

Figure 10: throughput Vs maximum turn-around time for 16 Mbps

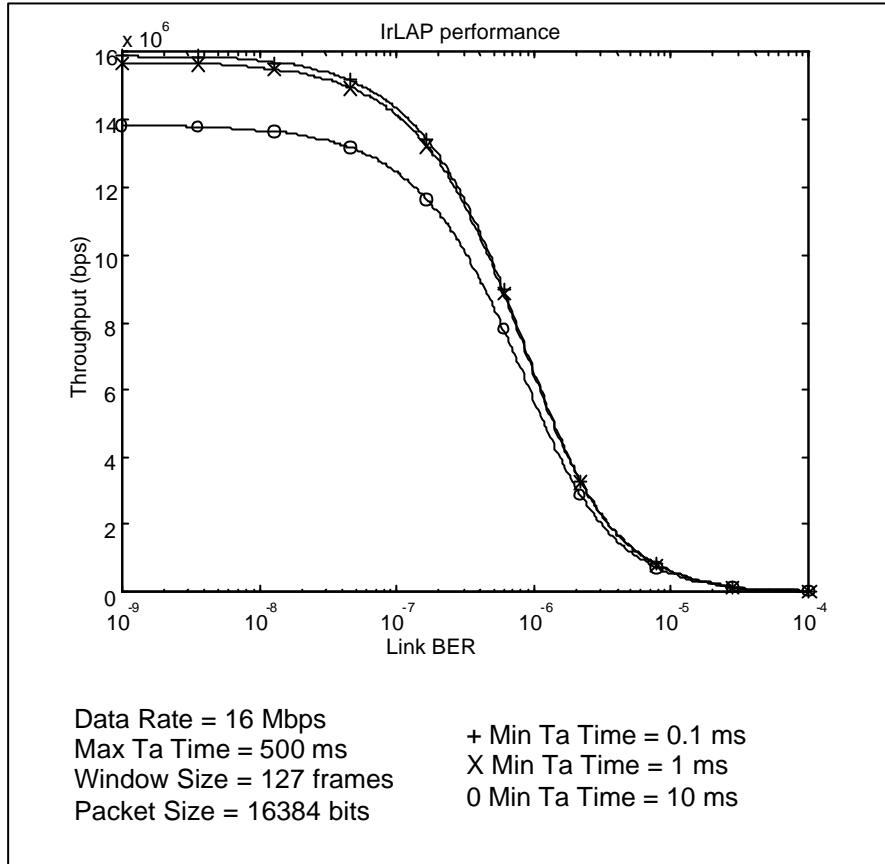


Figure 9: throughput Vs BER for 16 Mbps




Figure 8: throughput Vs packet data size for 4 Mbps

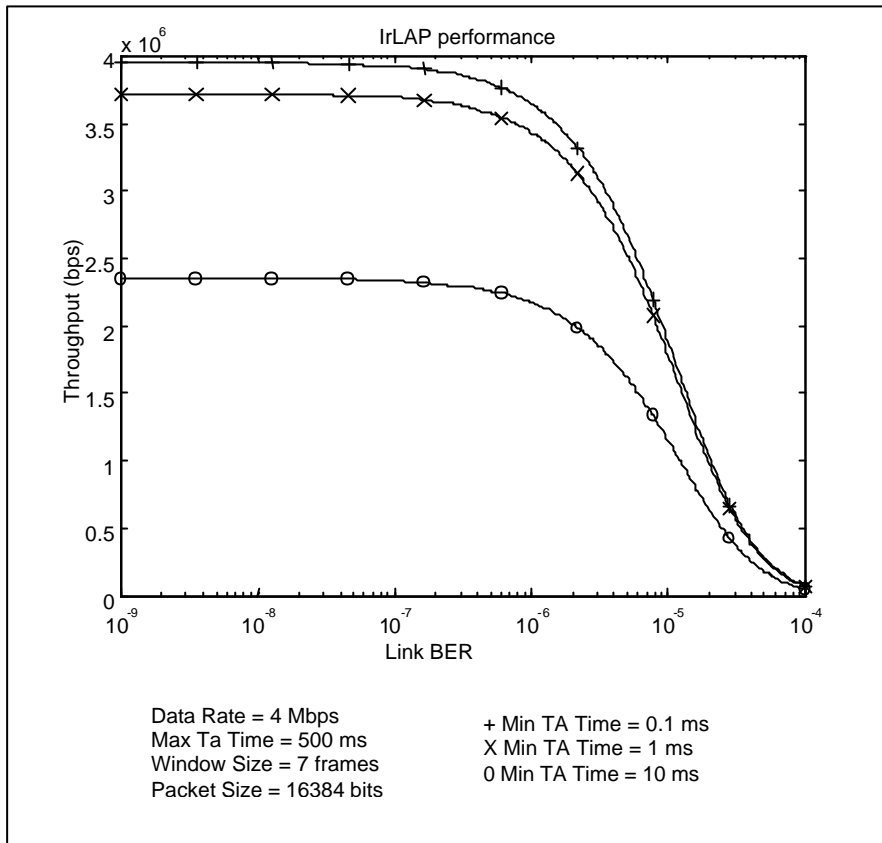


Figure 7: throughput Vs BER for 4 Mbps




Figure 6: throughput Vs packet data size for 115200 bps

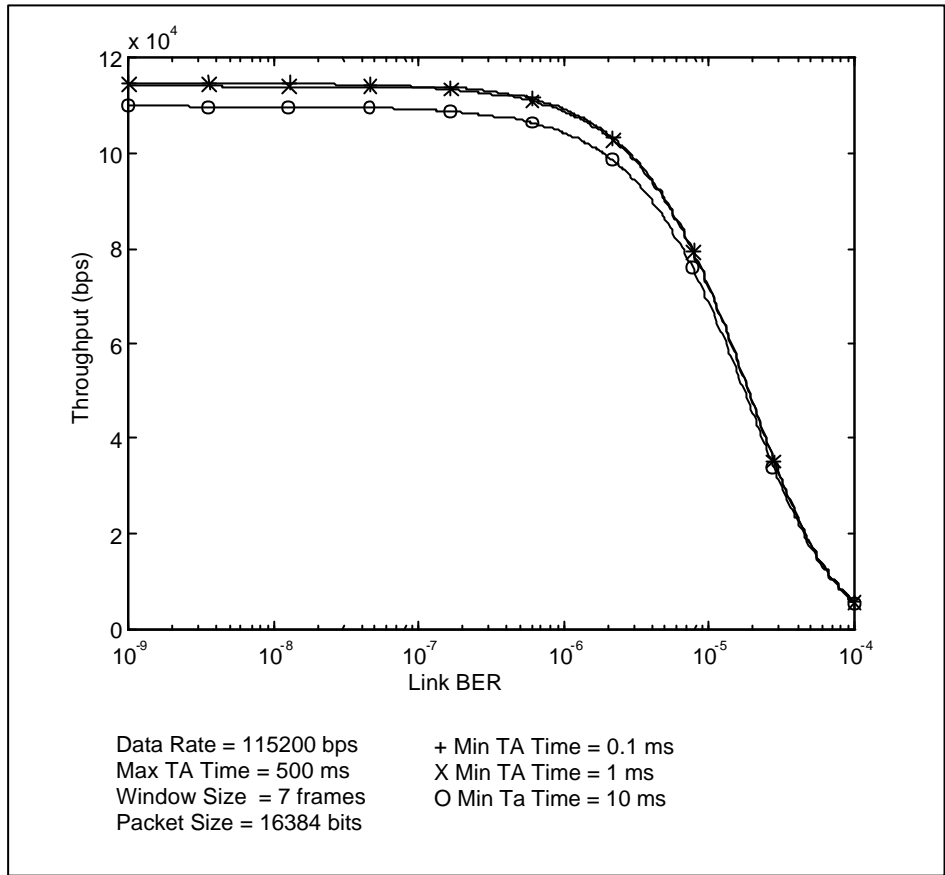





Figure 5 : throughput Vs BER for 115200 bps



a) Error free transmission at $w = 3$ and $w = N$



b) Retransmitted frames due to error frame at $w = 4$



c) Retransmitted frames and F-timer delay due frame error at $w = 4$ and $w = 1$

Figure 4: Determination of frame virtual transfer times




Figure 3: state transition diagram for the widow width process

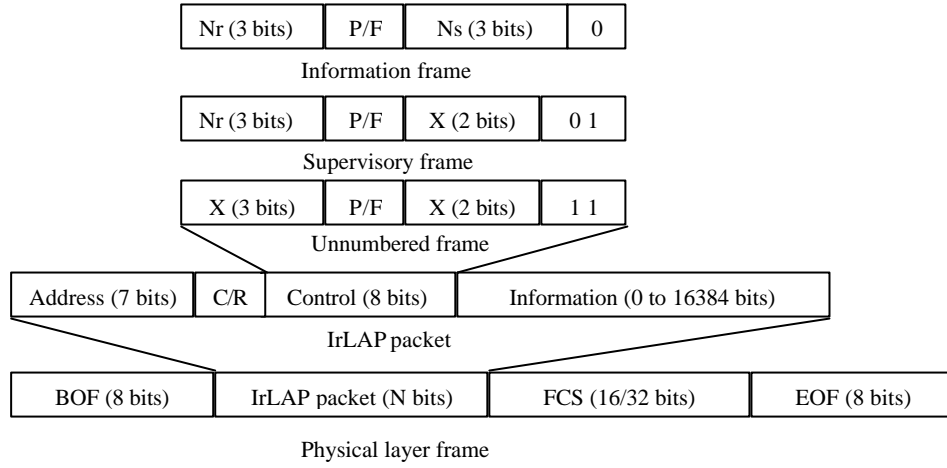


Figure 2: IrDA frame structure

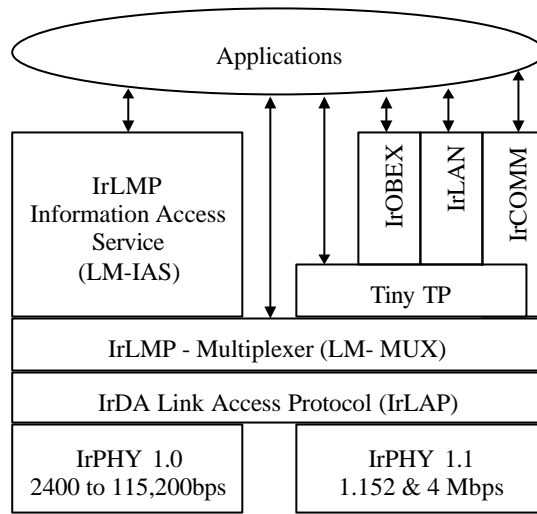


Figure 2: The IrDA SIR protocol stack

List of Figures

Figure 1: The IrDA SIR protocol stack

Figure 2: IrDA frame structure

Figure 3: state transition diagram for the widow width process

Figure 4: Determination of frame virtual transfer times

Figure 5 : throughput Vs BER for 115200 bps

Figure 6: throughput Vs packet data size for 115200 bps

Figure 7: throughput Vs BER for 4 Mbps

Figure 8: throughput Vs packet data size for 4 Mbps

Figure 9: throughput Vs BER for 16 Mbps

Figure 10: throughput Vs maximum turn-around time for 16 Mbps

- [11] P. Barker, A. C. Boucouvalas, 'Performance modeling of the IrDA protocol for Infrared Wireless Communications', *IEEE Communications magazine*, **36** (12), pp.113-117, (1998)
- [12] W. Bux & K. Kummerle, 'Balanced HDLC Procedures: A Performance Analysis'. *IEEE Transactions On Communications*, **com-28**(11), pp 1889-1898. (1980).
- [13] M. Schwartz, *Telecommunications Networks: Protocols, Modelling and Analysis*. Reading, MA: Addison Wesley. (1987).

5 Conclusions

The mathematical analysis presented shows that the concept of the 'virtual transmission time' which had been previously used to analyse full-duplex 'wired' protocols can be applied to the IrLAP protocol, being a derivative of the HDLC protocol in NRM mode. The model examines the time required for retransmission of packets and the probabilities that such retransmissions are required. Results are presented for throughput efficiency versus data message BER and other link parameter settings. The results indicate possible limitations of the present IrDA protocol in supporting higher bit rates (e.g. 10 Mbps and above) as the link turn-around frequency can severely impair the link performance. The larger window size proposed for the 16 Mbps link provides a solution but only if the link BER is optimised. If the BER becomes poor, the larger window size will impair the performance as more frame retransmissions will be required. The model presented can thus be useful in determining the link parameter settings for optimum performance of an IrDA infrared link in a particular link environment.

Acknowledgements

The authors would like to thank BT laboratories for their financial assistance in carrying out this research.

References

- [1] D. J. T. Heatly, D. R. Wisely, I. Neild, & P. Cochrane, 'Optical Wireless: The Story So Far'. *IEEE Communications Magazine*, **36**(12), 72-82 (1998).
- [2] A. C. Boucouvalas, 'Free space infrared optical links: Eye safety issues', *Proceedings of MEDIACOMM 95 - International Conference on Multimedia Communications*, Southampton, UK, 11-12 April 1995, pp. 205-209. (1995)
- [3] A. C. Boucouvalas, 'Eye safety Issues of Free Space Optical Links'. *Focus*, 23..pp 8 –23, (1995)
- [4] A. C. Boucouvalas, 'IEC 825-1 Eye Safety Classification of some consumer electronic products' *IEE Colloquium on Free Space Communication Links*, paper No. 13, 19th February 1996, IEE Savoy Place, London. (1996)
- [5] A. M. Street, K. Samaras, D. C. O'Brian, & D. J. Edwards, 'Indoor Optical Wireless Systems - A Review'. *Optical and Quantum Electronics*, 29, pp 349-378. (1997).
- [6] I. Millar, M. Beale, B. J. Donoghue, K. W. Lindstrom, & S. Williams, 'The IrDA Standard for High-Speed Infrared Communications'. *The Hewlett-Packard Journal*, February 1998.**49**(1) pp 10 -26, (1998)
- [7] IrDA, *Serial Infrared Physical Layer Link Specification – Version 1.1*, Infrared Data Association, (1995)
- [8] IrDA, *Serial Infrared Link Access Protocol (IrLAP) – Version 1.1*, Infrared Data Association, (1996)
- [9] D.W. Suvak, *Proposed Changes to IrLAP for VFIR*, Infrared Data Association, (1999)
- [10] A. C. Boucouvalas, P. Barker, 'IrLAP protocol performance analysis of IrDA wireless communications' *Electronics letters*, **34** (25), pp.2380-2381, (1998).

4.1 Analysis of 115200 Kbps and 4 Mbps links

Figure 5 shows the throughput versus BER for the 115200 bps link. This shows that a link BER of 10^{-7} or better is required for optimum link quality, and that a BER worse than 10^{-6} would be quite detrimental to the throughput. It also shows that at this data rate, the effect of increasing the minimum turn-around time is not very significant. Figure 6 shows throughput versus packet data size for two values of BER: 10^{-8} and 10^{-5} , therefore representing links of 'good' and 'bad' quality respectively. The plot shows that the throughput increases with increasing packet data size. This is because as the size of the frame increases, the frequency of the link turn-around required after every 7 frames decreases, thus decreasing the frequency of acknowledgement times and turn-around delays introduced into the throughput. It can also be seen that the effect of the increase in minimum turn-around time is greater for a lower data size because of the higher link turn-around frequency. The 'jerky' nature observed at the high end of the plot is due to the quantised reduction in frames that can be transmitted within the 500 ms maximum turn-around time, as the frame size increases. In fact at the 115200 bps data rate, the number of frames sent within 500 ms at the maximum data size of 16384 bits is 3 and not the maximum 7. For the BER of 10^{-5} plots it can be seen how the throughput increases with packet data size as before until around 10^3 bits after which it decreases. This is because as the packet data size increases, the probability of the frame being in error, given the same BER, also increases and begins to dominate over the improvement due to increasing data size after 10^3 bits.

Figures 7 and 8 show the same analysis for the 4 Mbps link. Figure 7 shows that a BER of nearer 10^{-8} as compared to 10^{-7} for the 115200 link is required for optimum performance. It can also be seen that the effect of increasing the minimum turn-around time is much more significant. This is because at this higher data rate, the frequency of link turn-around (required after every 7 frames sent) is much higher. Figure 8 shows throughput versus packet data size for BER values of 10^{-8} and 10^{-5} . This again shows the increased significance of the minimum turn-around time. It can also be seen how at the higher data rate a potentially larger package size than the maximum 2048 bytes could be sent. However it can also be seen that a larger packet size with a 'poor' BER will result in a higher probability of frame errors.

4.2 Analysis of 16 Mbps links

The proposed 16 Mbps link, as well as providing a higher data rate, also presents the issue of a maximum window size, of 127 frames. It can be seen from figures 9 that the effect of the increase in the minimum turn-around time is not as great as would be expected with the increase in data rate because the larger window size results in a reduced rate of link turn-around frequency. However it can also be seen that the increase in window size results in a greater sensitivity to the BER as a larger window size means a higher probability of errors within the window and a larger average number of frames that will require retransmission. This is further demonstrated by figures 10 where the maximum turn-around time is varied for the 'good' and 'bad' BER values. A similar effect as for the 4 Mbps link can be seen but the degradation in performance for a bad BER is much more severe.

$$P_n = p^n \quad n \geq 2 \quad (18)$$

and the value of the transmission time is given by:

$$E[T_n(w)] = T_2 \quad n \geq 2 \quad (19)$$

where T_2 is the time for each successive retransmission of the frame and is now independent of the window width w , given by:

$$T_2 = \begin{cases} Nt_I + p(t_{Fout} + t_S) + t_{ack} & N \neq 1 \\ t_I + t_{Fout} + t_S + t_{ack} & N = 1 \end{cases} \quad (20)$$

3.3 Overall window specific virtual transmission time

The overall virtual transmission time for a specific value of w can therefore be determined by using the expression in equation (12) and applying the derived expressions for the expectations of the individual cases as shown above, combined with the stated probabilities of each case arising. By using,

$$\sum_{n=2}^{\infty} p^n = \frac{p^2}{(1-p)} \quad (21)$$

it can be shown that.

$$t_v(w) = T_0(w) + pE[T_1(w)] + \frac{p^2}{(1-p)}T_2 \quad (22)$$

The equations (22) and (11) provide the components for equation (9) to determine the packets virtual transmission time and thus provide a mechanism to determine the average throughput of a link against any of the link parameter settings of BER, data rate, packet data size, minimum turn-around time, maximum turn-around time, and maximum window size. Calculations could also be made to determine the link BER from the link topology and transceiver characteristics, but that is outside the scope of this paper.

4 Results and analysis

The equations derived are used to plot the throughput of IrDA links using a range of system settings. Plots are produced for three data rates: namely 115200 bps, 4 Mbps and 16 Mbps, being the maximum data rates of the SIR, FIR and proposed VFIR standards respectively.

This case examines the situation where a frame transmitted needs to be retransmitted because it has been in error. This case has the probability of a single error, i.e. p . The time $T_1(w)$ denotes the time taken from the end of the initial transmission to the end of the 1st retransmission. We need to examine whether we have additional errors following the initial error frame, and whether this includes the last frame in the sequence.

i) Error not last in series

Here the frame is simply retransmitted at the end of the sequence and after the REJ has been returned. This is shown in Figure 4b.

$$T_1(w) = wt_I + t_{ack} \quad (14)$$

ii) Errors including last in series.

In this case the P-bit is lost, therefore the secondary will not know that it can now transmit. If the primary has not received a response when the F-timer has expired a P bit will be sent with an RR S-frame. This situation has an additional probability of p , i.e. the probability of a specific frame (the last) following the first corrupted frame being in error. This is shown Figure 4c.

$$T_1(w) = wt_I + t_{Fout} + t_S + t_{ack} \quad (15)$$

(iii) Single error being the last in the series

This case is the same as case (ii) above with $w = 1$ except that there is only one error in the sequence, thus there is only the single error probability.

Therefore by combining equations (14) and (15), with the stated additional probabilities the overall expectation value of $T_1(w)$ and the probability P_1 can be given by:

$$P_1 = p \quad (16)$$

$$E[T_1(w)] = \begin{cases} wt_I + p(t_{Fout} + t_S) + t_{ack} & w \neq 1 \\ t_I + t_{Fout} + t_S + t_{ack} & w = 1 \end{cases} \quad (17)$$

C. Secondary retransmissions due to subsequent errors of the same frame

As the initial resent frame will be the first in the series, all N frames will be retransmitted. Therefore the retransmitted frame will always be the first in the sequence. However as before if the last frame of the sequence is also in error, the P bit will not be received so the recovery only begins when the F-timer expires. The probability of n successive transmissions of the same frame being in error is p^n . Thus we have:

$$Q = RQ \quad (10)$$

which can be shown to give

$$f(w) = \begin{cases} \frac{p(1-p)^{N-w-1}}{1-(1-p)^N} & w = 1, 2, \dots, N-1 \\ \frac{p(1-p)^{N-1}}{1-(1-p)^N} & w = N \end{cases} \quad (11)$$

The virtual transmission time for a particular window width w (for equation 9) is calculated by the addition of the time to transmit the initial frame and the time required for subsequent retransmissions of the frame multiplied by the probabilities that these retransmissions are required.

$$t_v(w) = T_0(w) + \sum_{n=1}^{\infty} P_n E[T_n(w)] \quad (12)$$

where T_0 is the time to transmit the frame initially P_n is the probability that the n_{th} retransmission is required, and $E[T_n(w)]$ is the time expectation value of the time required for the n_{th} retransmission only.

3.2 Calculation of Virtual Transmission Times

In calculating the virtual transmission time we need to examine a number of different cases involving whether an error occurred, in what position in the sequence the error is located, whether there are additional errors in the sequence and in what positions, and whether the retransmissions of frames are themselves in error.

A. Initial frame transmission time

The transmission time in this situation is simply the transmission time for the I-frame unless the next frame to transmit is the first frame of a sequence (i.e. $w = N$) in which case the primary station must wait to receive an acknowledgement from the receiving station before continuing to transmit. This is shown in Figure 4a.

$$T_0(w) = \begin{cases} t_I & w \neq N \\ t_I + t_{ack} & w = N \end{cases} \quad (13)$$

B. Retransmission of a frame in error

The virtual transmission time for a particular frame is found to be dependent on the window width (w) seen by that frame. The window width is defined as the number of remaining frames (from the start of the current frame) that can be sent before transmission control is passed. With IrDA, the maximum window width is determined by the ‘maximum window size’ parameter, here denoted W_{max} , negotiated during the link establishment, and the maximum turn-around time combined with the data rate and packet data size which may limit the number of frames that can be sent to below W_{max} . We therefore use a value N to indicate the number of frames that can be sent before transmission is passed to the other station.

$$N = \min \left\{ W_{max}, \text{floor} \left(\frac{T_{max}}{t_I} \right) \right\} \quad (8)$$

where \min is ‘the lesser of’, and floor is ‘the largest integer not exceeding’. The overall virtual transmission time is determined by summing the frames virtual transmission time for all values of w , multiplied by the probability of the frame having each value of w . This is therefore given by:

$$t_v = \sum_{w=1}^N f(w) t_v(w) \quad (9)$$

where $t_v(w)$ is the virtual transmission time for a frame seeing a width w , and $f(w)$ is the probability that the window width has a specific value of w .

Each time a frame is transmitted error free (with probability $1-p$), the window width reduces by 1. When it reaches 1, the window width returns to its maximum value of N . However if a frame is in error (with probability p) then it and all the following frames in the sequence will need to be retransmitted. Therefore the window width effectively returns to $(N-1)$. The overall process can be considered as a Markov chain and represented as the state transition diagram shown in Figure 3. The value for $f(w)$ can then be determined by calculating the ‘limiting (i.e. steady-state) probability distribution’ for the window width state set. The state transition diagram can be represented by the state transition matrix R , shown below, where the probability of going from a state i to a state j is given by the element r_{ij} .

$$R = \begin{array}{c|ccccccc} & (N-1) & (N-2) & \dots & (w) & \dots & (1) & (N) \\ \hline (N-1) & p & 1-p & 0 & 0 & 0 & 0 & 0 \\ (N-2) & p & 0 & \ddots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & 0 & 0 & 1-p & 0 & 0 & 0 \\ (w) & p & 0 & 0 & 0 & \ddots & 0 & 0 \\ \vdots & \vdots & 0 & 0 & 0 & 0 & 1-p & 0 \\ (1) & p & 0 & 0 & 0 & 0 & 0 & 1-p \\ (N) & 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{array}$$

If we define the limiting probability distribution vector Q such that $Q = [f(1) \dots f(w) \dots f(N)]$, using matrix multiplication we can determine Q by using:

The following parameters in Table 1 are used in the analysis.

Parameter	Description	Unit
C	Link data baud rate	bits /sec
p_b	Link bit error rate	-
p	Frame error probability	-
l	I-frame message data length	bits
l'	S-frame length / I-frame overhead	bits
t_I	Transmission time of an I-frame	sec
t_S	Transmission time of an S-frame	sec
t_{ta}	Minimum turn-around time	sec
t_{ack}	Acknowledgement time	sec
T_{max}	Maximum turn-around time	sec
t_{Fout}	F-timer Time-out period	sec

Table 1: Analysis Parameters

The values for t_S , t_I and t_{ack} are calculated as follows:

$$t_S = \frac{l'}{C} \quad (1) \quad t_I = \frac{l+l'}{C} \quad (2) \quad t_{ack} = 2t_{ta} + t_S \quad (3)$$

The acknowledgement time t_{ack} here being the time to receive back an S-frame acknowledgement after the last frame in the sequence is transmitted, thus requiring a turn-around delay at the secondary, the transmission time of the S-frame and a turn-around delay at the primary. The frame error probability p is related to the link bit error rate p_b by:

$$p = 1 - (1 - p_b)^{l+l'} \quad (4)$$

This is the probability that there is at least one bit in an I-frame in error.

3.1 Throughput calculation

The frame throughput D_f in frames per second of the link is given by the reciprocal of the virtual transmission time. This can be multiplied by the frame data length l to give the data throughput in D_b in bits per second.

$$D_f = \frac{1}{t_v} \quad (5) \quad D_b = \frac{l}{t_v} \quad (6) \quad \text{Normalised Throughput } D_{Norm} = \frac{D_b}{C} \quad (7)$$

lost frame or frames (i.e. failed the CRC and were discarded). The value of V_r then remains unchanged and the packet is discarded. Upon replying with an acknowledgement frame the secondary sets the N_r field to the current V_r value. Upon receiving a frame from the secondary with the F bit set (thus returning transmission control to the primary) the primary checks the N_r value against that expected, which should be equal to the current V_s value. If it is not as expected, this indicates that a retransmission of the buffered frames beginning at the indicated N_r value is now required before any new frames can be transmitted. The secondary station operates in the same way as the primary except that it does not contain a P-timer or F-timer. The secondary will immediately return transmission control to the primary if data is not immediately available for transmission [8].

A proposal has been made by IrDA to increase the maximum data rate of an IrDA link from 4 Mbps to 16 Mbps. It was recognised that at this data rate the increased frequency of link turn-around after the maximum 7 frame window size would limit the performance of the link. It has therefore also been proposed to increase the maximum window size from 7 frames to 127 frames. This requires a restructuring of the IrLAP frame control field to constitute 2 bytes in length and to have 7 bit fields for N_s and N_r [9], [10], [11].

3. Mathematical modelling

The mathematical performance model uses the concept of the “virtual transmission time”, denoted t_v , of a packet in determining the throughput efficiency of a link. The value of t_v provides the average time taken from the start of a frame's transmission to the successful arrival of the frame in-sequence at the receiving station, incorporating the time for retransmission of frame errors, and other enforced delays. If we assume continuous data transmission, then the average throughput (frames /sec) is the reciprocal of the virtual transmission time, thus representing the arrival rate of packets at the receiving application above the IrLAP layer. This method was used in an analysis of the HDLC protocol operating in ABM with a full-duplex circuit by Bux [12]. A review of the method and applications to simpler systems was also given in Schwartz [13].

The model used here involves a primary and a single secondary station operating in the data exchange phase of the IrLAP procedure. All user data (i.e. I-frames) is transmitted from the primary to the secondary only. The link is assumed to be saturated, i.e. there is data continually waiting transmission. The secondary responds with RR (Receive Ready) or REJ (Reject) S-frames only. It is also assumed that an S-frame, whether from the primary or secondary, is small enough to be considered error free. The action of the P-timer is ignored in this model as we are assuming a saturated condition. However the action of the F-timer is incorporated when the primary forces a reply from the secondary if there is corruption of a P bit frame.

Upon link establishment devices are assigned primary and secondary station roles. There can be only one primary but potentially multiple secondaries. All data traffic must pass to or from the primary, i.e. the secondary stations cannot communicate between themselves.

The IrDA frame structure is as shown in figure 2. There are three types of frame used; Supervisory frames (S-frames), Unnumbered frames (U-frames), and Information frames (I-frames). S-frames are short control frames used to issue specific commands and responses (e.g. to acknowledge received frames) and contain no user data. U-frames are used in the establishment of a link to exchange connection data or for information exchange outside of an established link. I-frames are used for the exchange of information once a link has been established. They contain the user data in the information field and have sequencing numbers for error recovery.

The IrLAP packet consists of an address field, a C/R (Command /Response) bit, an 8 bit control field, and the information field which can contain up to 2048 bytes of data. The control field contains a frame type identifier, the P/F bit, and frame sequence numbers and /or command /response codes depending on the frame type. The P/F bit is used to pass transmission status between devices. For the primary device it is called the P bit (Poll) and for the secondary it is called the F bit (Final). The I-frame control field contains a send sequence number N_s which cycles through values 0 to 7. Both the I-frame and the S-frame control fields contain a receive sequence number N_r which acknowledges the receipt of frames with sequence number up to $(N_r - 1)$, thus indicating that N_r is the next expected frame sequence number. The IrLAP packet is encapsulated within the physical layer frame before transmission. The physical layer frame contains a BOF (Beginning Of Frame) and an EOF (End Of Frame) framer which both consist of the identifier '01111110' and the FCS (Frame Check Sequence) which contains the result of the CRC.

The IrLAP information exchange procedure is summarised as follows. Each device in a link has a send sequence variable V_s , a receive sequence variable V_r , and a window value indicating the number of frames remaining to be transmitted before transmission control is passed. In addition the primary device holds a P-Timer, which sets a limit to the transmission period of the primary device, and an F-Timer which sets a limit on the waiting time for the secondary to return transmission control. If the primary device has a data packet to send, the N_s value of the frame is assigned the current V_s value of the device and a buffer copy of the frame is made for possible retransmissions. Following transmission the value of V_s is then increased by 1 (modulo 8) and the window value decreased by 1. If the window value becomes 1, this indicates that the next frame is the last in the sequence before passing transmission control to the secondary device. This frame therefore has the P bit set to poll the secondary device. Following transmission of this frame the window value is re-set to the maximum and the F-Timer started. If the P-Timer should expire during the transmission control period of the primary due to sporadic or limited data, the primary will send an S-frame with the P bit set to poll the secondary. If the F-timer should expire following transmission of a P bit frame, this indicates that either the P bit frame has been lost, or the returning acknowledgement frame has been lost. In either case the primary must transmit an S-frame to re-poll the secondary.

Upon receiving a data frame from the primary, the secondary device checks the received N_s value against the expected value. N_s should be equal to the current V_r value of the secondary. If N_s is as expected, the data packet is extracted and passed to the user. The stored value of V_r is then increased by 1 (modulo 8). If N_s is not as expected, this indicates a

application to use the IR link, IrLAN which provides network access via a wireless infrared LAN port, and IrOBEX which provides object exchange facilities based on the HTTP protocol.

2.1 The physical layer

The IrDA physical layer provides the hardware for an IrDA infrared link, including optical transmitter and detector, and modulation and encoding hardware. The CRC (Cyclic Redundancy Check) and IrLAP packet framing are also considered to be part of the physical layer although often implemented in software. Data rates up to and including 1.152 Mbps use a 16 bit CRC, while the 4 Mbps employs a 32 bit CRC. Frames that fail the CRC, thus indicating bit errors, are simply discarded and not passed to the upper IrLAP layer. The standard specifies a link distance of at least 1m with a transmission cone half angle between 15° and 30° and a receiver cone half angle of at least 15°. The IR transmitter power output is limited by eye safety regulations [2]-[4]. The specified maximum output power is 500 mW/Sr but typical implemented values are around 100 to 200 mW/Sr. Version 1.0 of the physical layer provides a data baud rate of up to 115.2 Kbps. This uses an RZ modulation scheme with a 3/16 pulse duration. The hardware connects to the devices serial UART interface, either internally or using an external 'dongle'. The baud rate is thus limited to the maximum transmission rate of the devices UART chip. Version 1.1 of the physical layer provides a baud rate of up to 1.152 Mbps or 4 Mbps. Rates up to 1.152 Mbps use the same RZ encoding of the 1.0 standard but implement the CRC and framing function in hardware and connects to the devices data I/O interface. For the 4 Mbps data rate, a 4 PPM modulation scheme is used and also requires additional PLL circuitry for clock synchronisation [7].

2.2 The IrLAP layer

The IrLAP layer is the data link layer of the IrDA protocol stack and is based closely on the HDLC protocol operating in Normal Response Mode (NRM). The functions of the IrLAP layer include device discovery, link establishment, data exchange control and error recovery. Before a link can be established a device must first 'discover' the other devices within range that it can connect to. The user can then establish a link to one of these devices. During link establishment, a number of link parameter values are negotiated. These are given below. All negotiation traffic occurs at 9600 bps.

Data Rate: Both devices in a connection must negotiate the same data rate. This can be up to 115.2 Kbps for a the version 1.0 physical layer, or 4 Mbps with the version 1.1 physical layer.

Maximum turn-around time: This is the maximum time that a station can hold transmission control and has an upper limit of 500 ms.

Minimum turn-around time: This is a delay that must be implemented each time a device regains transmission control and is used to cover receiver latency. This has values from 0.01 ms to 10 ms.

Maximum window size: This is the maximum number of frames that can be sent before the passing of transmission control to another device. This can have values 1 to 7. The maximum turn-around time combined with packet data size and data rate have priority over the maximum window size and thus can limit its value to below that requested.

Packet data size: This is the maximum size of the information field of the IrLAP frame which contains the user data and control overheads from higher layers of the protocol. This has a maximum value of 2048 bytes (16384 bits).

1. Introduction

Mobile and portable computers have seen a rapid growth in recent years which has led to an increasing demand for wireless data connectivity [1]. The infrared (IR) optical medium has the benefits of using lightweight, cheap and readily available components, of being restricted to the room of operation, of being spectrally unregulated and of providing potentially very high data rates. However optical power output is limited by eye-safety regulations and a desire to limit power consumption [2]-[4]. Wireless IR is therefore particularly suited to short range indoor applications [5]. The Infrared Data Association was created in 1993 to establish an open standard for short range IR data communication. The resulting IrDA SIR (Serial Infrared) protocol standard was developed to provide a simple, low cost and reliable means of data communication between IrDA compliant devices using point-to-point half-duplex IR links. Version 1.0 of the protocol provides a data rate of up to 115.2 Kbps using a connection to the standard serial port interface, while version 1.1 provides a rate up to 1.152 Mbps and 4 Mbps with high speed extension hardware [6]. The higher speed extension is often referred to as IrDA FIR (Fast Infrared). IrDA have also proposed an extension to 16 Mbps, referred to as VFIR (Very Fast Infrared).

The data link layer of the protocol, IrLAP (IrDA Link Access Protocol), is based on the HDLC protocol operating in NRM (Normal Response Mode) which specifies primary and secondary station roles with all data traffic passing to or from a single primary station. This uses station polling and frame sequence numbering to indicate lost frames and force retransmissions of lost and resulting 'out of sequence' frames. The model presented here examines the performance of an IrDA link at the IrLAP level using the concept of a frame's virtual transmission time. This indicates the time from the start of an initial frame transmission to the successful receipt of an in-sequence frame by the receiving station, thus incorporate the time for retransmissions and other frame delays using the probability of frame errors based on a specified link bit-error-rate (BER). The performance is examined in relation to the BER and other link parameter settings including link baud rate, minimum turn-around time, maximum turn-around time, maximum window size and packet data size. The model can thus be beneficial in determining the optimum parameter settings for a particular link and media environment.

2. The IrDA protocol stack

The protocol stack of IrDA is shown in figure 1. There are three mandatory layers in the stack; the physical layer, the IrLAP layer and the IrLMP (IrDA Link Management Protocol) layer. The later consists of two components; the LM-MUX (Link Management Multiplexer) which provides access for multiple entities or application channels to independently use a single established IrLAP link, and the LM-IAS (Link Management Information Access Services) which provides a database of supported services which can be interrogated by a connecting device. Optional protocol components include Tiny TP, a lightweight transport component that provides credit based flow control and segmentation and reassemble functions, IrCOMM a serial cable emulation component which permits native serial

Performance Modelling of the IrDA Infrared Wireless Communications Protocol

P. Barker, A.C. Boucouvalas, V. Vitsas
Multimedia Communications Group
School of Design, Engineering and Computing
Bournemouth University
Talbot Campus, Fern Barrow,
Poole BH12 5BB, UK
Email: {pbarker, tboucouv, vvitsas}@bournemouth.ac.uk

Revised September 2000

Abstract

We present a performance analysis of the IrDA (Infrared Data Association) IR data communications protocol using a mathematical throughput model based on the virtual transmission time of uni-directional IrDA data packets. The IrDA standard specifies a protocol stack for reliable short-range indoor IR wireless data communication between devices with the data link layer (IrLAP) being a HDLC derivative. The virtual transmission time provides the average end-to-end transmission time of a packet incorporating the time for retransmissions if packet errors occur or other enforced delays. Analytic results are produced for the 115.2kbit/s 4 Mbit/s and 16 Mbit/s versions. Results are given for throughput with varying link parameters of packet size, BER, data rate, minimum turn-around time and maximum turn-around time.