AN IMPLEMENTATION OF X.25 OVER TCP/IP

A Thesis Presented in Partial Fulfilment
of the Requirements for the Degree of
Master of Technology in Computing Technology
at Massey University

Ajitkumar J. W. Rasiah

1991

ACKNOWLEDGMENTS

I am indebted to my supervisor, Dr. Peter Kay of the Department of Computer Science, Massey University, for his invaluable advice, assistance and helpful criticism throughout this study and the preparation of this manuscript. His patience and encouragement are greatly appreciated.

I also wish to thank:

Professor Robert M. Hodgson of the Department of Production Technology, Massey University, for providing the opportunity for this study.

Professor Mark D. Apperley of the Department of Computer Science, Massey University, and the Massey University Computer Centre for providing the facilities for this study.

I would also like to express my sincere thanks to my parents Dr. and Mrs. B. W. Rasiah for their patience, help and encouragement during the course of this study.

CONTENTS

LIST OF FIGURES

ABSTRACT

Computer Networks are being used increasingly around the world. More importantly, many of these networks are interconnected by gateways, allowing a user in one geographical location to send and receive messages from a host located elsewhere.

The International Standards Organisation (ISO) has proposed the Open Systems Interconnection (OSI) Reference Model as a basis for building computer networks and the protocols which are used on those networks. Developers of networks are encouraged to follow these guidelines so that their networks may have an 'open architecture'.

This thesis examines two networking protocols, CCITT's X.25 Recommendation and TCP/IP. It continuously refers back to the OSI Reference Model as it describes the design of these protocols. It then looks at an implementation of X.25 over TCP/IP, which will allow users on non-X.25 hosts to develop and run X.25 applications courtesy of a separate host that supports X.25.

CHAPTER 1

INTRODUCTION

1.1 The Rapid Growth of Computer Networks

Computer networking is an area of technology which has rapidly grown over the past decade. A few years ago networks were research tools used by a few specialist organizations. Over the period of a few years computer networks gained a reputation for usefulness and reliability and many more organizations recognized the necessity to invest in a public or private network of some kind for daily use with machines ranging from personal computers to supercomputers. A computer network offers a practical means of communication and eliminates the problem of isolation which many organizations have previously experienced. Organizations have come to rely on computer networking to an extent that they are almost helpless without one. World-wide electronic mail is used daily by millions of people. Over the past decade networks have evolved from being a research tool used by academics to an essential tool for users in business, government and research institutions.

As networking evolved, every computer manufacturer had a different networking architecture that was incompatible with the products of other manufacturers. This is now changing. Almost the entire computing industry has agreed on standardizing the architecture of networks. The spin-offs from the resulting standards are already apparent. The products of one manufacturer can communicate easily with the products of other manufacturers.

The advantages of using computer networks are immense. One of the most significant is resource sharing. This enables a community of users in one geographical location to share resources with another community somewhere else. Most universities and research institutions nowadays make their computing facilities available to remote users via resource sharing networks. This means that a lecturer or student at one university can employ the facilities at any other university on the network.

Some networks provide access to specialized facilities. For example, the New York Times has made available millions of abstract news items dating back to the early 1960's. These can be accessed by journalists using a suitable query language. Some corporations have a large number of computer solutions for engineering design or evaluation. These are made available to engineers throughout the corporation via a corporate network.

1.2 International Standards Organizations

Computer networks may be found in most countries today. The equipment and protocols that are used is varied and so it is necessary to define standards for obtaining maximum compatibility. The ISO, CCITT and IEEE organizations make major contributions towards standards in many areas.

The International Organization for Standardization (ISO) is a body which produces recommendations of all types. It has a computer technical committee which is responsible for all types of computer standards. One of its most important achievements is the Open Systems Interconnection (OSI) reference model.

The International Telegraph and Telephone Consultative Committee (CCITT) is a United Nations committee which makes recommendations about telephone and data communications services. The recommendations are revised every four years. Many of the CCITT's recommendations have been adopted to an extent that they have become the standard. X.25 is one such example.

The Institute of Electrical and Electronics Engineers (IEEE) is a large, professional organization which has comprises of a number of smaller interest groups to which members can belong. One of these groups deal with standardization and develop standards in the electrical engineering and computing areas. IEEE's 802 standard for local-area networks is widely used and has subsequently been taken over by ISO as the basis for ISO 8802.

1.3 Protocols

To enable the increasing variety of computers to communicate with one another, there must be a well-defined set of rules about how messages are exchanged between machines. These rules are known as protocols. A protocol is an agreed set of rules and procedures which, if followed by all participants, will allow the orderly transmission of information among the participants.

Many different protocols exist for use over computer networks. A protocol may be used on its own or as part of a 'layer' or hierarchy of different protocols which collectively make up a networking system. Two of the most well-known protocols used today are X.25 and TCP/IP.

1.4 Reference Models

To reduce the complexity in network design, most networks are organized as a hierarchy of levels or layers. This approach was conceived when it was realized that the functions required for data communications are best implemented in a hierarchical fashion. The purpose of each layer is to offer certain services to the layers above whilst shielding them from the details of how the services are actually implemented. A layer on one machine communicates with the corresponding layer on another machine using a protocol suitable for servicing that layer. In this instance, the protocol will not be suitable for use with other layers. The set of layers and protocols is called the network architecture. The network architecture is based on an abstraction called the Reference Model.

The most widely known reference model is ISO's OSI. Work on this was first completed in 1979. In June 1982 it was proposed as a draft international standard and it was promoted to an international standard at the end of 1983. The OSI reference model for network architectures contains seven layers. In reality information is not transferred directly from layer n on one machine to layer n on the peer. Instead, each layer passes data and control information to the layer below it until the lowest layer is reached. Below layer 1 is the physical medium through which the actual communication occurs. At the peer information which is received is taken from the physical medium and passed to the layers above until the originating layer is reached. International standards exist for all seven layers.

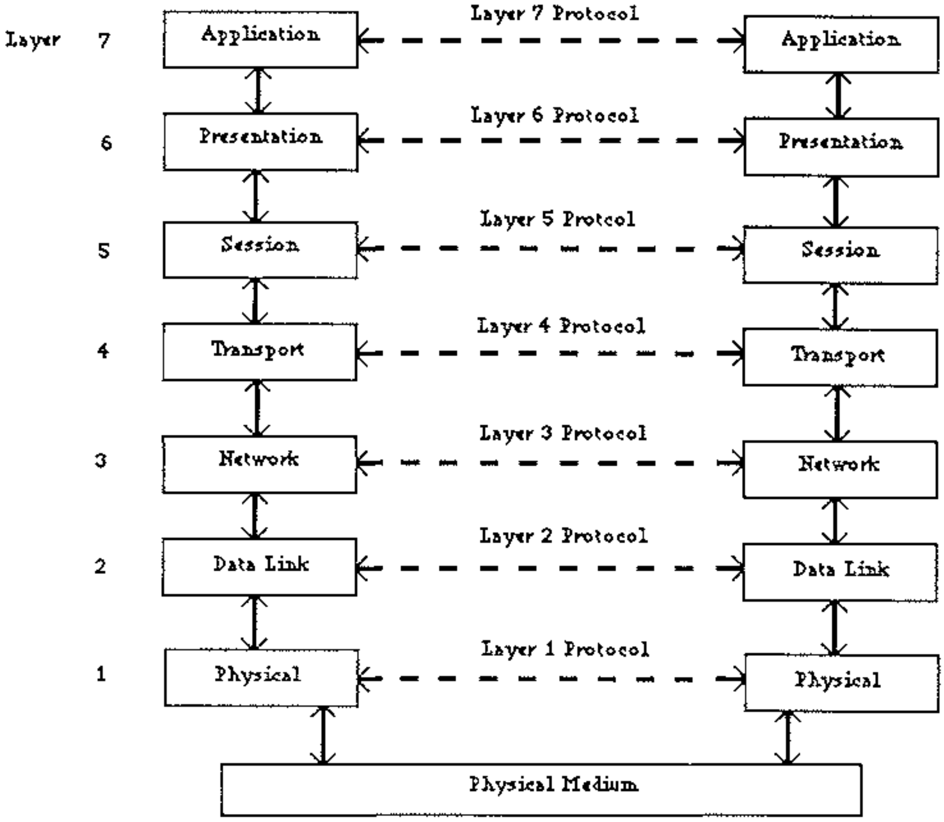

Figure 1.1 - The OSI Reference Model

Layer 1, called the physical layer, is responsible for transferring raw bits from one end of a communication channel to the other. The design issues in this layer are mainly concerned with the mechanical and electrical interfaces to the network and the physical transmission medium.

The data link layer (layer 2) is responsible for taking the raw transmission facility and transforming it into a line that appears free of transmission errors. This is done by breaking down the data that is passed to it by layer 3 above into data frames, typically containing a few hundred octets, and adding error checking information. Layer 2 then passes them to layer 1 which transmit the frames sequentially. The receiver typically uses the error checking information to check that the frame has not been corrupted during transit by noise on the line. It then sends back an acknowledgement frame to the sender, or requests the sender to retransmit the frame if data corruption has occurred.

Layer 3, the network layer, is concerned with how packets are to be routed through the network. It uses the error free channel provided by the data link layer below to communicate packets of data between different hosts on the network. Layer 3 must also control the maximum number of packets in the network at any time. If too many are in transit, the network will become congested and it's performance will degrade.

The transport layer (layer 4) is responsible for accepting data from the session layer, splitting it up into smaller pieces if necessary, passing these to the network layer, and ensuring that the pieces all arrive correctly at the other end. All this must be done efficiently, and in a way which makes the session layer independent of changes in the hardware technology. The transport layer must create network connections as required by the session layer. If the session layer requires a high throughput then the transport layer may create multiple network connections and divide up the data among the network connections to improve throughput. Alternatively, the transport layer may multiplex several transport connections onto a single network connection if creating and maintaining a network connection is expensive. The transport layer also determines the type of service to provide the session layer, such as message broadcasting to multiple destinations, message transmission with no guarantee of the order of delivery, or an error-free point-to-point channel which delivers messages in the order in which they were sent. Transport layer protocols have end-to-end (host-to-host) significance, meaning that this layer and the layers above it are implemented on the host, rather than in the network. The transport layer also provides packet resequencing and flow control to ensure that data arrives safely at the destination.

Layer 5, the session layer, allows users on different machines to establish sessions (connections) between them by providing procedures for initializing, controlling the dialogue and terminating the sessions. A session might, for example involve a user logging onto a time-sharing remote host to use a facility that it offers. One of the services that the session layer provides is dialog control, where it must determine whether the message traffic is exchanged in a full-duplex or a half-duplex fashion. Another essential session service is synchronization. If a large file transfer crashes, then retransmitting the entire file can be avoided since the session layer is able to provide checkpoints into the data stream and therefore only the data after the last checkpoint must be repeated.

The presentation layer (layer 6) provides commonly used functions such as text compression or conversion between different file formats. These functions are requested often enough by users to warrant providing solutions for them rather than letting them solve the problem each time. The presentation layer relies on the underlying layers to provide error-free data transmission. This way, the presentation layer is free to concentrate on the syntax and the semantics of the data.

Layer 7, the application layer, is the highest of the seven layers. This provides a means by which application programs can gain access to the OSI environment for the purpose of exchanging information with application programs on remote hosts.

The network layer and below are largely used by the CCITT X.25 definition. However there have been some alterations at layer 3 since the existing X.25 definition was not completely satisfactory for direct use in the OSI environment.

Variations of the OSI reference model exist under different names. One widely used variation is the TCP/IP Internet Layering model where applications reside directly above the transport layer and have application-specific session and presentation layers. This model forms the basis of the TCP/IP protocol suite.

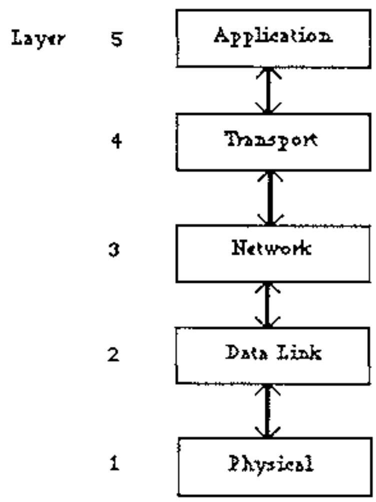

Figure 1.2 - The TCP/IP Internet Layering Model

1.5 Objectives of the Project

This thesis examines two widely used networking protocols, X.25 and TCP/IP. It covers much of the details pertaining to layers 1 to 3 of their architecture. It then describes the design and implementation of software for running X.25 over TCP/IP.

The implementation of X.25 over TCP/IP was stimulated by the networking requirements of the Department of Computer Science at Massey University. The department runs two Ethernet local-area networks which are used to connect a wide variety of devices. These devices include Sun workstations, Apple Macintosh and IBM Personal Computers, and a laser printer. TCP/IP is run over the Ethernet which allows an already extensive networking environment to grow even further. However, the computing facilities at the department do not include a connection to a public data network. Therefore software development in the networking area is restricted to the TCP/IP environment.

The Massey University Computer Center runs a Pyramid minicomputer, which is also connected to the Internet. However, this machine has the added feature that it is connected to a public data network. This provided the motivation to develop software that would allow users at the Department of Computer Science to write local X.25 software, use TCP/IP to reach the Pyramid, and use the Pyramid's connection to the public data network.

The implementation was developed in 'C' under the UNIX operating system. The project consisted of three parts. A programmer's interface in the form of an X.25 library was written to provide users with routines for opening and closing virtual circuits, sending and receiving data in the form of suitably formatted X.25 packets or data, and getting status information for connections. On the Pyramid side, two server programs were written to provide the X.25 services for the X.25 library.

Chapter 2 and Chapter 3 describe in detail the X.25 and TCP/IP protocols respectively.

Chapter 4 describes some underlying, medium-specific network protocols used for the data link layer which are essential for the operation of X.25 and TCP/IP. It describes how these protocols are used to transport X.25 and TCP/IP packets.

Chapter 5 describes the implementation details of this project and presents some examples of how the implementation may be used.

Chapter 6 goes on to propose further developments to the existing system.