust Header Received	RHR	indication	Robust Header (RH) portion of an incoming frame received
AIr PHY Main Body Received	MBR	indication	Main Body (MBR) portion of an incoming frame received
AIr PHY Frame Transmit	FTX	request	MAC requests a frame transmission
AIr PHY Frame Transmission Completed	FTC	indication	MAC requested frame transmission completed

Table 5.1 AIr Physical layer Service Access Point primitives

defer transmission, the third CAS is empty. Station A transmits an RTS at the beginning of the fourth CAS and successfully reserves the infrared medium. After the RTS/CTS exchange, station A transmits a window of SDATA frames and requests reservation termination by an EOB frame. The receiving station (station B) confirms reservation termination using an EOBC frame. As explained earlier, the next contention period starts a TAT delay after the EOBC frame and all stations are synchronized by the EXIT1 and EXIT2 timers. As station A has successfully reserved the medium, it decreases its CW value (CW_A=12), selects a new deferral period (7σ for example) and contention for medium access continues.

5.5 Physical layer service access point

AIr MAC layer uses the service primitives provided by the AIr physical layer through the AIr Physical layer Service Access Point (PSAP) [46]. AIr PHY passes frame reception information to MAC layer by PSAP service primitives called indications. AIr MAC calls for frame transmissions by invoking PSAP primitives called requests. PSAP indications and requests concerning frame transmissions are presented in Table 5.1. Fig. 5.6 (a) and (b) shows the PSAP primitives for a frame reception and a frame transmission respectively. When a station receives an incoming frame, AIr PHY layer indicates reception of PA, SYNC, RH and MBR fields separately by using the CSD, CSC, RHR and MBR indications respectively. AIr MAC requests a frame transmission by invoking a PSAP FTX request; when the frame transmission is completed, AIr PHY invokes a FTC indication.

The remaining sections present an AIr MAC simulator implementing the described

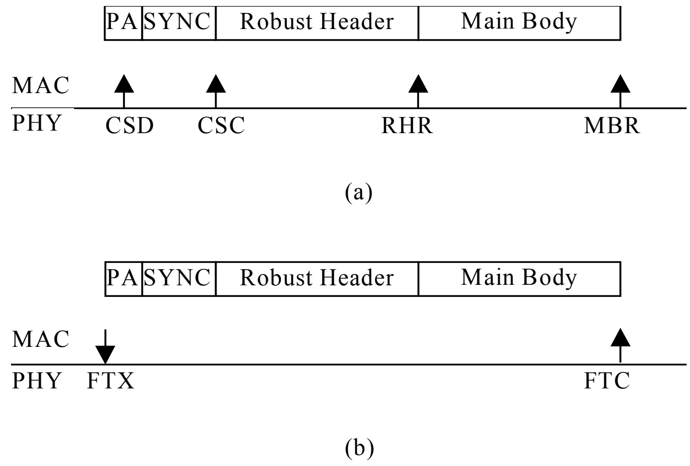

Figure 5.6. Physical layer Service Access Point primitives (a) frame reception (b) frame transmission

frame types, collision avoidance procedures and PSAP indications and requests. Simulation results are presented for the Unreserved transfer mode and for the Reserved transfer mode with sequenced data in an effort to evaluate the effectiveness of the presented collision avoidance procedures, PSAP primitives and transfer modes.

5.6 AIr simulator using OPNET

A simulator for the AIr protocol is developed using the OPNET™ package from MIL3 Inc. Infrared transmissions are emulated by altering the radio version of the OPNET modeler. Simulation models for the AIr PHY and AIr MAC layer operations are developed. The LM contention window adjustment algorithm is also implemented. The simulator emulates the real operation of a station as closely as possible, by implementing the collision avoidance procedures and all parameters such as frame transmission times and turn around times.

OPNET uses hierarchically linked domains to denote a network design. Stations are defined in the network domain, which is the top-level domain. Each station has a set of processes specified at the lower level, the node domain. Each process can represent a layer in the protocol stack or physical layer transmitters and receivers. A process can be defined by a finite state machine. User written C code to be executed when entering and exiting each state can be defined. Finally, the C code is accumulated and compiled. OPNET is an event driven simulator and provides a powerful graphical tool to display simulation statistics.

The OPNET AIr simulator closely follows all timer values and frame element

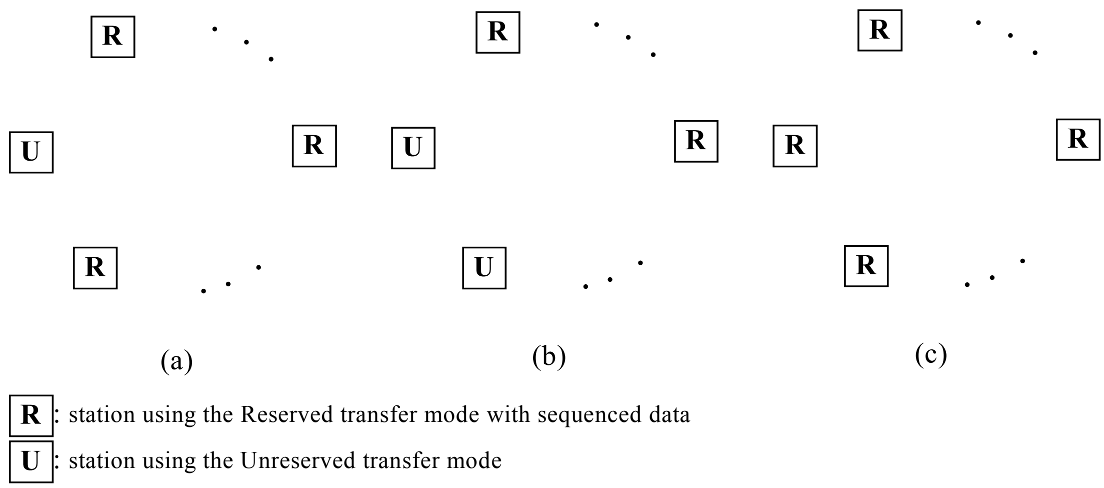

(a) one station use the Unreserved transfer mode
(b) two stations use the Unreserved transfer mode
(c) no stations use the Unreserved transfer mode

Figure 5.7. LAN simulation scenarios with stations employing the Unreserved transfer mode

transmission times defined by AIr specifications. However, a different value for the Wait For CTS (WFCTS) timer is implemented. WFCTS timer operation is explained in section 5.4. The implemented WFCTS timer value of 556 μsec is different from the protocol proposed value of 632 μsec but still obeys the protocol restriction $T_{WFCTS}>432\ \mu sec$. The implemented value was selected because it directs the colliding stations to contend again for medium access at the beginning of the next CAS [97].

A set of simulation runs was taken [7] to examine the performance of the AIr protocol under ideal channel conditions; an error free medium is assumed and no hidden stations are considered. Simulation runs that examine the performance of the Unreserved transfer mode are presented in section 5.7. The performance of the CA procedures and of the Reserved transfer modes is examined in section 5.8. All simulation runs consider a LAN of *n* transmitting stations operating at saturation conditions, i.e. all stations always have a frame ready for transmission.

5.7 Simulation results for the Unreserved transfer mode

According to the AIr MAC specification [45], Unreserved mode will be employed only for multicast and broadcast transmissions. Thus, it is expected that only a small number of stations will employ the Unreserved transfer mode in real life scenarios. Two different scenarios are considered; in the first (Fig. 5.7(a)) only one and in the second

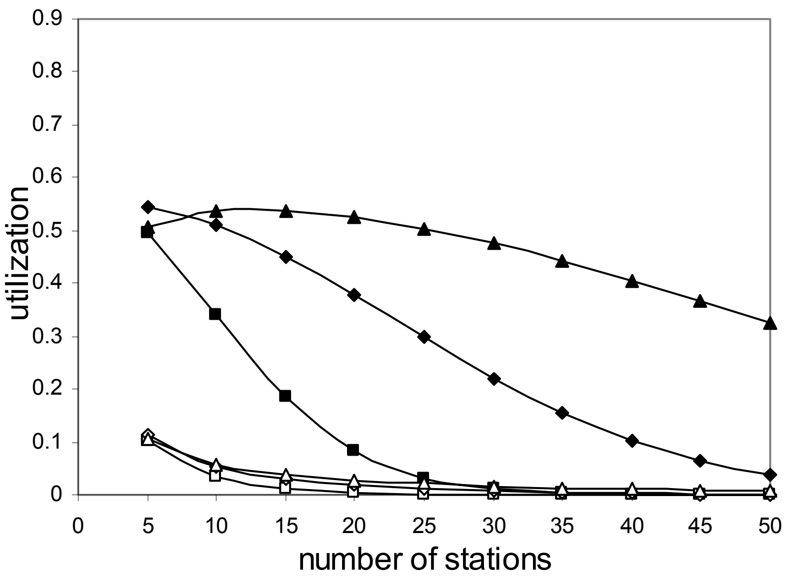

■ *CW=8, total utilization* □ *CW=8, UDATA utilization*
◆ *CW=16, total utilization* ◇ *CW=16, UDATA utilization*
▲ *CW=32, total utilization* △ *CW=32, UDATA utilization*

Figure 5.8 Utilization versus n for fixed CW, C=4 Mbit/s, l=16Kbits, 1 UDATA station

(Fig. 5.7(b)) only two stations implement the Unreserved transfer mode. All the remaining stations always employ the Reserved transfer mode with sequenced data. This reservation scheme was chosen because it incurs the least overhead among the AIr MAC reservation schemes and, at the same time, acknowledges frames at the MAC layer. The Unreserved transfer mode transmits only one UDATA frame at every transmission attempt. To allow comparison between UDATA and SDATA frame transmissions, stations employing the Reserved transfer mode transmit only one SDATA frame in every successful reservation. SDATA and UDATA frames always carry 16Kbits of payload data at RR=1 for the 4Mbit/s data rate. Simulations run for 15 sec and calculate utilization as defined in section 2.7. In particular, they calculate total utilization, which is defined as the time portion the medium transfers SDATA and UDATA payload data between stations. They also calculate UDATA utilization, which is defined as the time portion the medium transfers only UDATA payload data between stations.

Fig. 5.8 plots total (SDATA and UDATA) and UDATA utilization versus network size for different fixed *CW* values. The figure shows drastic utilization decrease for large network sizes due to the increased number of collisions. As collisions may involve

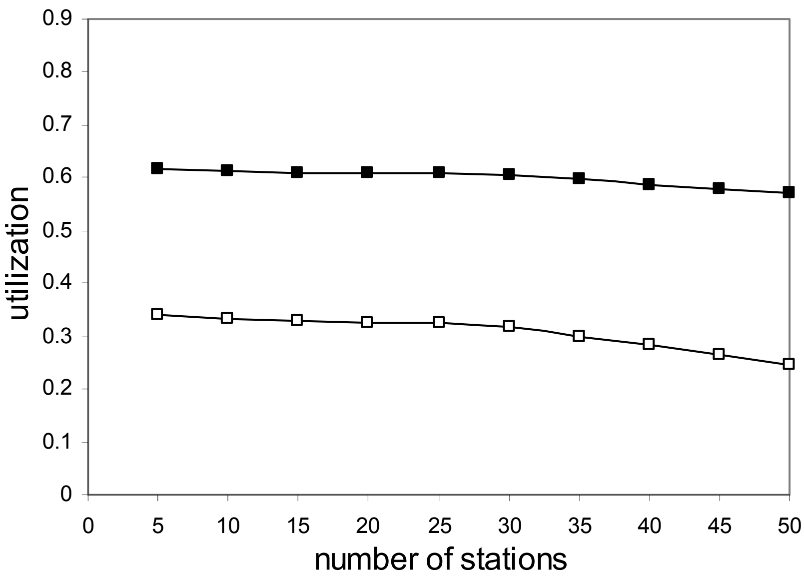

■ *total utilization* □ *UDATA utilization*

Figure 5.9 Utilization versus n, CW_{min}*=8,* CW_{max}*=256, C=4 Mbit/s, l=16 Kbits, 1 UDATA station*

a large UDATA frame, collision duration is very long and significantly affects LANs implementing small *CW* values because the collision probability is high in this case. As a conclusion, the selection of 'proper' *CW* value is of great importance when a number of stations employ the Unreserved transfer mode. The degradation of UDATA utilization for large network sizes can be easily explained by considering that the number of SDATA stations increases with network size increase but one station always employs UDATA transmissions.

Fig. 5.9 plots total (SDATA and UDATA) and UDATA utilization versus network size for the CW adjustment algorithm proposed by the protocol as described in section 5.4. Comparison of figures 5.8 and 5.9 indicates that the CW adjustment algorithm significantly increases UDATA utilization, especially for large network sizes where a large number of stations employ the Reserved transfer mode. The situation is explained as follows. When a UDATA frame is involved in a collision, the collision lasts for the time required for the complete UDATA frame transmission. The remaining stations are unaware of the collision existence since the AIr PHY layer, as shown in Table 5.1 and Fig. 5.6, does not send an indication to the MAC layer unless a valid PA portion of an incoming frame is detected. The remaining stations continue contending for medium access resulting in unsuccessful reservation attempts in the collision duration. As the

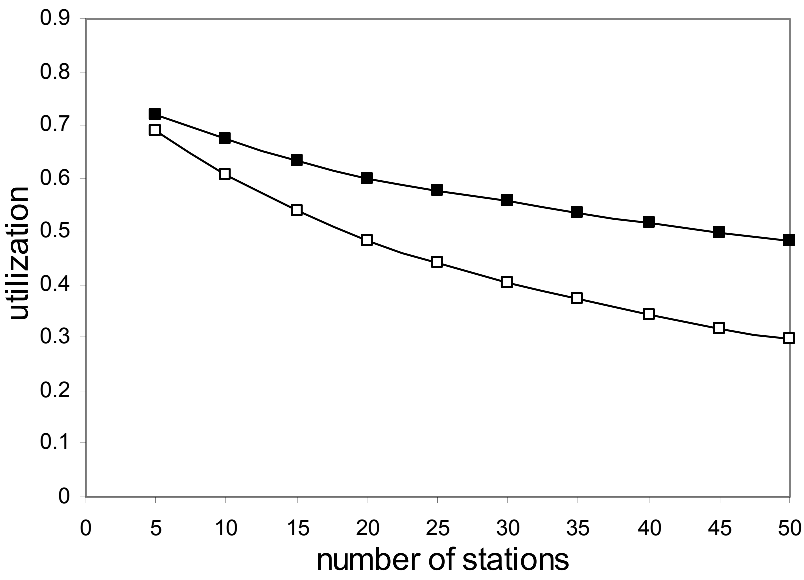

■ *total utilization* □ *UDATA utilization*

Figure 5.10 Utilization versus n, CW_{min}=8, CW_{max}=256, C=4 Mbit/s, l=16 Kbits, 2 UDATA stations

CW is increased after every unsuccessful reservation attempt, the stations employing the Reserved data transfer mode end up using large CW sizes and having a low probability of gaining access to the infrared medium. At the same time, the station employing the Unreserved data transfer mode uses the smallest CW size and transmits a UDATA frame often. The reason is that this station is unaware of whether its UDATA frame collides or not because it does not expect a CTS frame response or a MAC layer acknowledgment. As a result, no CW adjustment takes place. Upper layers are informed of the reception status of the UDATA frames but protocol specifications do not pass this information down to the LM layer, which is responsible for adjusting the CW.

Fig. 5.10 plots the same results when two transmitting stations employ the Unreserved transfer mode (Fig. 5.7(b)). It shows that total utilization is more sensitive to network size increase because the probability of long collisions is increased. It also shows that UDATA utilization is increased and UDATA frame transmissions dominate over SDATA transmissions because the Reserved stations increase their *CW* values more frequently. As a conclusion, an unexpected behavior is observed when the Unreserved transfer mode is employed because the PSAP does not include an indication to the MAC layer that the medium is busy. As a result, when a UDATA frame collides, the remaining stations contend for medium access during the collision. The resulting

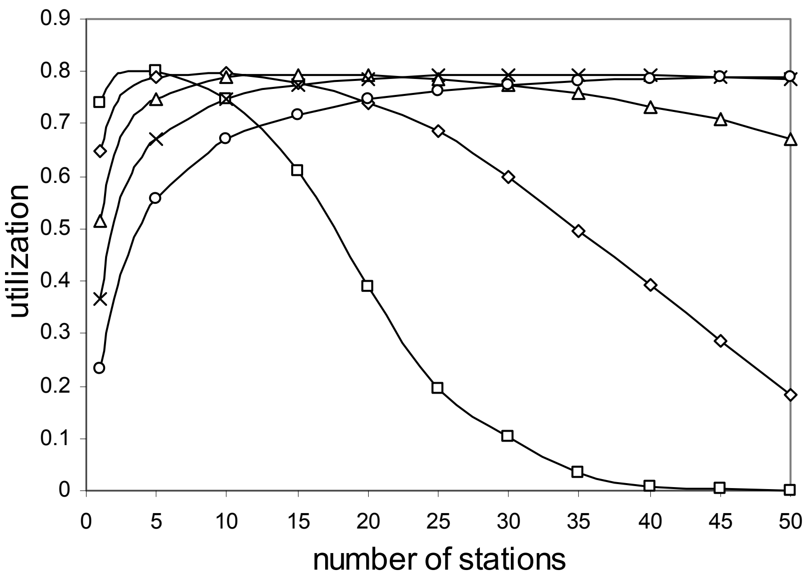

□ *CW=8* ◇ *CW=16*
△ *CW=32* x *CW=64*
○ *CW=128*

Figure 5.11 Utilization versus n for fixed CW, C=4 Mbit/s, l=16 Kbits, w=4

unsuccessful reservation attempts cause successive *CW* increments and a fairness problem. Employment of the Unreserved transfer mode should be avoided in LANs with many transmitting stations.

5.8 Simulation results for the Reserved transfer mode with sequenced data

Simulation results for LAN scenarios with all *n* stations employing the Reserved transfer mode with sequenced data are produced (Fig. 5.7(c)). Stations transmit *w* SDATA frames in every successful reservation attempt where *w* is the window size. SDATA frames always carry 16Kbits of payload data at RR=1 for the 4Mbps data rate. Simulations run for 15 sec.

Fig. 5.11 plots utilization versus network size for various fixed *CW* values. If a small *CW* is implemented, utilization seriously degrades when the number of stations increases due to the increased number of collisions. If a large *CW* is used, utilization degradation is observed for a few contending stations caused by the large number of empty CAS. Thus, optimum *CW* implementation becomes of key importance if maximum utilization is to be achieved. Fig. 5.11 reveals the interesting result that maximum utilization when the optimum *CW* value is employed for the specific network

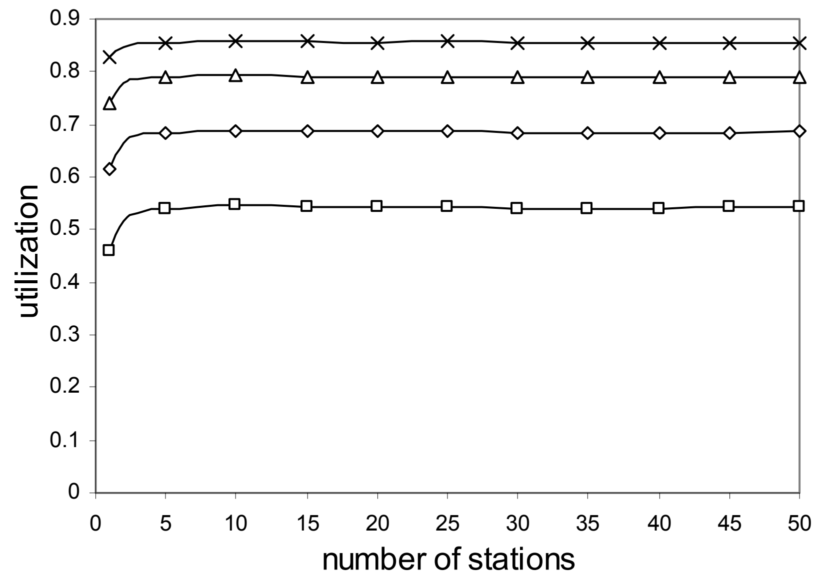

□ *w=1* ◇ *w=2*
△ *w=4* x *w=8*

Figure 5.12 Utilization versus n, C=4 Mbit/s, l=16 Kbits

size is practically independent of the network size.

Fig. 5.12 plots utilization versus network size for various w values when the proposed CW adjustment algorithm is implemented. It shows that utilization is relatively very low for low w values. This is an important conclusion because the considered SDATA frame transmissions incur the least overhead, the maximum frame size is used, no hidden stations are considered and, more important, an error free medium is assumed. As real life AIr networks may not be able meet these conditions, a lower utilization performance is expected. Fig. 5.12 also reveals that utilization is practically independent of network size when n is greater than 5 but lowers for LANs with less than five transmitting stations. Utilization reaches its lowest value when only one station is transmitting, a usual real life scenario. This is an unexpected result because no collisions occur in this case. These results call for further study of the collision avoidance procedures of the AIr protocol. Thus, an analytical model for these procedures is developed in the next chapter. This model is expected to reveal the effectiveness of PHY and MAC parameters on utilization. Please note that the same figures (Fig. 6.5 and 6.8) are produced by employing the proposed analytical model.