9-1990

Gateway design for LAN interconnection via ISDN

Xian-Yu ZHANG
Institute for Infocomm Research

Robert H. DENG
Singapore Management University, robertdeng@smu.edu.sg

Citation

ZHANG, Xian-Yu and DENG, Robert H.. Gateway design for LAN interconnection via ISDN. (1990). *Computer Networks and ISDN Systems*. 19, (1), 43-51.
Available at: https://ink.library.smu.edu.sg/sis_research/34

Published in Computer Networks and ISDN Systems
Volume 19, Issue 1, September 1990, Pages 43-51
https://doi.org/10.1016/0169-7552(90)90117-B

Gateway Design for LAN Interconnection via ISDN

Xian-Yu ZHANG and Robert H. DENG
Institute of Systems Science, National University of Singapore, Kent Ridge, Singapore

Abstract. Recently, the use of bridges/gateways to interconnect physically distant local area networks (LANs) has become increasingly popular. There are various ways of connecting these bridges/gateways. ISDN is one of them and an attractive one. In this paper we discuss our work in the gateway design for interconnecting LANs via ISDN.

Keywords. Gateway, IP router, LAN, ISDN, LAN interconnection.

1. Introduction

Local area networks (LANs) have been designed for communication using data rates of a few Mbps and covering small areas. With the large number of LANs now in use in universities, research laboratories, industry, and commerce, interconnection of LANs over several to hundreds of kilometers are becoming common interest. Traditionally, two technologies are used for long distance communications between LANs [1]: dedicated data networks based on high-speed leased lines which offer a data transmission rate of 64 kpbs and the Public Switched Telephone Network (PSTN) which offers the circuit switching without high speed. Leased lines are expensive to use and non-switchable. The data rate on the PSTN is often too low (seldom better than 1.2 kbps) to be used for LAN to LAN communications. However, in many situations the PSTN is often the only choice, it is switchable and it is available anywhere.

The Integrated Services Digital Network (ISDN) has the advantages from both the leased line and the PSTN. It provides a set of reliable (low error rate) bit pipes with considerable large transmission capacity than that available from most existing wide area networks and yet with the convenience and easy access of the PSTN. In this paper, we discuss gateway design problems for interconnecting LANs via the ISDN. Such a gateway is currently under development at the Institute of Systems Science, National University of Singapore. The gateway project is intended to connect LANs at various R&D institutions and academic departments in Singapore with the use of ISDN. We consider LANs which follow the ANSI/IEEE 802 standards [2]. Examples of such LANs include Ethernet, IBM Token Ring and 3COM LANs. The Transmission Control Protocol/Internet Protocol (TCP/IP) [3,4] is used on top of the Logical Link Control (LLC) layer as the network and transport protocols. TCP/IP, developed originally for use with ARPAnet, has been

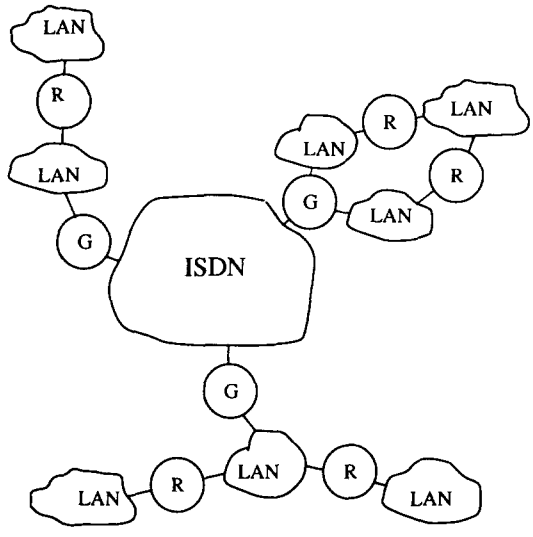

R: IP router
G: Gateway

Fig. 1. LAN interconnection via ISDN.

adopted as the internet and transport protocols by the U.S. Department of Defense (DoD). It has also over the years become a de facto standard in much of U.S. (as well as Singapore) university community involved in computer networks and has been incorporated as well by a number of manufacturers into intelligent systems designed to communicate over LANs. The IP is a datagram protocol designed for use in interconnected systems of packet-switched networks. The TCP is a connection-oriented, and end-to-end reliable protocol. The TCP interface can be visualized as a set of routines used by higher-level programs to achieve process-to-process communications.

The interconnected network is depicted in Fig. 1. At each location, a number of LANs are interconnected by IP routers [5] to form an LAN cluster. The ISDN functions as a switchable transparent network between gateways. We wanted a gateway that would:

(1) enable workstations within one LAN cluster to communicate with workstations in other LAN clusters;

(2) require minimal changes to the existing system hardware and software; and

(3) be compatible with the IP router operations and share the same internetwork management protocols and user interface.

As will be discussed in Sections 2 and 3, the IP router can be upgraded to meet our gateway requirements.

The remainder of the paper is organized as follows. Section 2 discusses selection of ISDN interface protocols to the gateway. Section 3 presents some gateway design issues. The gateway software architecture is given in Section 4 together with a detailed discussion on the gateway management functions. Finally, Section 5 is a summary.

2. Selection of Interface Layers

The basic building block of the ISDN is a 64 kbps channel, referred to as the "B" or bearer channel. This channel is used to transmit user information. Another channel, called the "D" channel, carries signaling or control information as well as user data. This information is used to setup, redirect, or terminate calls. The ISDN supports two standard interfaces to connect the customer premise equipment to the network. The basic rate interface is two B-channels and one 16 kbps D-channel, "2B + D", for a total bandwidth of 144 kbps. The primary rate is thirty B-channels and one 64 kpbs D-channel, "30B + D", for a total bandwidth of 1.984 Mbps [6,7]. An important feature of the ISDN is the recursive application of the OSI seven-layer protocol structure in modeling the two generic types of information flows, namely, user information flow over B-chan-

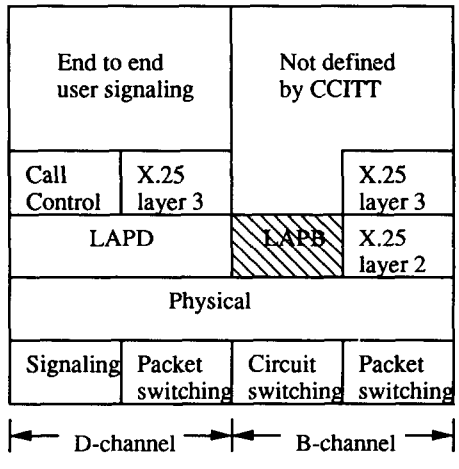

Fig. 2. Layered protocol structure at the ISDN user-network interface.

nel and control information flow over D-channel (see Fig. 2).

In the interconnected network, as mentioned before, the ISDN is used as a transparent network between the remote LANs. In such a design, the basic functions of the gateway are to wrap a LAN packet with ISDN protocol header to form an ISDN packet and to deliver the resulting packet to the destination gateway. At the destination gateway, the original LAN packet will be recovered by simply dropping the ISDN header and then sending the original LAN packet to the destination LAN where the receiving station is located. Given the above LAN and the ISDN protocol structures, the first step in the gateway design is to select layers at which the protocol interface will be performed. Protocol layer selection has a great impact on the gateway design and implementation. In the following we list several possible approaches on the protocol layer selection.

2.1. LLC-X.25 Interface

The first possibility is to choose the LLC layer in LAN and the X.25 layer in the ISDN as the interface layers. This approach should have a superior performance to the IP-X.25 approach discussed below, since fewer layers of protocol need be processed by the gateway. Also, because only LLC layer is involved in the gateway software, no constraint is exerted on the use of higher layer protocols.

Communication between two stations anywhere in the interconnected network requires the existence of a route between the individual LANs on which the two stations reside, i.e., the source gateway must decide which frame to forward to which destination gateway via the ISDN. This can be done by using station addresses contained in the LAN frames. The gateway maintains routing tables which represent up-to-date topology or station-location information. The routing table can be programmed into the gateway or transmitted to the gateway by a human or program responsible for management of the interconnected network. Frames on a LAN pass a gateway at high speed and a gateway must be able to execute a table lookup fast enough so that the incoming frames are forwarded in time; otherwise, frames destined to other LANs may be lost due to gateway congestion. Therefore, large routing tables at the gateway are not desirable. However, LLC uses flat addresses, the size of the routing tables can be on the order of half the number of stations in the interconnected network [8]. More importantly, this approach violates our gateway design requirements (2) and (3) described in the last section.

2.2. IP-X.25 Interface

IP is designed for use in interconnected packet-switched networks. The IP provides for transmitting blocks of data called datagrams from sources to destinations, where sources and destinations are stations identified by fixed length addresses called internet addresses. The internet address is in the form Net.Host, where Net is a LAN address and Host identifies a station in that LAN. Because of the hierarchical address of IP, the routing tables in the gateway can be made very small and consequently, the time required for table lookup is reduced significantly.

Besides addressing and routing problems discussed above, both the LLC-X.25 and the IP-X.25 approaches possess some common advantages and disadvantages:

(1) since multiple virtual circuits can be multiplexed over one B-channel, it is possible for one LAN to communicate with multiple remote LAN clusters simultaneously over a single B-channel;

(2) because of the packet-switching nature of X.25, long distance LAN to LAN communications can be achieved at lower cost than that using circuit-switching but at the expense of lower transmission speed and longer delay; and

(3) the X.25 packet length is usually smaller than that of the IP datagram; therefore, packet fragmentation and reassembly are needed to relay IP datagrams via the ISDN.

2.3. IP-LAPB Interface with Circuit-switching

To reduce the ISDN delay, a LAPB (Link Access Procedure-Balanced) [9] layer can be implemented on top of the B-channel with circuit-switching. The gateway protocol configuration with this approach is shown in Fig. 3. In this configuration, gateways are connected with the ISDN at the ISDN S/T reference point [7]. Circuit switched mode connection with 2B + D access structure is used between gateways. The LAPB is implemented on top of the B-channel circuit

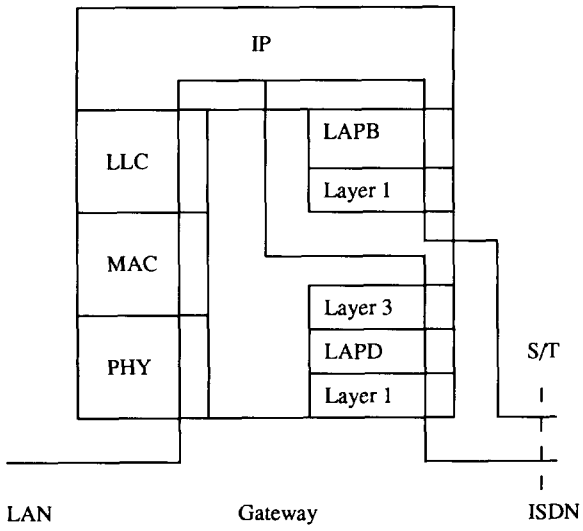

Fig. 3. Gateway protocol configuration.

switched physical channel to provide transparent delivery of the IP datagrams. Some of the interesting features of this approach are:

(1) small routing tables at the gateways due to the hierarchical IP address structure;

(2) simple gateway congestion control schemes; and

(3) full utilization of the B-channel transmission capacity.

Due to these considerations, this approach is used in our gateway design. In the following sections we examine design issues for the IP-LAPB approach. These are: address translation; B-channel connection setup and clearing; fragmentation and reassembly; congestion and flow control; and circuit management.

3. Design Issues

3.1. Address Translation

When an IP datagram arrives at the gateway, the gateway analyzes the IP header to determine whether this datagram contains control information intended for the gateway, or data intended for station on a remote LAN. In the latter instance, the gateway carries out address translation by performing a table lookup. The format of the table is shown in Table 1.

Table 1
Address translation table

IP network address	Gateway ISDN number
NetX1	NumberY1
NetX2	NumberY2

3.2. B-channel Connection Setup and Clearing

The destination gateway ISDN number obtained via address translation is provided to the D-channel layer three "call control procedures for circuit switched calls" [7] to establish and terminate B-channel circuit-switched connection and LAPB data link connection between the calling gateway and the called gateway. Simple procedures for connection setup and clearing are given in Figs. 4 and 5, respectively.

3.3. Circuit Management

As mentioned above, before data can be sent between interconnected LANs the B-channel physical circuit switched connection must be setup. However, since IP is a connectionless protocol, it gives no indication when to open and close a connection. One approach to solve this problem is to let the gateway monitor the incoming IP datagrams, and the first datagram destined for a re-

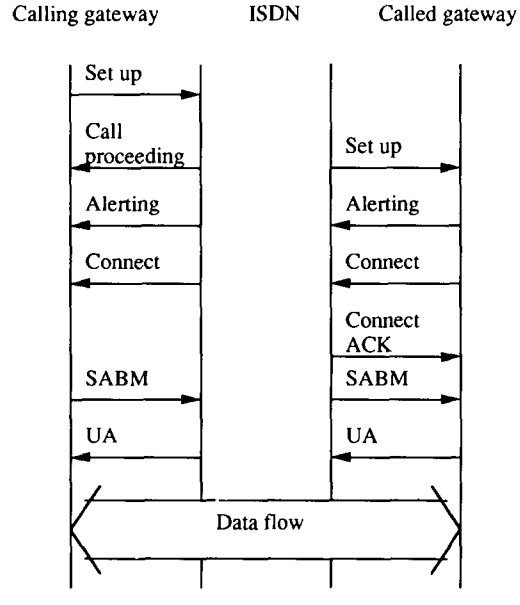

Fig. 4. Establishment of circuit-switched B-channel connection and LAPB data link connection between gateways.

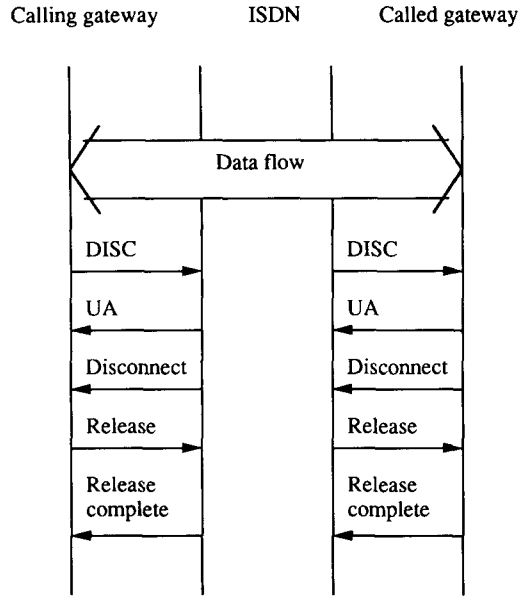

Fig. 5. Clearing of circuit-switched B-channel connection and LAPB data link connection between gateways.

mote LAN triggers the circuit to initiation of a B-channel setup to the corresponding remote gateway [1]. Once the circuit connection is established, it is used by the gateway to delivery subsequent datagrams to the remote LAN. However, since IP provides no disconnect information, the gateway must make a decision as to when to terminate the circuit based on its perception that traffic on the circuit has ceased.

Circuit management depends on the tariff structure of the ISDN. Over eager disconnection may cause excessive circuit setups with consequent cost and connection setup delays, while maintaining an idle circuit connection for a long time is obviously wasteful. Let S be the cost of opening a B-channel connection between the two gateways. Let t be the elapsed time since the last use of the open B-channel. Let P be the cost per time unit of maintaining an open circuit. A simple circuit management algorithm can be devised as follows:

Algorithm: If $t \geqslant T$, close the circuit;
else, leaves the circuit open,

where T is a time threshold lower bounded by S/P. Selection of the values of T must take specific application environments into consideration. Sometimes a tradeoff need be made between the circuit setup and circuit disconnect delays, and the ISDN tariff structure. The above circuit release algorithm is similar to the method proposed in [10] for call clearing, where the time-out interval is selected on the basis of economic and implementation specific criteria.

Since TCP is a connection oriented protocol, both connection and closing messages are provided in TCP. Another approach in circuit management is to maintain a finite state machine in the gateway which keeps tracking the states of the TCP OPEN and CLOSE entities. According to these states the gateway will be able to determine when to open and close a circuit connection. This method is very efficient in the usage of the ISDN B-channel, but its implementation is quite complicated.

3.4. Fragmentation and Reassembly

One problem often encountered in the network interconnection is packet fragmentation. If the maximum data field length in the ISDN frame is less than the maximum IP datagram length, IP datagram must be fragmented at the source gateway, and reassembled at the destination gateway. However, fragmentation is considered harmful [11]. The harm comes, in our case, mainly from the following:

(1) Fragmentation causes inefficient use of B-channel transmission capacity and gateway processing power. Poor choice of fragment sizes can greatly increase the cost of delivering a datagram. Additional bandwidth is used for the additional header information (a minimum of 20 bytes is required for an IP header), destination gateway must reassemble the fragments.

(2) Loss of fragments leads to degraded performance. Reassembly of IP fragments is not very robust. Loss of a single fragment due to B-channel transmission errors and gateway congestion requires the higher level protocol to retransmit all of the data in the original datagram, even if most of the fragments were received correctly.

(3) Efficient reassembly is hard at the destination gateway. Given the likelihood of lost fragments and the information present in the IP header, there are many situations in which the reassembly process, though straightforward, yields lower than desired performance [11].

In our gateway design, since the maximum frame length of the LAPB is at designers choice, we avoided fragmentation by choosing the maxi-

mum data field length of the LAPB frame equals the maximum IP datagram length.

3.5. Flow and Congestion Control

Flow control is concerned with a pair of nodes trying to ensure that the rate of transmission of packets from the source shall not exceed the capacity of the destination to receive the packets. The sender can create packets faster than the destination is able to receive or to process them. In an internetwork environment, flow control is essentially between a LAN station and the gateway. The gateway may not be able to process datagrams as fast as it receives them because of lack of resources such as buffers or CPU time.

On the other hand, congestion control is used to control the number of datagrams arriving from all sources (LAN stations) at the gateway and preventing it from being overloaded. As internet traffic increases, a gateway could encounter a period of severe congestion. A slow gateway will not be able to cope with the sudden burst of traffic. A gateway must have greater processing power and buffer capacity to encounter the sudden traffic situation. Congestion may also occur because the adjoining network (the ISDN in our case) may be slow to accept datagrams from the gateway. The gateway is bound to be a bottleneck if there exists much internet traffic. There are several efficient congestion control methods but they are all complex schemes [12].

To meet the high internet throughput requirement at reasonable cost in our system, the functions performed by the gateway must be simple, which excludes its performing any complex flow or congestion control. Gateway congestion can be remedied partly by ICMP (Internet Control Message Protocol), as described in the next section. However, ICMP cannot prevent gateway buffer from being overflowed. Since IP is a connectionless protocol, both flow and congestion control are performed by simply discarding datagrams in case the gateway buffer is full. Discarded datagrams will be recovered through the higher level protocol, i.e., TCP. In such a design, the gateway buffer size has a great impact on the system performance. To study this, an simple analytical model is formed for the gateway in [13] and is shown here in Fig. 6. In Fig. 6 the ISDN is modeled by two FCFS servers corresponding to the B-channels handling the two directions of data flow, and the gateway is modeled by a receiving buffer R and a transmitting buffer T. Buffer R receives packets from an ISDN B-channel and transmits them onto the LAN to which it is attached. Since the transmission rate of the B-channel is much lower than that of the LAN, congestion at buffer R is highly unlikely. Buffer T receives datagrams from the source LAN whose destination is a remote LAN and forwards them over an ISDN B-channel. The size of the buffer T is assumed to be K datagrams. An approximate analysis of the gateway can be carried out as follows. We assume that the datagrams arriving at buffer T follow Poisson distribution with a mean arrival rate of λ datagrams per second. We also assume that datagram length has an exponential distribution with an average length of L bits. Then the transmitting buffer T and the B-channel can be modeled as a $M/M/1/K$ queue (see Fig. 6) with a mean service rate μ ($= 64$ kbps$/L$ bits) datagrams per second. Define the B-channel offered traffic intensity as $\rho = \lambda/\mu$, the probability of buffer T overflow, or the probability of discarding a datagram at buffer T, is given by [12]

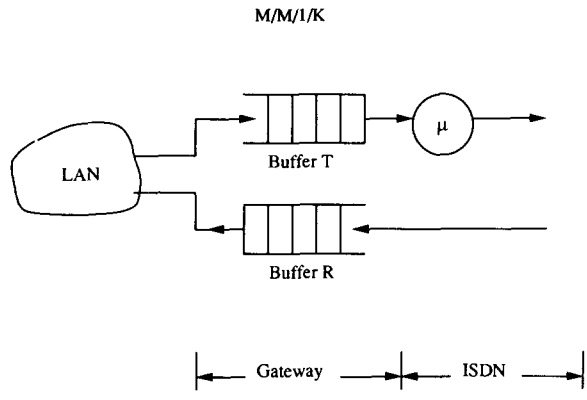

Fig. 6. An analytical model.

$$P_d = (1-\rho)\rho^K/(1-\rho^{K+1}),$$

and the internet throughput normalized by the B-channel service rate is given by

$$\prod = (1-P_d)\rho.$$

The overflow probability P_d and the normalized throughput $\prod$ versus the B-channel utilization are plotted in Figs. 7 and 8, respectively. These two figures can be used as design curves in the gateway design. For a detailed study on the perfor-

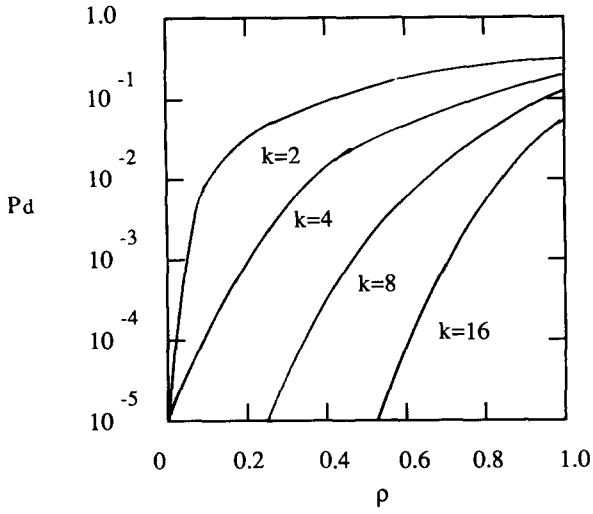

Fig. 7. The gateway buffer overflow probability versus the B-channel offered traffic intensity.

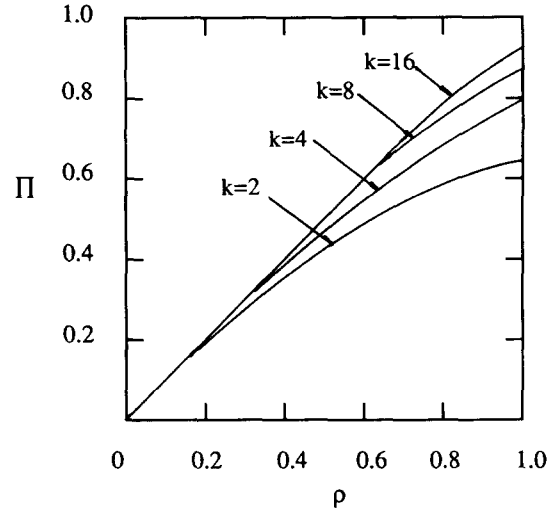

Fig. 8. Normalized internet throughput versus the B-channel offered traffic intensity.

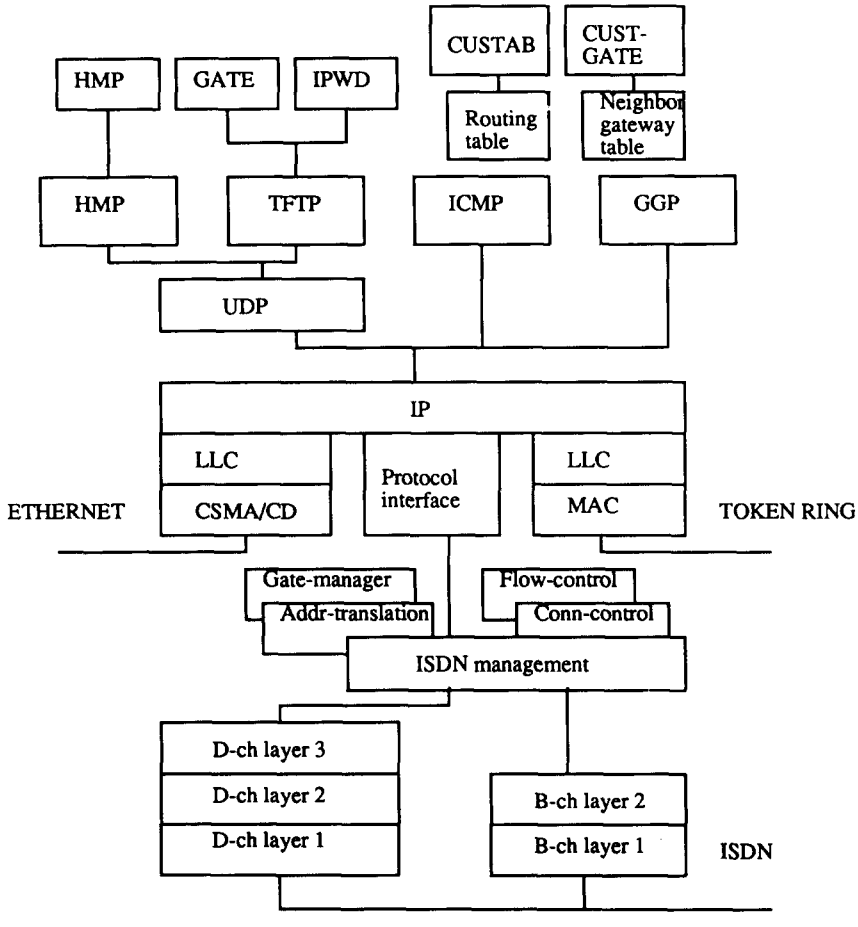

Fig. 9. Gateway software architecture.

mance of the interconnected network please refer to [13].

4. Gateway Software Architecture

The gateway runs on a dedicated IBM PC/AT system. Its software architecture, a modification of the IP router structure [5], is depicted in Fig. 9. Besides IP protocol software and protocol interfacing software, there is a gateway management software which is fully compatible with the IP router management functions. The gateway management software carries out such important tasks as initialization of the gateway environment, gateway status monitoring and maintaining, and updating routing table. The architecture of the gateway software is summarized in the following.

4.1. The Gateway Protocols

The gateway supports the following protocols: TFTP, IP, ICMP, GGP, and HMP.

TFTP (Trivial File Transfer Protocol): It is a short file transfer protocol supported by IP. It is used to update files on the gateway disk.

IP: The IP module functions like a switchboard. It forwards IP datagrams from one interface to another according to the contents of its routing table. IP passes outgoing datagrams to the ISDN interface and the incoming datagrams to the LAN interface. Once an IP datagram is received, the IP verifies the datagram with the least amount of overhead. Items checked are: IP version number, IP header length, IP message length, IP header checksum, and time to live. If any of the items is not checked the datagram is dropped. This prevents bad datagrams from propagating throughout the network. Once the datagram has passed the tests, the IP destination internet address is examined. If the destination address is not found in the routing table, the gateway sends an ICMP destination unreachable message back to the source address. If the destination address is found, the datagram is forwarded to the appropriate interface. In case the datagram is addressed to the gateway itself, the internet data field is examined to find the higher level gateway management protocols. The protocols supported are: ICMP, GGP, and HMP.

ICMP (Internet Control Message Protocol): ICMP supports a number of internet control messages. The messages most used by the gateway are:

(1) messages answering to ICMP echo request from stations on the internet. The source usually sends the echo request message with the intent of interpreting the response as yes if the gateway is alive or as no if the gateway is not;

(2) ICMP source quench message indicates that the number of free buffers in the gateway transmitting queue had reached its minimal acceptable level;

(3) ICMP destination unreachable message indicates that the destination address is not in the routing table; and

(4) ICMP redirect message which is sent to the source address indicating a more direct route should be taken. The source station then should retransmit its data using the shorter path.

GGP (Gateway-to-Gateway Protocol): The gateway should have prior knowledge of at least one neighbor gateway/IP router. This allows neighbor gateways/IP routers to exchange routing information. The list of neighbor gateways/IP routers and the list of networks will be initialized by the management commands. When the gateway is initialized, the status of all known gateways/IP routers is down. Every 15 seconds the router sends its neighbor gateways/IP routers GGP echo messages. To ensure that the gateway and its neighbors are communicating without difficulty, each neighbor gateway/IP router must respond to at least three out of four GGP echo messages. When less than three responses are received, the gateway changes the status of that neighbor to down.

HMP (Host Monitoring Protocol): HMP collects statistics information within the gateway. HMP permits the monitoring of the gateway from remote locations. The gateway never has to write anything to the screen, so a monitor at the gateway is not necessary. With an HMP client station, an internet administrator can perform all the functions listed below even when the gateway is physically located at a different place: query (neighboring gateway's/IP router's internet address and status, communicating or not communicating, current routing table, throughput since last query); add or delete routes from table; add or delete neighbor gateways/IP routers; shutdown gateway; reboot gateway; transfer files to and from the gateway.

4.2. Administrator's Programs

Some programs in the gateway can be used by the gateway administrator to perform gateway management, e.g., CUSTAB, CUSTGATE, IPWD and GATEWAY. These programs are supported by the protocols mentioned above. The CUSTAB is used to modify the binary file ROUTTAB.SYS which the gateway uses to determine explicit routes. These routes are redirections to machines which have no dynamic routing capabilities. The CUSTGATE is used to modify the binary file GATES.SYS which the gateway uses to determine its neighboring gateways/IP routers. The gateway reads the GATES.SYS from the current directory. The GATEWAY accomplishes the initialization of the gateway software and runs it. The IPWD is a routine which provides the administrator with the ability to change password that will be used by the gateway. The HMP command is also used to manage the gateway password.

4.3. Environment and System Files

The gateway requires that an environment variable GATE = be set to determine how many and which LAN, and ISDN interfaces are to be used. The following is an example of a statement in an autoexec.bat file that would be used to start the gateway:

GATE = UTI
GATE

The letters following GATE = indicate that this gateway had one Ungermann-Bass Ethernet card, one IBM Token Ring card (or IBM PC Network) and an ISDN card. These letters also describes the specifications found in NETDEV.SYS, NETDEV1.SYS, NETDEV2. SYS. These files are the device driver files which must be specified in CONFIG.SYS file:

DEVICE = DRIVER:PATH \ NETDEV*.SYS
...

The CUSTOM command can be used to modify any of these network device drivers.

The gateway is dependent on two binary files to load the initial routing table and the neighboring gateways/IP routers: ROUTTAB.SYS and GATES.SYS. These files reside on the diskette or in the directory from which the gateway program is executed. Another system file is the SECURITY.SYS which is used by the IPWD program and HMP command to change the password.

5. Summary

In this paper we have discussed our pre-trial experience in gateway design for interconnecting physically distant LANs via ISDN. The ISDN is seen as an attractive alternative to the conventional ways for long distance LAN communications. The gateway is an upgrade of the IP router and therefore requires minimal hardware and software changes in the existing system. More importantly, the gateway management software is fully compatible with the IP router management software which ensures that the internet be managed in an unified manner.

The current project uses the 2B + D basic rate ISDN interface; however, the gateway architecture described here applies to other ISDN interface structures also, e.g., 30B + D interface.

References

[1] G. Knight, D. Deniz and J. Fan, Gateways to the ISDN, in: *Conference Proceedings ISDN 87* Middlesex, UK (1987) 141–152.

[2] ANSI/IEEE Standards 802.2–1985, 802.3–1985, 802.4–1985, and 802.5–1985, Institute of Electrical and Electronics Engineers, Inc.

[3] Defense Advanced Research Projects Agency, Internet Protocol. RFC: 791, Sept. 1981.

[4] Defense Advanced Research Projects Agency, Transmission Control Protocol. RFC: 793, Sept. 1981.

[5] C. Douglas, *Internetworking with TCP/IP: Principles, Protocols, and Architecture* (Prentice-Hall, Englewood Cliffs, NJ, 1987).

[6] S. Kano, Layers 2 and 3 ISDN recommendations, *IEEE J. Select. Areas Comm.* **4** (1986) 355–359.

[7] CCITT Red Book, Recommendations of the Series I, 1984.

[8] J.O. Limb and C. Flores, Description of Fasnet–A unidirectional local-area communications network, *Bell Systems Tech. J.* **61** (Sept. 1982).

[9] CCITT Red Book, Recommendations X.25, 1984.

[10] ISO 8473: Information processing systems–Data communications protocol for providing the connectionless mode network service, 1988.

[11] C.A. Kent and J.C. Mogal, Fragmentation considered harmful, in: *Proc. SIGCOM'88 Symposium: Communications Architectures & Protocols*, Stanford, CA (1988) 390–401.

[12] D. Bertsekas, and R. Gallager, *Data Networks* (Prentice-Hall, Englewood Cliffs, NJ, 1987).

[13] W.C.L. Chiew and R.H. Deng, Performance analysis of a LAN-ISDN-LAN interconnected network, in: *Proc. SICON'89*, Singapore (1989) 389–393.