Modeling IrDA Performance: The Effect of IrLAP Negotiation Parameters on Throughput

Michael G. Robertson, Scott V. Hansen, Franklin E. Sorenson, Charles D. Knutson
{michaelr, scott, sorenson, knutson}@cs.byu.edu

Computer Science Department
Brigham Young University
Provo, UT 84602, U.S.A.

ABSTRACT

The Infrared Data Association's (IrDA) infrared data transmission protocol is a widely used mechanism for short-range wireless data communications. In order to provide flexibility for connections between devices of potentially disparate capabilities, IrDA devices negotiate the values of several transmission parameters based on the capabilities of the devices establishing the connection. This paper describes the design and implementation of a software tool, `Irdaperf`, to model IrDA performance based on negotiated transmission parameters. Using `Irdaperf`, we demonstrate that for fast data rates, maximizing window size and data size are key factors for overcoming the negative effects of a relatively long link turnaround time. At slower speeds (especially 115.2 Mbps and below), these factors have a less pronounced effect.

KEYWORDS
IrDA, Infrared Data Communications, Performance Analysis, Wireless Data Communications

1 INTRODUCTION

IrDA infrared devices are widely used for short-range wireless data communication. They require little power and can achieve high transmission rates at close range. IrDA devices are small and robust, which makes them ideal in portable, mobile electronic devices, such as laptop computers, cell phones, and personal digital assistants (PDAs). Despite the simplicity of their design and implementation, IrDA devices can transmit data at up to 4 and 16 Mbps.

In order to make data transfers as efficient as possible, IrDA devices initiate transmission by negotiating certain parameters. Some parameters (such as data rate) must be agreed upon by the two devices, while other parameters (such as packet size) must be respected by the other device. This allows more powerful devices to interoperate with smaller, less powerful ones. Depending on the negotiated values, throughput for a given file transfer can vary greatly. In this paper, we examine the IrDA negotiation parameters in order to determine which have the greatest impact on throughput and to specify the values for those parameters that maximize performance.

We have developed an analytical model of the IrDA protocol and have implemented that model in an application named `Irdaperf`. We have verified this model and the corresponding application with empirical data obtained from testing infrared-enabled devices. Using `Irdaperf`, we simulated IrDA data transfers using different values for negotiation parameters and compared the

resulting throughput, allowing us to determine the parameters that affect throughput most significantly and the values that are ideal for those parameters.

1.2 Related Work

The most relevant published research in the field of IrDA performance was conducted by Barker and Boucouvalas in [2] and [3]. They created a mathematical model with which they could calculate IrDA performance based on certain parameters. They concluded that the best performance is achieved with a small minimum turnaround time and large window size. Ozugar, et. al. investigated the IrDA protocol in [5] and concluded that the device's processor speed determines the parameters that have the biggest impact on performance. Other papers, such as [1], provide a summary or history of the protocol or a vision of its future.

1.3 IrDA Negotiation Parameters

The IrLAP (Infrared Link Access Protocol) layer of the IrDA protocol specifies seven parameters that infrared devices must negotiate before data transfer may commence [4]. The parameters govern the size of the packets, the speed at which they are sent, and the timing of their transmission. The negotiation process frees infrared devices from having to support all possible configurations; instead, two devices may choose the best set of parameters that are mutually supported.

Negotiation parameters are divided into two groups: type 0 and type 1. Two devices must agree on the same value for type 1 parameters but may use different values for type 0 parameters, provided that they respect the other's chosen value. Baud rate and link disconnect/threshold time are the only type 0

parameters. The type 1 parameters include minimum and maximum turnaround time, data and window size, and the number of additional beginning of frame (XBOF) bytes.

2 RESEARCH

This study describes a model of the IrDA protocol stack for use in performance analysis. Our research focuses on the IrLAP layer with data transfer rates between 576 Kbps and 4 Mbps. We have excluded slower transmission speeds because varying negotiation parameters has minimal effect on throughput at those speeds. We have also excluded the 16 Mbps VFIR specification as it is not yet widely available in IrDA devices. Our model provides a means of analyzing the effect of negotiation parameters on throughput and thereby determining those that play the most significant role in improving performance. Using our model, we have discovered several scenarios in which selecting appropriate negotiation parameter values may compensate for other deficiencies that diminish throughput.

2.1 Model

In an IrDA data transfer, a block of data is first passed from the application initiating the transfer to the top level of the IrDA protocol stack. This top layer may be one of several application-layer protocols, depending on the particular usage model being employed. For point and shoot object exchange [7], the appropriate top layer is the IrDA Object Exchange protocol (IrOBEX), which sits on top of the IrLMP layer. For other applications the top layer may be some other high-level protocol where that protocol is the appropriate layer for the chosen usage model. For the purpose of this study, we have selected the point and shoot object

exchange usage model since it is the most commonly used among IrDA devices. As a consequence, our study assumes that applications write directly to IrOBEX as the top layer of the IrDA protocol stack.

Our model focuses on two metrics: the number of bytes transmitted and the transfer time. Transfer time is calculated by determining the time required to send bytes at the Physical layer and the time spent waiting for the receiving device to acknowledge received data according to the IrLAP protocol. The number of bytes represents the data to be sent plus the header information added by the protocol layers. Our model accounts for the additional headers added by the upper layers but ignores them when computing transfer time. As our goal is to predict throughput, we focus on only the data transfer time, not the time required for an entire IrDA transmission (including device discovery, link creation, etc.).

2.2 Results

From a performance standpoint, minimum turnaround time is among the most important of the IrDA negotiation parameters. Boucouvalas and Baker showed in [2] and [3] that turnaround time should be kept as low as possible during IrDA data transfers to achieve the greatest throughput. Furthermore, our research indicates that the negative effect that a high minimum turnaround time has on throughput increases as the data transfer rate increases. To overcome this loss in throughput, it is necessary that data and window sizes be made as large as possible. We will quantify the throughput loss due to a high minimum

turnaround time and show the benefit of using large window and data sizes. Other IrDA negotiation parameters have less significant effect on throughput.

2.2.1 Minimum Turn Around and Data Rate

While a longer minimum turnaround time always degrades throughput, the effect is more pronounced at higher data transfer rates. Figure 1 shows the percentage of maximum throughput for various data rates and minimum turnaround times. For data rates up to 115.2 Kbps, near maximum throughput is achieved for all minimum turnaround times. At higher transfer rates, however, throughput drops dramatically when the minimum turnaround time is greater than 1 ms. At higher data rates, more data can be sent per unit time. This means that time spent idle represents a greater proportional loss in throughput at higher data rates than at slower ones.

As an example, at 9600 bps, 96 bits can be transmitted in 10 ms. At 4 Mbps, however, 40,000 bits can be transmitted in the same amount of time. Likewise, a 9600 bps device requires 6.7 ms to transmit a 64 byte frame while a 4 Mbps device requires only 16 microseconds. Assuming a 5 ms turnaround time, the 9600 bps device will spend approximately 43% of its time idle. The 4 Mbps device, however, will spend over 99% of its time idle. Thus, an infrared device capable of transmitting at 4 Mbps will suffer a greater proportional loss in throughput if it is forced to wait 5 ms between packets than will a 9600 bps device.
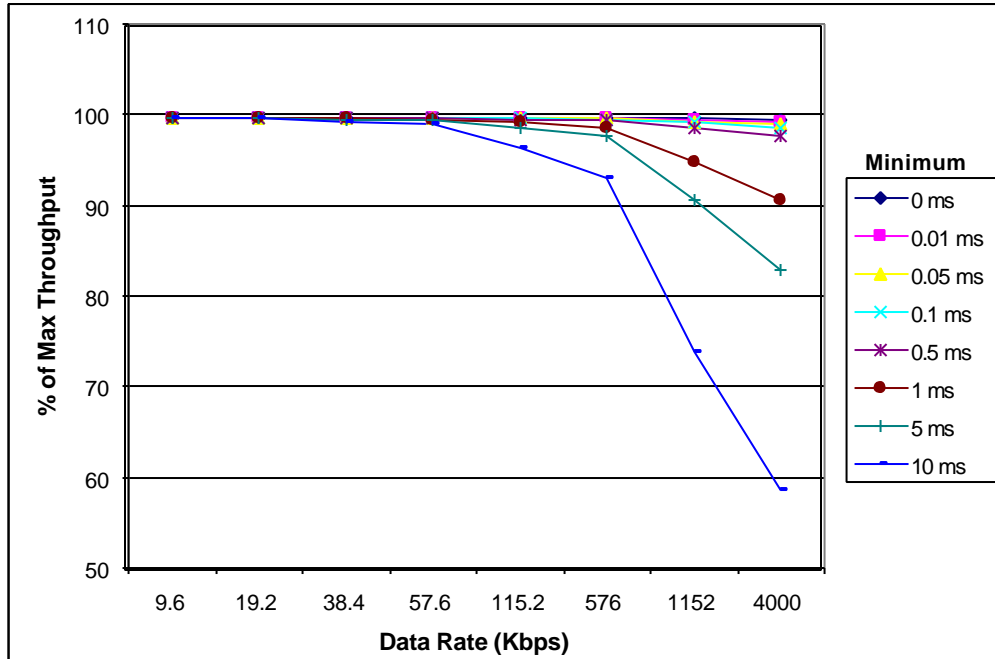
Figure 1: Effect of Minimum Turnaround Time on Throughput at Different Data Rates

2.2.2 Compensating for High Minimum Turn Around

At higher data rates, minimizing the length of the turnaround period between data transmission is essential for high throughput. If a device requires a long minimum turnaround time, it is best to reduce the number of times link turnaround must occur. This can be accomplished by using large windows or frames, which ensure that as much data as possible is sent before the sender must wait for an acknowledgement.

When the data and window sizes are small, throughput decreases significantly as the turnaround time increases. Reducing data and window sizes helps minimize the loss in throughput. This happens for two reasons. First, as the window size increases, the number of times the link must be turned around decreases, enabling the sending device to spend a higher percentage of its time transmitting data. Second, sending data

in larger frames or windows reduces header overhead.

Figures 2 and 3 show throughput for data transfers at 1.152 Mbps and 4.0 Mbps, respectively, with various window sizes and minimum turnaround times. In both charts, when the window size is small, throughput drops as the minimum turnaround time increases. When a large window size is used, however, throughput decreases much more slowly as the minimum turnaround time approaches 10 ms. Furthermore, the drop in throughput is much more pronounced at higher data rates. For example, at 1.152 Mbps, with a window size of 1, throughput with the longest minimum turnaround time is roughly 58% of the throughput with the shortest minimum turnaround time. At 4 Mbps, however, the percentage drops to 17%. When a window size of 7 is used, the percentages are 91% and 59%. Although throughput always drops as the data rate increases, the decrease is less when the larger window size is used.
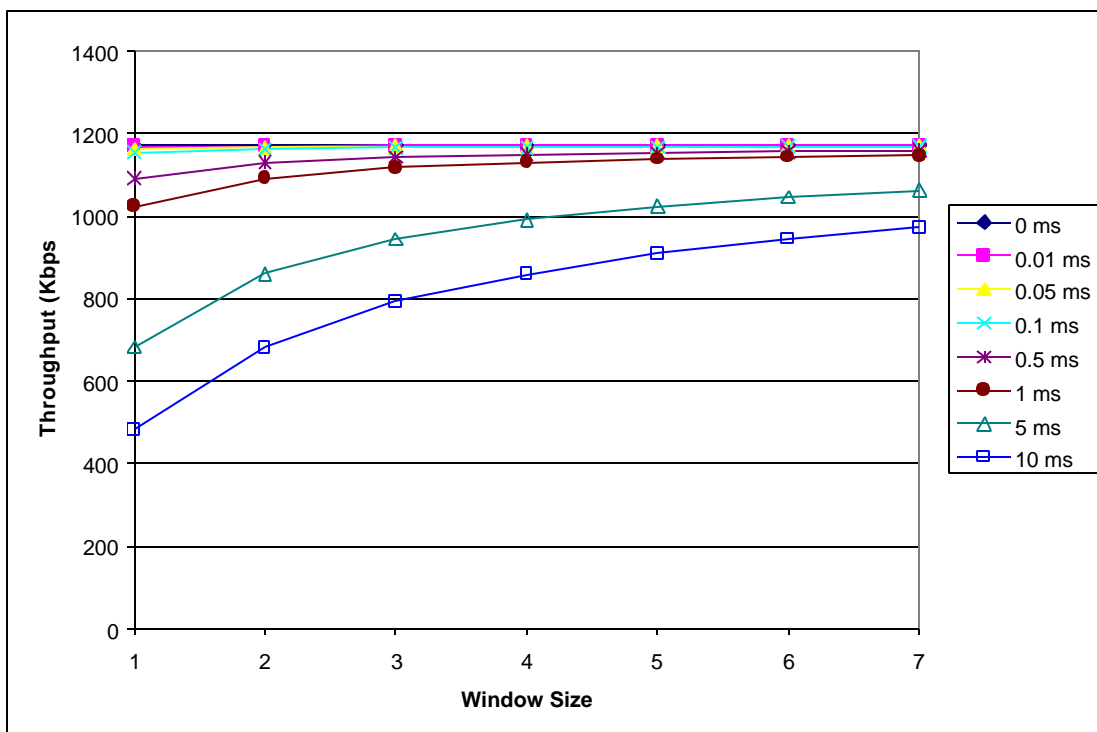
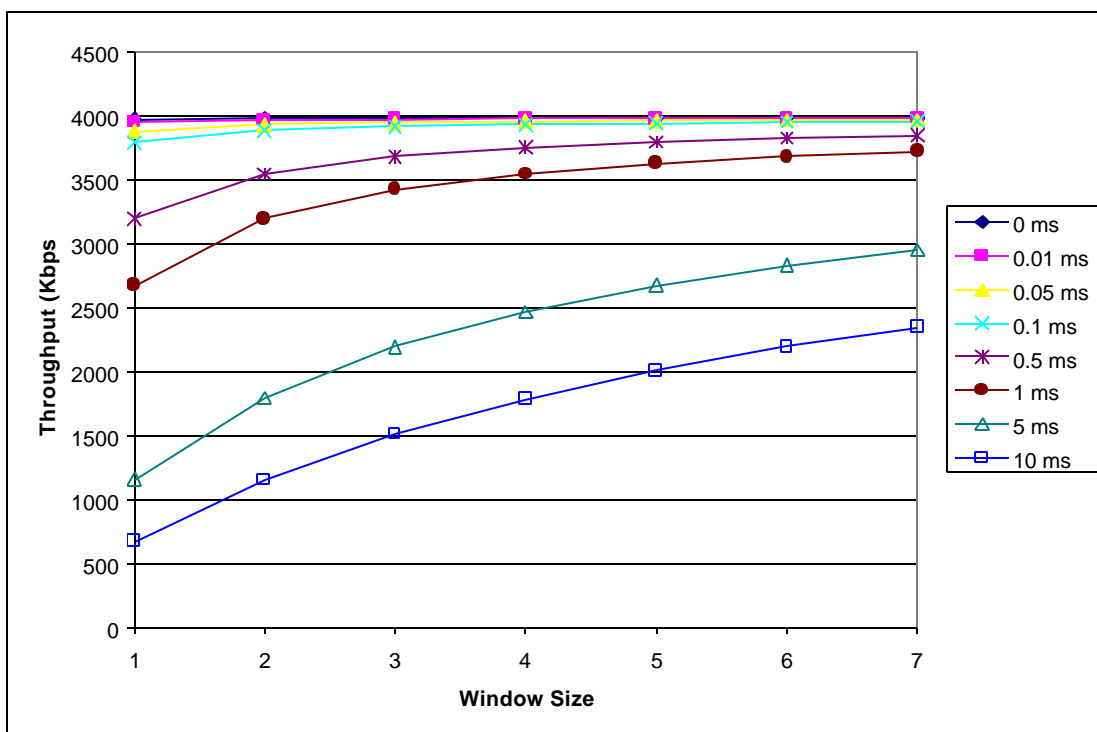**Figure 2: Throughput at 1.152 Mbps Data Rate for Various Window Sizes and Minimum Turnaround Times**

Figure 3: Throughput at 4.0 Mbps Data Rate for Various Window Sizes and Minimum Turnaround Times

When a short minimum turnaround time is used, throughput remains fairly constant no matter what window size is used. When a longer minimum turnaround time is necessary, however, throughput suffers considerably, especially at higher data rates. The loss in throughput can be alleviated through the use of large window and frame sizes. While throughput still drops at high data rates when a long minimum turnaround time is needed, the loss is not as substantial for large window and frame sizes.

3 CONCLUSIONS

The key to attaining high throughput in infrared data transfers is to minimize the amount of time a device is not transmitting data. Since IrDA is a half-duplex protocol, by necessity a device must stop transmitting data and wait for an acknowledgement from the receiving device before it can continue. At low data rates, the time lost to turnaround is not significant. At higher data rates, however, the loss in throughput is much greater. Specifying a small minimum turnaround time, then, will increase throughput. In situations where this is not possible, using a large window and data size can alleviate the negative effect of a large minimum turnaround time by increasing the amount of data sent between turnarounds. Doing so also decreases the amount of header data that must be transmitted. High throughput may then be maintained despite the long minimum turnaround time.

5 REFERENCES

[1] S. Williams. "IrDA: Past, Present and Future." *IEEE Personal Communications*. February 2000.

[2] P. Barker, A. C. Boucouvalas. "Performance Modeling of the IrDA Protocol for Infrared Wireless Communications." *IEEE Communications Magazine*. December 1998.

[3] A. C. Boucouvalas, P. Barker. "IrLAP Protocol Performance Analysis of IrDA Wireless Communications." *Electronics Letters*. December 10, 1998.

[4] "Serial infrared link access protocol (IrLAP) - Version 1.1." Infrared Data Association. Walnut Creek, CA. 1996.

[5] T. Ozugur, P. Kermani and M. Naghshineh. "Comparison of Go-Back-N and Selective Reject ARQ Modes of HDLC Over Half-duplex and Full-Duplex Infrared Links and the Effect of Window Size and Processor Speed in Utilization." *Proceedings IEEE PIMRC '98 Boston, MA*, pp. 708-12.

[6] "IrDA Point and Shoot Profile revision 1.0." Infrared Data Association. Walnut Creek, CA. January 12, 2000.