European Space Agency
European Space Research and Technology Centre
ESA Computer and Network Operations Department

Delft University of Technology
Department of Electrotechnical Engineering
Telecommunications Group

Title: Evaluation of the availability, quality, performance and possible applications on narrowband-ISDN

Author: H.A. Lisman

Subject: Thesis

Size: 107 pages (94 + xiii)

Period: 1 September 1993 - 31 May 1994

Date: 30 May 1994

Mentors: Dipl.Ing. M. Boele, Ir. J.A.M. Nijhof, Prof.Dr. R. Prasad

Code: A 518

Period: September 1993 - May 1994

Acknowledgement

I wish to thank the following people for their help and kindness:

Mr. E. de Jong, head of the ESTEC division of the ESA Computer department, who offered me to do my thesis work in the ESA Computer and Network Operations Department (ECNOD).

Mr. K. Blank, head of the Network Engineering and Implementation Sector of ECNOD, also for offering me the position in the group, and for his comments and suggestions on my work.

Mr. M. Boele, my supervisor, for his extensive support, comments, and suggestions on the work I have done.

Summary

The ISDN (Integrated Services Digital Network) is becoming increasingly important for both data and non-data communications. For the evaluation of ISDN functionality and services in the area of ESA's communications requirements the ESA Computer and Network operations Department has built an ISDN testbed at ESTEC.

This report describes results of testing and evaluation of availability and quality of national and international ISDN bearer channels. ISDN channels seem an efficient and flexible alternative for leased lines. A cost-comparison between ISDN and dedicated leased lines is also presented.

Secondly, several applications using the ISDN network have been evaluated. The digital ISDN lines offer a bigger capacity than existing analog telephone lines and therefore several new applications can be used. Typical applications are the interconnection of LANs and desktop applications like filetransfer, fax and videoconferencing.

Measured performance of LAN-to-LAN connections has been compared to theoretical values that have been calculated.

Indexing terms: *ISDN, leased line, WAN*

Table of contents

Glossary

Basic Access. A type of user-to-network interface supporting combined data rates up to 144 kbit/s.

Basic Rate Channel Structure. A composite communications capacity across the user-to-network interface of an ISDN consisting of two B channels and one 16 kbit/s D channel.

B channel. Bearer channel. An information carrying channel with a capacity of 56 kbit/s (US and Japan) or 64 kbit/s (Europe), providing synchronous data connectivity between a specific pair of end users.

Bearer Service. A type of telecommunication service that provides the capability for the transmission of signals between user-to-network interfaces.

Bit Stuffing. The process of inserting zeros into a sequence of ls in order to prevent the occurrence of too many 1s in sequence.

Broadband Channel Structure. A composite communications capacity across the user-to-network interface of a broadband ISDN.

Broadband ISDN. An Integrated Services Digital Network providing a range of data transmission rates up to several hundred Mbit/s.

Calling Line Identification. A service whereby the calling subscriber's number can be identified to the called subscriber prior to answering the call.

Circuit Switching. An information transfer mode in which switching and transmission functions are accomplished by permanent allocation of channels or bandwidth between the connections.

Common Channel Signalling. A method of signalling in which signalling information relating to a multiplicity of circuits is conveyed over a separate channel by addressed messages. Also used for transfer of network management and operations messages.

Common Channel Signalling Network. A network used to transfer signalling information and consisting of signalling points, signalling transfer points, and interconnecting signalling links.

Common Channel Signalling System. A set of protocols used to convey signalling information over a common channel signalling network.

Customer-Premises Equipment. Telecommunications equipment located on the premises of the subscriber to a telecommunications service.

Cyclic Redundancy Check. A scheme used to detect errors that may occur during the transmission of a frame of data. At the sending side a code is generated that is transmitted with the data. The code is compared with another code generated at the receiver.

Data Circuit-Terminating Equipment. Data communications equipment such as a modem that terminates the circuits provided by the network and connects a DTE to the network.

Data Link Layer. The second layer of the OSI Reference Model that ensures reliable and efficient transmission of information between adjacent stations in a network.

Data Terminal Equipment. End user devices such as terminals or computers that convert user information into data signals for transmission, or reconvert received data signals into user information.

D channel. Data channel. A portion of the communications capacity across the user-to-network interface with a data rate of either 16 kbit/s (for BRI) or 64 kbit/s (for PRI). It is used to transfer signalling information.

Digital Subscriber Signalling System. A set of protocols used for conveying signalling information across the user-to-network interface of an ISDN.

Digital Section. Transmission facility between the network termination on the customer premises and the local exchange.

Digital Transmission System. Physical transmission medium between the network termination on the customer premises and the local exchange .

Echo Cancellation. Elimination from the receiving path of a sender of the portion of a transmitted signal that has been reflected back to the sender. Cancellation is accomplished by subtracting a delayed and attenuated portion of the transmitted signal from the total received signal.

Frame. A block of consecutive time slots whose positions are identified by reference to a marker indicating the start or end of the frame. Each time slot may contain one unit of data.

Frame Check Sequence. An error-detecting code placed in the transmitted frame that checks for errors when the data are received.

High-Level-Data Link Control (HDLC). Bit-oriented data link control procedure. The international standard for data link control developed by ISO.

Integrated Services Digital Network. A telecommunications network providing integrated services over digital connections between user-to-network interfaces.

ISDN User Part. A protocol of Signaling System No. 7 that provides the functions required for the control of basic bearer services and supplementary services in an ISDN.

Line Termination. The physical termination of the subscriber loop on the network side.

Link Access Procedure D Channel (LAPD). A layer 2 protocol used to control the reliability and efficiency of the exchange of messages over the D channel of the user-to-network interface in an ISDN.

Multipoint. A communications configuration in which one channel interconnects more than two stations, only two of which can communicate with each other at any one time.

Narrowband ISDN. An Integrated Services Digital Network providing data transmission rates up to 2 Mbit/s.

Network. A collection of links and nodes that provides connections between two or more defined points to facilitate telecommunication between them.

Network Layer. The third layer of the OSI Reference Model that manages routing, flow control, and sequence numbering of messages that are transferred over a network connection.

Network Termination. A functional group containing functions required for the physical or logical termination of a network on the user-to-network interface.

Packet Switching. A method of transmission in which small blocks of data called packets traverse in a store-and-forward method from source to destination through the intermediate nodes of the communications network and in which the resources of the network are shared among many users.

Physical Layer. The lowest layer in the OSI Model whose function is to provide a direct connection between neighbouring nodes in a network.

Point-to-Point. A transmission circuit that connects two devices directly, without any intermediate devices.

Primary Access. A type of user-to-network interface supporting combined data rates up to 2 Mbit/s.

Primary Rate Channel Structure. A composite communications capacity across the user-to-network interface of an ISDN consisting of a combination of B channels, and one 64-kbit/s D channel.

Private Automatic Branch Exchange. A switching system usually located on the customer premises for establishing connections between subscribers and public transmission facilities. Also used for on-premises switching and for enhanced communications services.

Pulse Code Modulation. The conversion of an analogue signal to digital form where the signal amplitude is sampled periodically, each amplitude is quantized into a finite number of levels, and each level is represented by a pattern of binary digits.

Rate Adaption. The change of the rate at which a device transmits or receives data to or from the rate required for the transmission of the data on the network.

Reference Configuration. An arrangement of functional groups and reference points.

Reference Point. A conceptual point of demarcation between two non-overlapping groups of functions .

Signalling. The exchange of information between network nodes or between network nodes and subscribers for the purpose of establishing and controlling connections and for management in a telecommunications network.

Signalling Link. A transmission circuit consisting of a signalling data link and its transfer control functions. Used for the reliable transfer of signalling messages.

Signalling System. A set of protocols for the exchange of signalling information.

Supplementary Service. A telecommunications service that is supplementary to a basic bearer service or teleservice.

Subscriber-Access Network. The total configuration of terminal equipment, subscriber-access line, and exchange termination between the customer premises and the network.

Terminal Adapter. A conversion device that allows the attachment of non-ISDN compatible user terminals to an ISDN interface.

Terminal Equipment. A functional group that includes functions required for protocol handling, maintenance, interfacing, and connection to other equipment.

User-to-Network Interface. The point of demarcation between the terminal equipment and the network termination.

Virtual Circuit. A communications path through a network established prior to information transfer, over which the resources are shared by multiple users.

X.21. A CCITT recommendation that specifies the interface between a DTE and a DCE for synchronous operation over public circuit switched data networks.

X.25. A CCITT recommendation that specifies the interface between a DTE and a DCE for operation in virtual circuit mode over a packet switched network.

List of symbols and abbreviations

ABM	Asynchronous Balanced Mode
ADPCM	Analog-to-Digital Pulse Code Modulation
B-ISDN	Broadband Integrated Services Digital Network
BRI	Basic Rate Interface
CCITT	Consultative Committee for International Telephone and Telegraph
CCSN	Common Channel Signalling network
CITAM	X.21-to-ISDN terminal adapter, manufactured by Controlware Gmbh.
CPI	Customer Premises Installation
CRC	Cyclic Redundancy Check
DCE	Data-Circuit Terminating Equipment
DLR	Deutsche gesellschaft für Luft- und Raumfahrt
DS	Digital Section
DSS1	Digital Subscriber Signalling System Nr. 1
DTE	Data Terminal Equipment
DTS	Digital Transmission System
ESA	European Space Agency
ESOC	European Space Operations Centre
ESTEC	European Space Research & Technology Centre
ET	Exchange Termination
ETSI	European Telecommunications Standards Institute
EURIE	Pan-European ISDN conference to show availability of ISDN throughout Europe. This conference was held between 14 and 16 December, 1993.
HDB3	High-Density Bipolar Code of Order 3
IEN	Interexchange Network
ISDN	Integrated Services Digital Network
ISO	International Organization for Standardization
ISPBX	Integrated Services Private Branch Exchange
ITU	International Telecommunications Union
ITU-TSS	ITU Telecommunications Standardisation Sector
IDN	Integrated Digital Network
LAN	Local Area Network
LAPD	Link Access Protocol D Channel
LET	Logical Exchange Termination
LT	Line Termination
MAMI	Modified Alternate Mark Invert
MOU	Memorandum of Understanding
N-ISDN	Narrowband Integrated Services Digital Network
NLR	Nationaal Lucht- en Ruimtevaartlaboratorium (National Aerospace Laboratory).
NT	Network Termination
PABX	Private Automatic Branch Exchange
PCM	Pulse Code Modulation
PDH	Plesiochronous Digital Hierarchy
PRI	Primary Rate Interface. The user-to-network interface for the ISDN synchronous channel structure consisting of 23 (US and Japan) or 30 (Europe) B channels and

	one D channel.
PSTN	Public Switched Telephone Network
SAN	Subscriber-Access Network
SCA	Synchronous Channel Aggregation
S_0	The CCITT I.420 'S' reference point associated with the basic rate interface.
S_{2m}	The CCITT I.421 'S' reference point associated with the primary rate interface.
SS#7	Signaling System No. 7
TAXI	ISDN inverse multiplexer, manufactured by Controlware Gmbh, with the capability of ISDN dial up and multi-channel synchronization of time-delayed B channels.
TE	Terminal Equipment
TIA/EIA	Telecommunications Industries Association/ Electronics Industries Association
UNI	User-to-Network Interface
WAN	Wide Area Network

Chapter 1. Introduction

1.1 Datacommunications and ISDN

The present state of communication technology is characterized by two features, namely the digital representation of all signals transmitted and processed, irrespective of information type - voice, text, data, or images - and the integration of systems and services, this integration being only completely possible using digital technology. The boundaries between switching and transmission are shifting, and functions are being redefined and redistributed between terminals and communication networks. Multiservice terminals - unlike telephones, teleprinters, video data terminals - are designed to handle more than one information type. Lastly, the communication network allows voice, text, data and video information to be transmitted on the same circuit; the user obtains access to this network via a non-dedicated "communication socket".

The essential features of this Integrated Services Digital Network (ISDN) have been standardized over the last 13 years by experts from all over the world under the aegis of the Consultative Committee for International Telephone and Telegraph (CCITT), today called ITU-TSS, International Telecommunications Union, Telecommunications Standardisation Sector, the international standardizing body of the carriers of public communication networks. All the leading network carriers are working towards ISDN implementation because of the substantial benefits it will offer to users, network carriers and manufacturers alike:

- Users will obtain additional and advanced services, most of them designated to cater for the growth in non-voice traffic. The ISDN subscriber access will also enable users to operate existing systems more cost-effectively than via various dedicated networks.

- Network carriers will benefit from the universal network idea that allows them to introduce new services and supplementary services without large specific capital investment. Economic advantages will also flow from service integration, especially at subscriber line level, and from standardization of operation and maintenance for the range of services provided.

- Manufacturers will welcome above all the standard network access, enabling one terminal to access several communication services, i.e. opening a wider market.

The driving forces behind the ISDN are thus cost-effective communication in various information types via the same subscriber line, and sophisticated new services and supplementary services, especially in the increasingly important non-voice area.

1.2 The European Space Agency (ESA)

1.2.1 ESA - European Space Agency

The ESA organisation was founded in May 1975 as the outcome of the merge between the European Space Research Organization ESRO and the European Launcher Development Organization ELDO. At the moment thirteen member states contribute financially and technically to all the agencies basic activities (studies on future projects, technological research, scientific satellites). In addition, specific programmes concerning telecommunications, earth observation, space transportation systems (Ariane), space stations and platforms (Spacelab, Eureca) and microgravity research receive the contribution of those of the member states wishing to participate.

ESA now employs 2040 staff members from 15 different nationalities. In addition thousands of scientists and engineers from universities and industry give indirect support.

The main objective of the organization is the promotion and coordination of the space research in Europe for exclusively peaceful purposes. The activities of the agency have already resulted in the launch of many satellites for telecommunications, earth observation and scientific purposes. Currently, research is done for several big projects such as the heavy lift launcher Ariane 5.

1.2.2 ESTEC - European Space Research and Technology Centre

ESTEC is located in Noordwijk, the Netherlands and is the largest establishment of ESA. Its principal functions are:

- Project management of space programmes concerning science, communications, earth observation, microgravity, space stations and platforms.

- Execution of studies for future satellite programmes.

- Execution of the Space-Science programme.

- Conception and management of the Space Technology Research Programme.

- Testing of spacecraft (from full systems down to individual components).

- Provision of specialized people and laboratory facilities in all technical space disciplines (e.g. structures, data handling, product assurance).

Approximately 1200 international ESA staff plus some 450 non ESA staffs are employed in ESTEC.

1.2.3 The ESA Computer and Network Operations Department

The Computer and Network Operation Department is part of the Directorate of Operations that gives support to all the other Directorates of ESA.

It is in the communications division of this department that the author has been working to complete his thesis at the Delft University of Technology.

1.2.4 ISDN in the European Space Agency

Flexible communications for the support of manned space flight and other ESA scientific projects are becoming increasingly important. The data distribution and telecommand connectivity, previously limited to specific ESA sites has to be extended to remote centres and user sites outside ESA. This requires extensive bandwidth which up to now has been provided via dedicated lines.

The concept of using leased lines has proven to be expensive and inflexible, as the implementation has long lead times and is fairly static in nature. The upcoming national and international ISDN-based switched digital networks seem to offer a flexible and economical alternative for temporary wide area connectivity to remote user sites.

ESA is evaluating the possibilities of ISDN for datacommunications. For this purpose, a testbed has been set up at ESTEC. The author has been working on this test facility.

1.3 Outline of the report

The intent of this report is to describe the research that the author has done on ISDN at the European Space Research & Technology Centre (ESTEC).

Chapter two gives an introduction to ISDN. It describes how the ISDN has been evolved from its precedessors, it describes the architecture and protocols of the ISDN and it concludes with specific information on the development of ISDN within Europe.

Chapter three deals with the availability and quality of international ISDN lines. For specific space projects, data has to be transported between different ESA sites in Europe. Up to now, leased lines have been used to serve this purpose. Although leased lines are inflexible and expensive, they have proven to be a reliable way of transmission. In considering a change to ISDN connections, extensive tests had to be done on these connections concerning reliability and performance. The chapter gives the results of these tests and a comparison to the performance figures as specified by CCITT in 1988.

Chapter four contains a cost analysis between ISDN and digital leased lines.

In chapter five, the emphasis is put on the applications that can be used over the ISDN. Typical applications evaluated are the interconnection of geographically spread LANs, dial-in connectivity to LANs from a remote single user workstation (to simulate homeworking), videoconferencing over ISDN lines using inverse multiplexers, PC-to-PC applications such as file transfer, interactive working, point-to-point videoconferencing, and fax Group IV [1].

To evaluate the results of the performance tests of ISDN applications, chapter six gives a theoretical analysis of network performance over ISDN circuits and transparent file transfer over ISDN lines.

Chapter 7 contains conclusions from the research done and recommendations for future activities.

Appendix A contains the detailed test results of tests described in chapter three.

Appendix B gives a description of all network protocols that have been used over the LAN-to-LAN links that have been tested.

[1] Fax Group IV is a new facsimile standard defined to be used over 64 kbit/s digital channels, offering higher quality (up to 400 DPI) and higher speed than the existing fax Group III.

Chapter 2. Introduction to ISDN

2.1 Evolution from PSTN to ISDN

2.1.1 The user and the telecommunications network

The telecommunications network is originally designed to deliver a limited number of services (mainly telephony), using a limited number of technologies. Gradually this has changed into a network with many services, using various different technologies. It becomes harder to define the borders of the network. There are different means of integration of infrastructure and end-equipment. The used technologies keep changing. The infrastructure for example has changed from open, unamplified lines to optical fibres during the last twenty years.

The result of this is, that the user can concentrate on the service instead of the infrastructure formed by the network. The service has roughly been the same; it is the technology that has changed. A good example of this is the service 'telephony' which did not really change during the last 100 years, while the used technology has changed a lot. The service provider, usually the PTT, has been confronted with a situation of competition instead of monopoly during the last decade. Therefore they concentrate more on the customer, i.e. the services instead of the infrastructure. Important requirements from the users to the services that are delivered via the network are:

- uniform standardised access
- application-support
- high speed and high quality
- connectivity

Simultaneously a good price-performance ratio is expected. These requirements can be fulfilled by the Integrated Services Digital Network by:

I (Integrated)	integrated and standardised access
S (Services)	standardised services
D (Digital)	end-to-end digital connections
N (Network)	

ISDN has evolved from the digitalisation of the telephone network. The most important service is still telephony, using a worldwide network with more than one billion connections.
Besides telephony, facsimile over the telephone network has grown recently. With the ISDN, both services can be more efficient with a higher quality. The fastest growing service however is datacommunications.

2.1.2 Tele-informatics and ISDN

The eighties have shown an evolution of tele-informatics or telematica: the coupling of data processes, using a telecommunications network. Speech and data will appear together more often. Signal transport is handled more digitally: Speech, picture and data are hard to distinguish within the communications network. The confluence of the voice and data world and the increasing information exchange over large distances result in a need for integrated digital networks.

2.1.3 The development of ISDN

Since 1950, public telephone exchanges have been automated. Telex traffic got its own network with separate exchanges. In the beginning of the eighties, PTTs started delivering a separate network for datacommunications. In the Netherlands this was the Datanet 1.

Although for the users more different networks became available, the PTTs were integrating the transmission media for these networks. A key role in this process is digitalisation. With the availability of digital technology, very reliable and compact transmission and switching circuits can be realised.

Analog signals are transmitted digitally using Pulse Code Modulation. Within the telephone network, different signals are being multiplexed to create an aggregated bitstream. Until 1975, the maximum bandwidth for these bitstreams was 2 Mbit/s in the Netherlands. With the availability of optical transmission systems, more bandwidth became available, up to several hundreds of megabits per second.

Since 1984, the Dutch PTT has been implementing digital switches. At the end of 1994, all switches in the Netherlands will be digital [1]. With this network, it is possible to create an Integrated Digital Network (IDN) with integration of transmission and switch systems. There is a separate signalling network between the switches, the well-known CCS7 or C7 (Common Channel Signalling System no. 7). C7 also transports information about network management.

ISDN is being developed in stages, with the initial emphasis on the so-called narrowband ISDN, which is designed to provide facilities for the transmission of information requiring bitrates of less than several megabits per second. Narrowband ISDN can be implemented with relatively minor changes and additions to the existing circuit and packet switched network transmission circuits and switches. The design of the N-ISDN is based upon the use of the existing twisted copper pairs from the PSTN. The length of the local cable has a maximum of about 10 km. The length and quality of the cables are limiting factors to the throughput that can be achieved. Experiments have proven that a throughput of 200 kbit/s full-duplex is possible on the existing copper cabling [2]. With 64 kbit/s as standard coding for speech (4 kHz bandwidth including guard bands, sampling rate of 8000 Hz) there is bandwidth for more than one channel. For the standard so-called Basic Rate Access two 64 kbit/s data channels are available as well as one 16 kbit/s signalling channel. In addition there is the Primary access which offers thirty 64 kbit/s channels and one 64 kbit/s signalling channel. Primary access makes use of two metallic pairs.

In contrast, the so-called broadband ISDN (B-ISDN), which is intended to offer services at bit rates up to several hundred megabits per second, depends on the large-scale availability of wideband transmission media such as optical-fibre cables in both the interexchange and local exchange plant and the deployment of very high speed switching fabrics. Its full-scale implementation will therefore require much more time than for narrowband ISDN. Current expectations are that no significant broadband impact will occur before the late-1990s.

2.1.4 Benefits of the ISDN

From the point of view of the network, there are three key note advantages:

- A network configuration based entirely on digital transmission and switching offers the reduced cost, low power consumption and easier maintainability that are general characteristics of digital equipment.

- The high level of functionality of an ISDN and the integration of multiple connection types, services and user applications on a single network is expected to result in a higher utilisation of the network's resources, with a corresponding increase in revenues.

- ISDN provides a variety of data communications services over a digital subscriber line.

From the point of view of the customer, there are also several advantages:

- Quality. The transfer of digital information suffers much less from noise on the line than analog information.

- Speed. On the PSTN, datacommunication is only possible with the use of modems that limit the speed to 9600 bits/s. The ISDN offers two (or 30) 64 kbit/s lines. This increased speed opens the way to new applications like videoconferencing. Besides the higher dataspeed, also the time to establish a connection is reduced to some two seconds (for national connections).

- Integrated access. Via the ISDN, there is connectivity to all public networks like the PSTN, the public packet switched network, telex network etc.

- Standardisation. ISDN is a worldwide development. In different standardisation committees, the ISDN has been defined into details, to create an international compatible network. However, still a lot of different protocols are in use, but in Europe the so-called Euro-ISDN is being implemented which offers guaranteed international connectivity and equipment portability.

2.1.5 ISDN applications

The most important application on the narrowband ISDN will be telephony. Included here are the transfer of standard digital voice at 64 kbit/s, compressed digital voice at various submultiples of 64 kbit/s, high-fidelity voice, and a host of new and traditional supporting call control procedures. Among these are calling party identification, call forwarding and call charging indication.

The datacommunications requirements of the modern office lead to applications like text and graphics message communications, including telex, teletex, videotex, electronic mail, facsimile and videoconferencing. A very important area is also the interconnection of LANs, remote access to LANs and communication between PCs and workstations.

2.1.6 ISDN standardisation

The fundamental concept of a public telecommunications network based on digital transmission of information and offering an integrated access to multiple network connections and services has evolved largely through the work of the CCITT (Consultative Committee for International Telephone and Telegraph, today called ITU-T or ITU-TSS (International Telecommunications Union, Telecommunications Standardisation Sector). Starting in the late 1970s, Study Group XVIII, in close cooperation with Study Groups VII and XI, was given the task of developing the recommendations that define the concepts and principles of the ISDN and specify its service capabilities, network features, user-to-network interfaces, internetworking interfaces, and maintenance aspects. All these recommendations are contained in the I-series recommendations approved by the IXth Plenary Assembly of the CCITT in November 1988 and published in Volume III of the Blue Books with some updates after, especially concerning supplementary services. Various other organizations have made major contributions to the standardization of the ISDN. Most important of these for Europe is the ETSI (European Telecommunications Standards Institute), that developed standard versions of two types of user-to-network interfaces known as *basic rate interface (NET3)* and *primary rate interface (NET5)*. Since there are many organizations involved in the specification of ISDN, many options incorporated in the specifications, and divergent economic and technical interests of public and private network providers, no single globally compatible version of ISDN exists. However, the European public network providers have agreed to harmonize their versions of the ISDN following the ETSI specifications. More information on this so-called Euro-ISDN is given in chapter 2.7.

2.2 ISDN architecture and protocols

The ISDN can be divided into three major parts shown in figure 2.1.

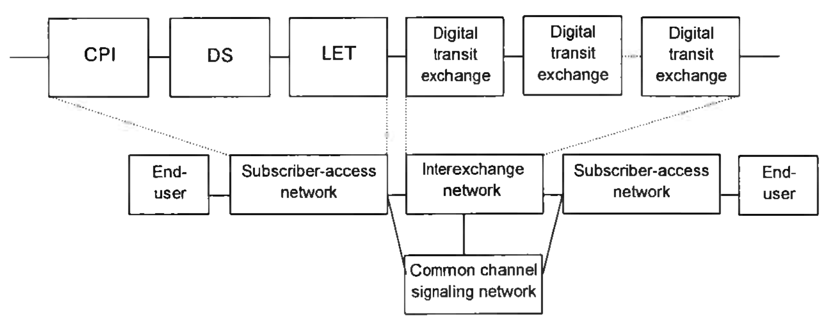

Figure 2.1 *ISDN decomposition*

Central to any ISDN is the so-called interexchange network (IEN), which consists of the physical and logical components of the backbone transmission network, including several network transit exchanges and the transmission trunks connecting these exchanges. As far as the end-user is concerned, the IEN's main purpose is to provide physical and logical transmission and switching facilities across which user information flows may be conveyed.

Superimposed on the IEN and interacting with it is the common channel signalling network (CCSN), which combines the functions required for the control, management, and maintenance of the ISDN. It provides the physical and logical transmission capacity for the transfer of connection control signals between the components of the IEN, for the management and allocation of network resources, and for the performance of maintenance functions.

The last major part of the ISDN -the subscriber-access network (SAN)- consists of the part of the ISDN between the end-user or subscriber and the IEN and CCSN. It can be divided into three components, namely the customer-premises installation (CPI), the digital section (DS), and the logical exchange termination (LET). The CPI combines those aspects of the SAN that are directly under the control of the subscriber. The DS consists of the local loop or subscriber-access line and the physical line termination equipment. It provides the transmission capacity for carrying information between the local exchange and the customer premises. The purpose of the LET is to terminate these transmissions in a logical sense. To specify the physical and logical properties of the IEN, CCSN, and SAN, and to aid in the development of standard ISDN implementations, it is useful to define a decomposition of the ISDN's total capability. This is accomplished by dividing it into an abstract topological arrangement of groups of functions that interact with each other across so-called reference points. Such a decomposition is known as a reference configuration. The functional groups consist of certain combinations of physical and logical functions that are required for the transmission and control of the signals across the IEN, CCSN, and SAN. They do not necessarily correspond in a one-to-one manner to specific physical devices. In particular, individual functions in a group may be implemented in one or several pieces of equipment, and one piece of equipment may perform functions from more than one group. The reference points define conceptual points of demarcation between pairs of functional groups. They may correspond to physical interfaces between separate pieces of equipment implementing the two functional groups, or virtual interfaces in cases where both functional groups are found in the same equipment.

2.2.1 The ISDN subscriber-access network reference configuration

Figure 2.2 shows the reference configuration for the subscriber access network (SAN).

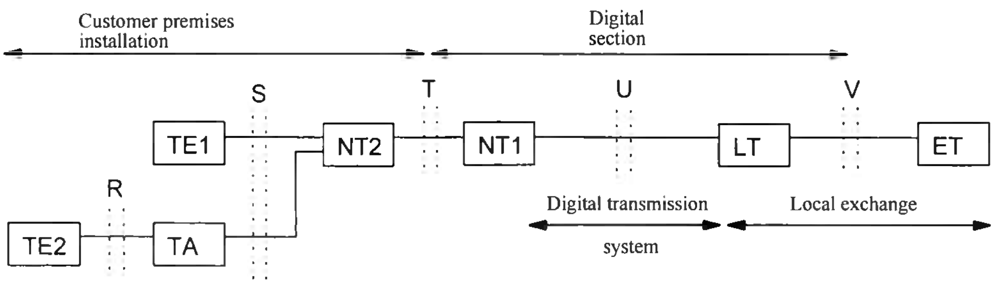

Figure 2.2 *Subscriber-access network reference configuration.*

According to the figure, the configuration is decomposed into six general groups of functions that together describe the capability of the SAN to carry information between the end-user and the interexchange network.

The exchange termination (ET) provides the functions necessary for the logical attachment of the SAN to the interexchange network (IEN).

The physical aspects of terminating the digital transmission system on the network premises are contained in the line termination (LT). Specific functions include the feeding of power across the digital transmission system (DTS) to the customer installation, fault location through the transmission of loopback signals, the generation and regeneration of baseband signals and the conversion from one baseband code to another.

The network termination 1 (NT1) forms the intermediary between the network and the equipment on the customer premises in a physical sense. Its main purpose is the termination of the DTS on the customer-premises side. It also performs conversions between the electromagnetic signals generated and received by the customer equipment and those transmitted over the subscriber line, transmission timing and the feeding of power to the customer-premises equipment.

The functions of the network termination two (NT2) usually correspond to those of a switching device located on the customer premises; they extend over the physical, link and network layers of the OSI architecture. The key features are multiplexing, concentration and switching of multiple information streams, and protocol handling at layers two and three. In actual implementations the entire capability of an NT2 may be incorporated in a PABX or LAN.

The terminal equipment (TE) corresponds to customer terminals such as digital telephones, integrated voice/data terminals, pc's and workstations. It incorporates the functions required for protocol handling, physical connection to the equipment associated with the NT1 and NT2 functions, and maintenance. Two types of TE's are identified. The terminal equipment one (TE1) functions completely in conformity with the ISDN specifications. A terminal equipment two (TE2) corresponds to the existing data communications equipment that does not follow the ISDN conventions, for example devices that use RS-232C, X.21 and X.25 protocols.

The terminal adapter (TA) provides the functional capability of converting the layer 1,2 and 3 protocol functions of a TE2 into those of a TE1, so that the combination of a TE2 and a TA appears as a TE1. TAs are intended to allow the attachment of conventional data communications equipment to an ISDN and may therefore eliminate the need for modems.

Reference points

The functional groups defined in the subscriber-access network reference configuration interact with each other across five distinct reference points that correspond to physical or virtual interfaces between the groups.

The S reference point defines the demarcation between the TE1 or TA on one side and the NT2 on the other.

The T reference point marks the interface between NT1 and NT2 and allows the separation of these functions into different groups. This creates the possibility that the NT1 and NT2 functions are provided by different entities.

The R reference point between the TE2 and TA accommodates existing non-ISDN terminal equipment designed to other ITU-TSS standards such as X.25 or X.21, and interface standards not conforming to ITU-TSS recommendations.

The U reference point defines the interface between the DTS and the NT1. In some implementations of ISDN it serves as the boundary between the network provider and the subscriber. However, the U reference point is not specified by ITU-TSS and therefore the properties of the signals traveling across it may differ from country to country.

The V reference point separates the physical and logical aspects of terminating the subscriber-access network on the network premises. Different versions have been defined to provide different individual user-to-network access capabilities, such as basic-access and primary access. Like the R reference point, the V reference point is not specified by ITU-TSS.

Within the scope of the general subscriber-access network reference configuration, many different physical arrangements of the physical groups and reference points are possible. Some of the most important cases are shown in figure 2.3.

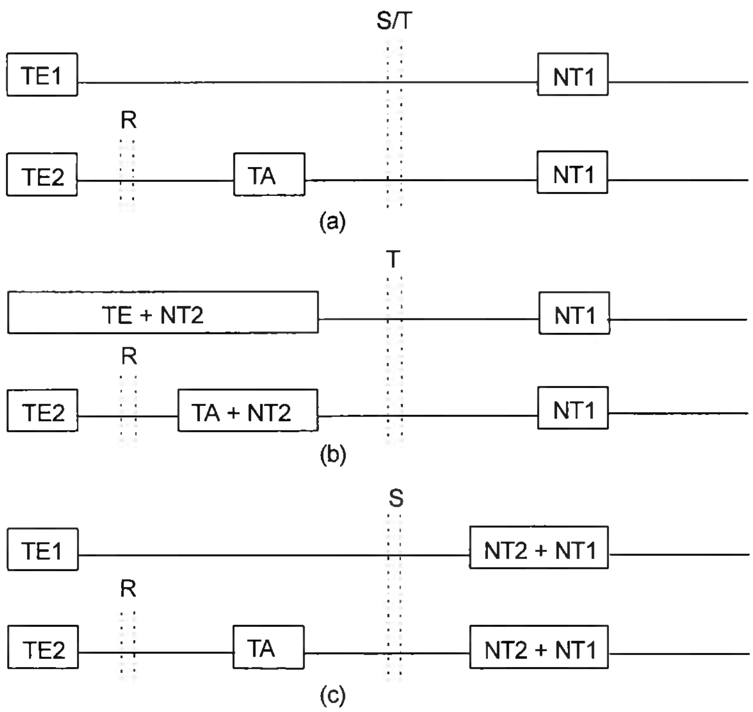

Figure 2.3 *Point-to-point subscriber-access network reference configurations*

From (a) it follows that the physical and electromagnetic characteristics of the interface at S and T have to be identical.

2.3 The ISDN communications architecture

2.3.1 User-to-network interface types

The ISDN is completely described to the end-user by specifying the following characteristics:
- The mechanical and electromagnetic properties of the devices in which the functional groups are implemented and the interconnecting media by means of which these devices are attached to each other.

- The signalling structures and logical protocols of interaction employed by the functional groups in communicating with each other across the S or T reference points.

- The operational, performance and maintenance properties of the functional groups and interconnecting media.

One of the most important objectives of the ISDN development is to standardize these characteristics of the user-to-network interface to obtain several major benefits. First, individual ISDN components can be designed and produced to well-known and universally accepted specifications by multiple suppliers, resulting in economical implementations and widespread availability. Second, end-user devices can evolve independently from the network equipment and network configuration. Finally, end-user devices may be relocated from one interface point to another on either the same network or a different network, without any concerns for compatibility.

Considering that the standardization process is very complex and difficult, it follows that to balance the requirements of efficiency, universality, flexibility, low complexity, and low cost on the one hand, and the wide variety of existing and evolving devices and applications on the other, several distinct standard user-to-network interfaces should be defined. These should offer the potential of matching the user application to the network capability optimally. Three of these - with increasingly higher levels of capability- are identified in figure 2.4, with examples of the intended applications. They are known as basic-access user-to-network interface, primary-access UNI and broadband-access UNI.

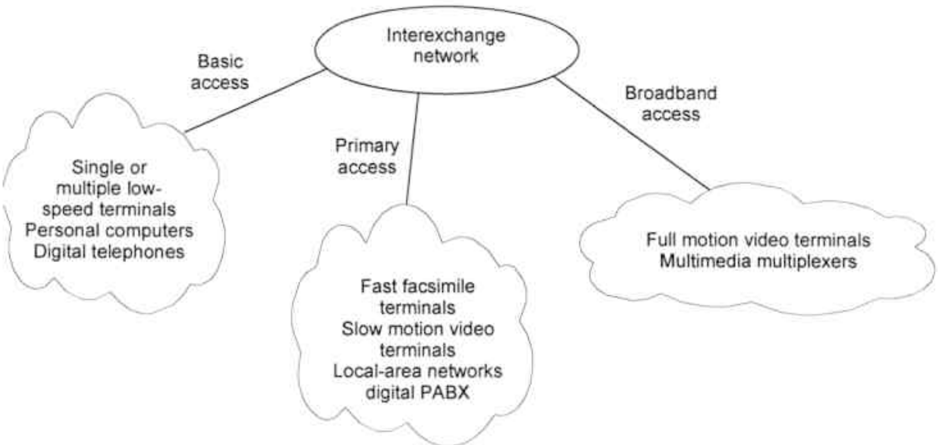

Figure 2.4 *Basic access, primary access and broadband access*

The basic access UNI is primarily designed for the kinds of devices typical of the residential or small-business user. The peak data rates generated in applications involving these devices typically are below 100 kbit/s.

The primary access UNI is designed for higher data rates up to 2 Mbit/s. Its most important application is probably in the connection of a digital PABX to the terminals on the one side and the network on the other.

The broadband access UNI provides the capability required for the transmission of television and very high speed file transfer among other things. Here the aggregate data rates may extend up to several hundred Mbit/s.

2.3.2 Channel structures and component channels

As mentioned in the previous section, there are three types of synchronous channel structures. A more detailed description of these is given below. The combining of the component channels into a synchronous channel structure is normally accomplished by synchronous time division multiplexing. The channel structure consists of a continuous and periodic sequence of frames separated by framing channels and synchronously transmitted across the S or T reference points of the UNI at a certain transmission rate.

B channel

The B channel operates at a synchronous data-rate of 64 kbit/s in full duplex mode. Its primary purpose is to carry information between a specific pair of end-users across the S or T reference points of the UNI, without imposing any restrictions on the binary representation of the data. Three classes of end-user information are generally considered appropriate for transmission over the B channel:

1. Digital voice encoded at 64 kbit/s, using pulse code modulation techniques according to CCITT Recommendation G.711.

2. Data streams corresponding to the synchronous user classes of service 3-12 and 30, as defined in CCITT Recommendation X.1 [3], with data rates of 0.6, 1.2, 2.4, 4.8, 9.6, 48, and 64 kbit/s. These streams may exist in the B channel either individually or in combination with other data streams at the X.1 rates.

3. Digital voice encoded at bit rates less than 64 kbit/s, either alone or in combination with other voice or data streams.

In the second and third class the aggregate data rate of the combined bit streams must be adapted to equal 64 kbit/s.

D channel

Two types of D channels, operating at synchronous data rates of either 16 kbit/s or 64 kbit/s in full duplex mode, have been defined. Their major function is to carry the signalling information for the control of circuit switched connections involving one or more B-channels between the user and the network. During periods when it is not needed for this purpose, a D channel may also be used for the transmission of user-to-user signalling information, low bit rate packet switched data, and telemetry signals, using statistical multiplexing techniques and priority access for call control signals to resolve any contention.

H channel

H channels are designed to carry the types of user information requiring data rates over 64 kbit/s and up to several 100 Mbit/s. Four types of H channels have been defined. They are summarized in table 2.1 [2].

Component channel	Data rate kbit/s	Multiple of B-channel rate	Multiple of H0-channel rate
H0	384	6	1
H11	1536	24	4
H12	1920	30	5
H21	32768	512	-
H22	44160	690	115
H4	135168	2112	352

Table 2.1 *H channel data rates.*

2.3.3 The basic rate channel structure

The first of the synchronous channel structures, corresponding to basic access, is referred to as the basic rate channel structure and consists of two B channels and one 16 kbit/s D channel (2B+D), for an aggregate data rate of 144 kbit/s. The B channels are used simultaneously and independently of each other, in either different connections or in the same connection, and carry the end-user data. Call control information for both B channels is transmitted on the logically separate D channel, creating an out-of-band signalling arrangement and allowing the entire B-channel capacity to be used for the transmission of user information.

The primary intent in defining this arrangement of channels is to create the possibility of two simultaneous and independent voice or data calls, either by the same or by different terminals, over a single UNI. Alternately, one B channel may be used to carry a voice call, while the other is engaged in the transmission of data, facsimile, low-speed video, or any other signal with a data rate that does not exceed 64 kbit/s.

2.3.4 The primary rate channel structure

The second type of synchronous channel structure, the primary rate channel structure, is associated with the primary-access interface. It is designed to accommodate applications requiring more than two simultaneous connections and/or data rates over 64 kbit/s. Two distinct configurations, based on the two ITU-TSS primary PCM transmission standards, have been defined.

Generally, in the 1544 kbit/s version (used in North America and Japan) any combination of B and H0 channels, with or without a 64 kbit/s D channel, is allowed, as long as the aggregate data rate does not exceed 1536 kbit/s.
In the 2048 kbit/s version (used in Europe) of the primary rate channel structure, the general arrangement is given by

$n\mathrm{H0} + m\mathrm{B} + \mathrm{D}$

where the integers n and m are constrained by the relations
$0 \leq n \leq 5 \qquad 0 \leq m \leq 30 \qquad 6n + m \leq 30$

The most important of these arrangements combines thirty B channels and one 64 kbit/s D channel (30B + D) into one data stream at an aggregate rate of 1984 kbit/s. Another important channel structure combines five H0 channels and one D channel, for the same aggregate data rate of 1984 kbit/s as before.

2.3.5 The broadband channel structure

The specification of synchronous channel structures for the broadband-access is still somewhat tentative, and considerable uncertainty exists about whether such channel structures are even appropriate in the broadband environment. Although several options are under consideration, no specific details will be given.

2.4 The ISDN user-to-network layer 1

Layer 1, the physical layer, is responsible for the bidirectional transmission across the R, S, and T reference points of the electromagnetic signals that represent the logical quantities in the B and D channels.

2.4.1 Layer 1 specifications for basic access

2.4.1.1 Wiring configurations

The devices containing the communicating TE and NT functional groups must be interconnected across an S or T reference point by an electromagnetic medium. With respect to the 2B + D basic-access channel structure, the medium must accommodate an aggregate data rate of 144 kbit/s, with an overhead incurred by the peer-to-peer protocol. The CCITT chose to specify the use of two unshielded metallic pairs, one for each direction of transmission. Two configurations, shown in figures 2.5 and 2.6, are defined in Recommendation I.430 [4].

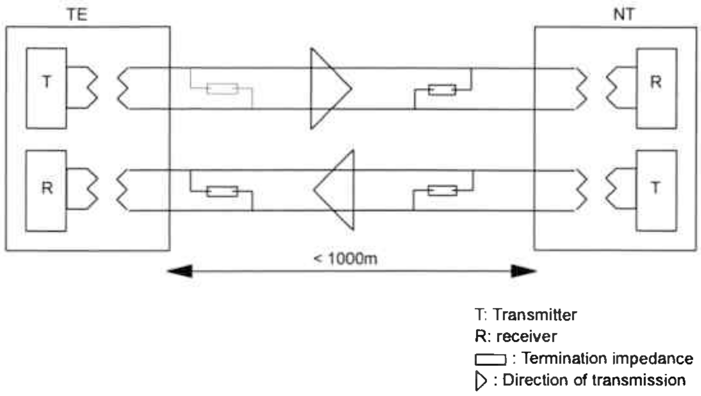

Figure 2.5 *Point-to-point wiring configuration*

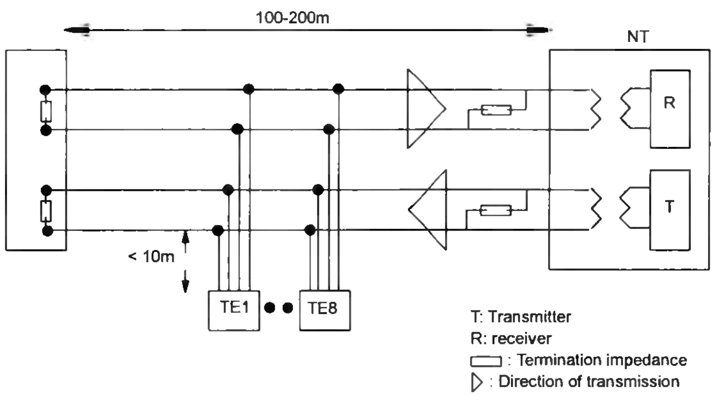

Figure 2.6 *Short passive bus wiring configuration*

ITU-TSS has specified the use of two unshielded metallic pairs, one for each direction of transmission, to physically carry the information flows in the channel structure across the S and T reference points. Typical values of these metallic pairs are given in table 2.2.

Propagation delay [μs/km]	7
Characteristic impedance [Ω]	100
Attenuation at 100 KHz [dB/km]	8

Table 2.2 *Characteristics of wires used for ISDN basic access*

The pairs may be arranged in a point-to-point configuration, connecting one active transmitter to one active receiver at each end of the cable, or a point-to-multipoint topology, which allows the connection of several simultaneously active devices at one end to one active device at the other.

2.4.1.2 Frame structure

The bitflows are organized into recurring frames of 48 bits, with each frame transmitted synchronously in 250 μs, for an effective transmission rate of 4000 frames/s, or 192 kbit/s. Of this total bit rate, 144 kbit/s correspond to the information obtained from the higher layers of the protocol architecture, with the remaining 48 kbit/s constituting layer 1 peer-to-peer protocol overhead and control information. The binary structure of the frame is shown in figure 2.7 [4].

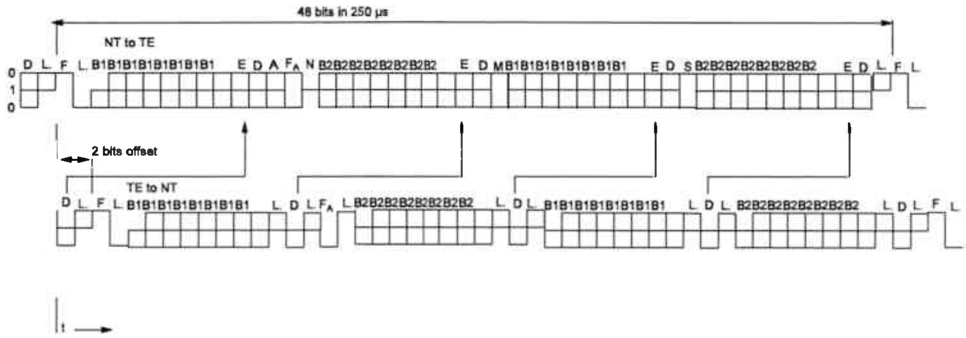

Figure 2.7 *Structure of S/T basic-access frame.*

The L bit's logical value is chosen so as to create a line signal with a zero average voltage. The B1 bits form the first 64 kbit/s B-channel, the B2 bits the second B-channel, and the D-bits the 16 kbit/s D-channel. The F and F_A bits are used to provide frame synchronization. The E-bit contains a copy or echo of the D-channel bit most recently received from the other side. The S-bit is unspecified and its value is set to binary 0. The N bit is used as a part of the framing procedure and the A bit plays a role in the activation and deactivation procedure of the interface. The M bit, known as the multiframe bit, allows the identification of a group of frames that form a multiframe.

2.4.1.3 Line code

A pseudoternary line code, also known as bipolar or modified alternate mark invert (MAMI), is used to physically represent the bits in a frame. Figure 2.8 gives an example.

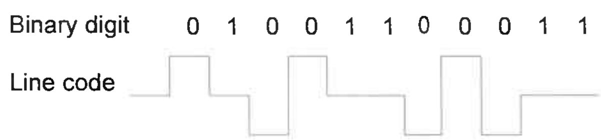

Figure 2.8 *Example of pseudoternary line code*

2.4.2 Layer 1 specifications for primary access

2.4.2.1 Wiring

Due to the difficulty of providing common timing to multiple terminals on a bus, only a point-to-point wiring configuration between the TE or NT2 and NT1 is provided. The cabling consists of two symmetric metallic pairs with a characteristic impedance of 120 Ω or two coax cables with a characteristic impedance of 75 Ω as specified in the G.703 [5] recommendation of ITU-TSS. However, the symmetric metallic pair is preferred.

2.4.2.2 Frame structure

The frame structure for the 2048 kbit/s version of the primary-access interface is shown in figure 2.9 [5]. It consists of 32 groups of 8 bits each, numbered from 0 to 31. The frames are transmitted in 125 μs, for a transmission rate of 8000 frames/s or 2048 kbit/s. Each 8-bit group is transmitted in a time slot of 3.91 μs, yielding a data rate of 64 kbit/s per time slot. The D-channel information is assigned to time slot 16. Time slots 1-15 and 17-31 may be allocated to the B-channels. Time slot 0 provides synchronisation and multiframe information and a CRC (optional).

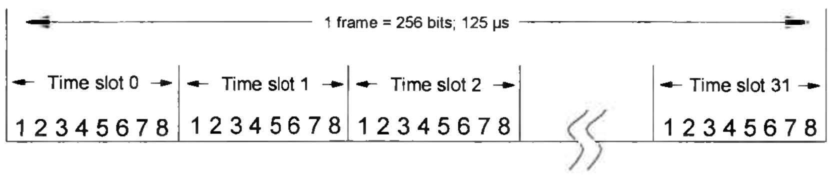

Figure 2.9 *Structure of S/T 2048 kbit/s primary-access frame*

2.4.3.3 Line code

On the 2048 kbit/s version of primary access, a high-density bipolar code of order 3 (HDB3), in which each block of four consecutive 0s is replaced by either the sequence 0 0 0 V or by the sequence 1 0 0 V. V represents a violation of the bipolar coding rule, and the particular sequence is determined by the requirement that successive V pulses alternate in polarity.

2.5 The ISDN user-to-network layer 2

Layer 2, the data link layer, provides a reliable and efficient information transfer service. The transfer taking place across the S or T reference point is either between a TE on one side and the NT2 or ET on the other or between the NT2 and ET. The NT1 is not involved in this exchange, since its functions extend only to layer 1. CCITT has specified a set of services and procedures known as the Digital Subscriber Signalling System No.1 (DSS 1) [6] that govern many of the interactions across the subscriber-access network of the ISDN. ETSI has used these specifications to define its E-DSS1, which is the base for the so-called Euro-ISDN. E-DSS1 will guarantee one uniform compatible ISDN signalling scheme within Europe.

The protocol used is based on the asynchronous balanced mode (ABM) of the high-level data link control procedures developed by ISO. The parts applicable to the information flows carried on the D-channel are known as LAPD [7]. Currently, LAPD is only used on the D-channel. However, ETSI has made specifications to adapt LAPD to the B-channel.

The purpose of the ISDN data link layer is to convey in a reliable and efficient manner the layer 3 information flows and the layer management information flows between peer entities. Services offered by the data link layer entity to the layer 3 and layer management entities are requested and provided through the exchange of service primitives across the appropriate service-access points. These exchanges represent in abstract fashion the logical interactions between the data link layer entities, the layer management entities, and the entities in layer 3. Generally the four service primitives Request, Indication, Response, and Confirm are used.

2.6 The ISDN user-to-network layer 3

The function of layer 3 of the ISDN communications architecture is to establish, operate, and terminate network connections, transfer information flows across them, and control the invocation of supplementary services. For the transfer of user information, a connection usually consists of an association between pairs of B-channels, one on each side of the connection, and either a circuit-switched or packet-switched connection over the interexchange network. Although a number of existing protocols such as the data transfer part of CCITT recommendation X.25 [8] can be adapted to the control of the user information flows, to date no single unified ISDN protocol serving this purpose has been defined. Above the physical layer many types of user information flows are carried transparently between end-users and therefore do not require any layer 3 protocols across the user-to-network interface.

The control information flows needed to establish, maintain, and terminate circuit switched or packet switched connections between the users, and to provide control over supplementary services are conveyed across the UNI over the D channel. They require a packet switched connection through the common channel signalling network. The D channel, together with a packet switched connection over the CCSN, also provides for the transfer of the user-to-user signalling information flows. Finally, the D channel, in conjunction with either a packet switched connection over the CCSN or the packet switched interexchange network, may carry packet switched user information flows. The exchange of these three types of flows over the D channel is governed by layer 3 of the Digital Subscriber Signalling System No. I (DSS 1) [6] that has been specified by the CCITT.

During establishing and clearing of calls, the D-channel provides information elements containing the state of the call (e.g. call proceeding, setup, connect, etc.) and eventual error codes (e.g. call rejected, number changed, user busy, etc.).

The D-channel also provides the exchange of I.451 [9] messages, or supplementary services. Supplementary services are services supplementary to a basic telecommunications service. An example is the so-called CLIP: Calling Line Identification Presentation. The receiver receives the number of the sender through the D-channel as the sender calls him. For more information about the supplementary services, see chapter 2.7 about EURO-ISDN.

2.7 Euro-ISDN

2.7.1 Description of Euro-ISDN

The 1988 CCITT recommendations left many options for implementation of ISDN, especially on the level 2 and 3 layers of the D-channel protocol. Therefore, the implementation of ISDN has differed between countries. This has led to incompatibilities between services in different countries and between some equipment and some services. Euro-ISDN, the result of a major pan-European effort, is designed to be a common ISDN implementation and all European operators plan to implement it. Euro-ISDN has two key elements:

- The ISDN standards drawn up by ETSI.

- The Memorandum of Understanding (ISDN-MOU) on the implementation of a European ISDN service, concluded in 1989 [10].

The ISDN-MOU committed its signatories to open a core set of ISDN services, conforming to standards drawn up by ETSI, by December 1993 at the latest. The term 'Euro-ISDN' is used to describe this common ISDN implementation. The core set of services consists of the following components:

- *Two types of ISDN access*

1) ISDN basic access consisting of two 64 kbit/s B channels and one 16 kbit/s D channel (2B+D).

2) ISDN primary access consisting of thirty 64 kbit/s B channels and one 64 kbit/s D channel (30B+D). B channels are used for voice and/or data; D channels are used for signalling and, if needed, for packet-switched data communications by the user. (Note that in the USA and Japan, ISDN primary access presents 23 B channels to the user rather than 30.)

- *Two ISDN bearer services*

Bearer services are telecommunication services that provide the capability for the transmission of signals between user-network interfaces. They are defined at the interface between the terminal and the network The two circuit mode bearer services in the core set specified in the ISDN-MOU are:

1) 3.1 kHz audio: corresponds to the service currently offered in the PSTN. It provides for the transfer on a B channel of speech and of 3.1 kHz bandwidth audio information such as voice band data via modems

2) 64 kbit/s unrestricted: provides unrestricted information transfer over a B channel and may therefore be used to support a range of user applications.
In practice, operators are also implementing one other bearer service simultaneously: the 64 kbit/s service usable for speech. ISDN packet mode and other bearer services are also being developed, although none are included in the ISDN-MOU core set of services.

- Five supplementary services

Supplementary services may be provided in association with bearer services. Supplementary services cannot be offered to users on a standalone basis. As the name suggests, supplementary services provide extra facilities to users. Many supplementary services have been proposed in Europe. The ISDN-MOU includes the following in its core set of services:

1) Calling line identification presentation (often abbreviated to CLIP): indicates the number of the calling party to the called party.

2) Calling line identification restriction (CLIR): enables the calling party to prevent presentation of the ISDN number to the called party.

3) Direct dialling in (DDI): enables a user to call directly via a public ISDN a user on a private ISDN (for example an ISDN PBX) without attendant intervention.

4) Multiple subscriber number (MSN): provides the possibility of assigning multiple ISDN numbers to a single public or private interface.

5) Terminal portability (TP): allows a user to move a terminal from one socket to another within one given access installation during a call.

Most Euro-ISDN implementations are being introduced with a much larger set of supplementary services. The important ones include:

- User to user signalling (UUS): allows a user to send or receive a limited amount of user information over the ISDN D channel either before or during a call.

- Sub-addressing (SUB): in a clustered terminal configuration, allows specific terminal selection beyond the ISDN access number with a supplementary address added to the number.
- Closed user group (CUG): enables users to form a group to and from which access may be restricted.

Beyond access, bearer and supplementary services, one other set of ISDN services has been defined: teleservices. A teleservice is a telecommunications service providing the complete capability, including terminal equipment functions at the human interface, for communication between users according to protocols established by agreement between telecommunications operators. Teleservices for which definitions have been drawn up include: telephony (3.1 kHz and 7 kHz), telefax (Group 4), videotex and videotelephony. The ISDN-MOU core set of services does not include any teleservices.

2.7.2 Availability of Euro-ISDN

The following countries have signed the ISDN-MOU: Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, UK.

Almost all these countries will have fully commercial Euro-ISDN services widely available during 1994 [11]. In every country except Greece, the operators have defined a programme for rolling out Euro-ISDN throughout the country. The programme always entails the full Euro-ISDN core set of services being accessible at the customer's request from at least half of the country's telephone lines by the middle of 1995. For Belgium, Denmark, Germany, Ireland, Luxembourg, Portugal, and Switzerland, this milestone will have been reached by the middle of 1994. Most international Euro-ISDN connections will also be implemented within this timescale.

Migration to Euro-ISDN

All the countries listed above had a Euro-ISDN service, at least on a pilot basis, by the end of 1993 [11]. All are committed to a full commercial service, though the speed at which Euro-ISDN will become available nationwide varies between countries.

Across Europe as a whole, 47% of PSTN lines will be within reach of Euro-ISDN by January 1994, rising to 84% by January 1995.

The picture that emerges is of impressive progress towards the creation of a uniform digital communications platform within Europe. There remain, however, several barriers to be overcome before the goal is fully accomplished.
They include:

- Limited geographic availability. In the early years, Euro-ISDN will not be available in many areas of Europe. The numbers given show, however, that progress will be faster than many have expected.

- Limited international functionality. Although many international interconnections of Euro-ISDN services were in place by the end of 1993, some support initially only a limited set of Euro-ISDN supplementary services.

- Different tariff structures in different countries, with a divergence of approaches to the positioning of Euro-ISDN with respect to existing telecommunications services.

By April 1994 there were some 350,000 ISDN basic accesses and 40,000 ISDN primary accesses in Europe. Over 70% of the ISDN basic accesses, and over 40% of the primary accesses, are in Germany, where ISDN has proved extremely successful. Only in Japan, where the number of ISDN basic accesses has also passed 200000, has the ISDN take up been comparable.

Over 90% of the basic accesses in Europe, and over 95% of the primary accesses, are country-specific ISDN rather than Euro-ISDN, but all the operators plan to make Euro-ISDN available and expect all their existing ISDN customers to migrate to Euro-ISDN during the next four years.

2.7.3 Additional planned supplementary services and teleservices

The operators are planning to introduce more supplementary services in the period from 1994-1996. However, not all of these services will be provided by all of the operators. Some of the services may be delivered without charge, but for most of them an additional charge will be invoked. The supplementary services are:

- Connected line identification presentation (COLP): enables the calling party to identify the national ISDN number of the connected party, including direct dialling-in and multiple subscriber number digits.

- Connected line identification restriction (COLR): enables the connected party to restrict presentation of his national ISDN number.

- Advise of charge (AOC): provides the calling party with charging information, i.e. the tariff information at the start of the call, a running total during the call, or the total charge at the end of the call.

- Closed user group (CUG): enables users to form a group to and from which access may be restricted.

- Call waiting (CW): (basic access only) - notifies the user of an incoming call when the line is already in use. The user has the option to accept or to ignore the call within a limited period of time.

- Completion of calls to busy subscribers (CCBS): provides automatic redial of a call when the called number changes from busy to free.

- Conference call, add on (CONF): provides the ability to have a call between more than two parties in different locations.

- Meet me conference (MMC): allows advance booking of conference calls.

- Call forwarding, unconditional (CFU): enables the forwarding of incoming calls to another number; if required, the forwarding may only be applied to calls on the speech bearer service.

- Call forwarding, busy (CFB): enables the forwarding of an incoming call (all calls, or speech calls only) to another number if a busy signal is returned.

- Call forwarding, no reply (CFNR): enables the forwarding of an incoming call (all calls, or speech calls only) to another number if the call is not answered.

- Call deflection (CD): enables a user, at the moment of receiving a call, to choose whether to answer it or to forward it to another number.

- Freephone (FPH): allocates the call charges partly or fully to the called party.

- Malicious call identification (MCI): enables a user to request the network to register the source number of an incoming call for further identification.

- Sub-addressing (SUB): allows an individual terminal behind an ISDN access to be addressed with a supplementary number.

- Three party service (3PTY): enables the user to have a simultaneous or alternate three way call.

- User to user signalling (UUS): allows a user to send or receive a limited amount of information over the D channel. Three different levels are defined. With level 1, a message is sent at the start of the call, and the called party can read the message before answering the call.

- Call hold (CH): allows a user to put a call on hold, make another call and when that has been completed, resume the original call.

There are also some teleservices that will be launched between 1994 and 1996. These services are Telephony 3.1 KHz, Facsimile Group 4, Teletex, Telephony 7 KHz, ISDN syntax-based videotex and Video telephony.

2.7.4 Euro-ISDN international connections

The Memorandum of Understanding of 1989 stated that all the Euro-ISDN services should be interconnected by December 1993. However, this objective has only partially been fulfilled. At this moment, most of the European interconnections support the two bearer services of the MOU, and some of them also support the five supplementary services from the MOU. This incompatibility is due to the differences in the exchange-to-exchange signalling. Most of the implemented SS#7 signalling systems still use the so-called TUP (Telephone user part) to provide the basic call control functions for the establishment, operation, and clearing of connections over the inter-exchange network. In the near future, all operators will upgrade to the ISUP which stands for ISDN User Part. With ISUP, in principle all supplementary services will be supported, although there may still be limits because of various national implementations. The first version of ISUP is only meant to provide the stated functionality within Europe, whereas the version 2 will be a global standard. ISUP will be implemented eventually on all Euro-ISDN international links. As an example, from the Netherlands most international connections within Europe will support the ISUP-1 by the end of 1994, the remaining by end 1995, and also countries like USA and Australia. Japan, Hong Kong, Canada, New Zealand and Singapore can be reached via ISUP by the end of 1995 [11].

Chapter 3. Availability and quality of international ISDN bearer channels

To gather information about the reliability and interconnectivity of the ISDN networks within Europe, extensive tests have been done on ISDN connectivity between different European countries.

3.1 Scope of the research

Data connectivity between different ESA and non-ESA sites in the Netherlands and Germany using the narrow band ISDN 'bearer service' was tested during December 1993 and January 1994. An assessment of the availability in terms of call setup & cleardown and information transport efficiency was done. The investigations were done empirically by setting up connection between the different sites that demand connectivity using inverse multiplexers with basic rate (BRI/S_0) interfaces and primary rate (PRI/S_{2M}) interfaces. The availability and quality of connections between the following sites have been tested:

- ESTEC - ESOC (Noordwijk/NL - Darmstadt/D)
- ESTEC - DLR (Noordwijk/NL - Köln/D)
- ESTEC - NLR (Noordwijk/NL - Amsterdam/NL)
- ESOC - DLR (Darmstadt/D - Köln/D)

In February 1994, some additional tests have been done between ESTEC and ESOC on one side and ESA/CNES Toulouse and ESA/HQ Paris on the other side.

The results of these investigations were considered important for the IML-2 spacelab flight mission, as it is currently planned to use ISDN switched circuits instead of the traditional leased lines to distribute data across Europe, and to use Synchronous Channel Aggregation (SCA) techniques for video, voice and data services.

3.2 Objectives of the test program

The objective of the test program was to evaluate the performance of ISDN switched 64 kbit/s data circuits between several European sites.

The performance was measured in terms of call connectivity and dis-connectivity efficiency, circuit availability and error rate performance. It includes the assessment of synchronous channel aggregator devices to verify the suitability of ISDN 64 kbit/s circuits to carry multiplexed datastreams such as found with 384 kbit/s video conferencing systems.
Further, it explored the capability of CISCO multiprotocol routers to provide 128 kbit/s data connectivity between ESTEC and NLR as part of the Columbus simulation workshop requirements scheduled for February 1994 and May 1994.

3.3 Test description

A total of 7 configurations were defined. These provided for evaluation of the ESTEC local ISDN switch (Siemens HICOM 300 ISPBX) and the Dutch PTT connection to Leiden, and combinations of selected German sites. The test configurations assessed ISDN and terminal equipment capabilities and performance. Single and multiple 'B' channel calls were attempted, and if successful, bit error rate transfer efficiencies were measured. Time for connect and synchronization was measured. Finally, the call(s) were requested to clear and results of the test process noted. A typical test configuration is given below.

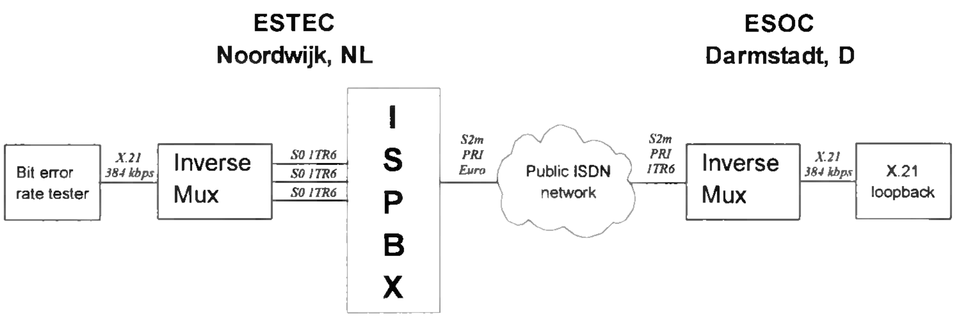

Figure 3.1 *Test configuration*

On the ESTEC side, there is a Euro-ISDN primary rate interface. The ISPBX is connected to this PRI. However, the ISPBX delivers the German 1TR6 protocol on its basic rate ports, which means that it does an internal translation. On the ESOC side, there is a 1TR6 primary rate interface. The inverse multiplexer is directly connected to this PRI. Inverse multiplexers usually have ISDN interfaces on one side and an X.21 port on the other side. Note that inside the ISDN network also a protocol translation is done between the Netherlands and Germany.

The inverse multiplexers were manufactured by Controlware. They provide 3 S_0 interfaces with 1TR6 signalling on the ISDN side and an X.21 interface on the other side to provide a transparent data channel with a bandwidth between 64 kbit/s and 384 kbit/s.

3.4 Performance objectives

To draw conclusions from the tests, it is necessary to know the performance objectives of the ISDN network.

Results of bit error rate tests are presented in terms of bit error rate, errored seconds, severely errored seconds and degraded minutes. The last three terms are found within CCITT G.821 recommendation [12] that states:

Performance classification	**Objective**
Degraded minutes	Fewer that 10% of one-minute intervals to have a bit error ratio worse than 1e-6
Severely errored seconds	Fewer than 0.2% of one-second intervals to have a bit error ratio worse than 1e-3
Errored seconds	Fewer than 8% of one-second intervals to have any errors

Table 3.1 *Error performance objectives for international ISDN connections.*

The performance objectives are stated for each direction of a 64kbit/s circuit-switched connection.

The objectives relate to a hypothetical connection of 27500 km. However it is recognized that many real connections will offer a better performance than the limiting values given above. The allocation of the overall performance objectives over the hypothetical connection of 27500 km are also specified. Apportionment is based on the assumed use of transmission systems having qualities falling into one of a limited number of different classifications. Three distinct quality qualifications have been identified representative of practical digital transmission circuits and are independent of the transmission systems used. These classifications are termed local grade, medium grade and high grade and their usage generally tends to be dependent on their location within a network. Figure 3.2 outlines their meaning.

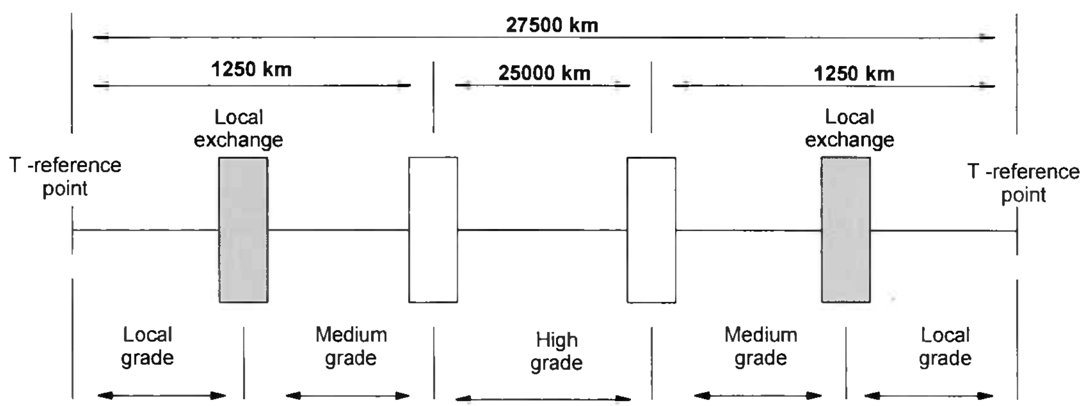

Figure 3.2 *Circuit quality demarcation of longest hypothetical reference configuration.*

The apportionment of the permitted degradation, i.e. 10% degraded minutes and 8 % errored seconds, is given in table 3.2.

Circuit classification	**Allocation of the degraded minutes and errored seconds objectives given in table 3.1**
Local grade (2 ends)	15% block allowance to each end (Note 1)
Medium grade (2 ends)	15% block allowance to each end (Note 2)
High grade	40% (equivalent to conceptual quality of 0.0016% per km for 25000 km)

Table 3.2 *Allocation of the degraded minutes and errored seconds objectives for the three circuit classifications.*

Note 1 - The local grade apportionment is considered to be a block allowance, i.e. an allowance to that part of the connection regardless of length.

Note 2 - The medium grade apportionment is considered to be a block allowance, i.e. an allowance to that part of the connection regardless of length. The actual length covered by the medium grade part of the connection will vary considerably from one country to another. Transmission systems in this classification exhibit a variation in quality falling between the other classifications.

The total allocation of 0.2% severely errored seconds is subdivided into each circuit classification (i.e. local, medium, high grades) in the following manner:

a) 0.1% is divided between the three circuit classifications in tbe same proportions as adopted for the other two objectives. This results in the allocation as shown in table

Circuit classification	Allocation of severely errored seconds objectives
Local grade	0.015% block allowance to each end
Medium grade	0.015% block allowance to each end
High grade	0.04%

Table 3.3 *Allocation of severely errored seconds*

b) The remaining 0.1% is a block allowance to the medium and high grade classifications to accommodate the occurrence of adverse network conditions occasionally experienced (intended to mean the worst month of the year) on transmission systems. This 0.1% is not taken into account onthe tested connections.

With the values of tables 3.1 and 3.2, the allocation of the percentage of degraded minutes intervals and errored seconds objectives can be calculated. Table 3.4 shows the results.

Circuit classification	Network performance objectives at 64 kbit/s	
	% degraded minutes	% errored seconds
Local grade	1.5	1.2
Medium grade	1.5	1.2
High grade	4.0	3.2

Table 3.4 *Allocation of % degraded minutes intervals and errored seconds objectives*

On the performed tests, inverse multiplexing equipment was used. Therefore bit error rates were measured often over a 384 kbit/s connection instead of a 64 kbit/s connection. Since many practical bit error rate tests will take place on such connections, CCITT has also specified guidelines for the assessment of the performance of higher bit rate systems. They have defined some formulas to provide a normalized estimate to the 64 kbit/s parameters of the error performance. However to calculate the percentage errored seconds normalized to 64 kbit/s the distribution of the errors has to be known. The tests done only provide overall performance numbers and thus the formula cannot be used. For the percentage of degraded minutes and severely errored seconds CCITT states that the percentages can be taken directly from measurements at primary bit rates (i.e. X% degraded minutes at the primary bit rate yields X% degraded minutes at 64 kbit/s).

The tested connections were mostly connections from the Netherlands to Germany. Such a connection consists of two times local grade and one time medium grade. Given the numbers of table 3.4, not more than 3.6% errored seconds are expected.

Controlled slips, which exist with PDH (Plesiochronous Digital Hierarchy) networks and may be perceived as short bursts of errors, are not included in the calculations of error performance objectives in recommendation G.821. The performance specification on controlled slips is found within CCITT G.822 [13] which states:

Performance category	**Mean slip rate**	**Proportion of time**
Configuration of the international digital network	$\leq$ 5 slips in 24 hours	$\succ$ 98.9%
National timing control arrangements	$\succ$ 5 slips in 24 hours and $\leq$ 30 slips in 1 hour	$\prec$ 1.0%
Wander due to extreme temperature variations	$\succ$ 30 slips in 1 hour	$\prec$ 0.1%

Table 3.5 *Controlled slip performance on a 64 kbit/s international connection or bearer channel.*

Again these objectives refer to a hypothetical connection of 27500 km.

A period of unavailable time begins when the bit error ratio (BER) in each second is worse than 10^{-3} for ten consecutive seconds. These ten seconds are considered unavailable time. A new period of available time begins with the first second of a period of ten consecutive seconds each of which has a BER better than 10^{-3}.

3.5 Test results

For a detailed overview of all test results see Appendix A.

3.6 Discussion of test results

The availability of international 64kbit/s ISDN circuits proved to be satisfactory. Only very few of the call attempts failed. The reason for non-succeeding calls was not specified, except for one specific call, where the message 'network congested' appeared. At any moment that a call failed, a second try within 1 minute always succeeded.

The quality of 64kbit/s ISDN circuits is good. The performed bit error rate tests show that the network is errorfree for more than 99.9% of the time. Errors occur often in bursts. The error performance is well within the limits given by ITU-TSS, even if the large bursts of errors that sometimes occurred are considered to be 'normal' errors instead of controlled slips that do not count to the error performance objectives.

Some established calls were dropped during bit error rate tests. The reason for these call drops is unknown. Since monitoring by the German PTT did not show any line or network errors, it is assumed that the call drops were caused by the Controlware TAXI inverse multiplexers. One reason for the call-drop may be that the Controlware TAXI inverse multiplexers cannot handle short bursts of errors, which appeared on the international ISDN connections that were used for the tests. A test with Controlware CITAM terminal adapters between ESTEC and DLR showed no call drops, which strengthens the idea that the Controlware TAXI inverse multiplexers cannot always cope with bit error bursts and consequently sometimes fail.

The Controlware TAXI inverse multiplexers still show some shortcomings in terms of reliability. After intensive use, they sometimes need a hardware reset to function properly again. Investigations on the reason for call drops have to be done.

The bit error rate tests over an internal PABX link showed that some the bit errors (appr. 1 in 10^9) are caused by the combination of the terminal equipment and the PABX. These bit errors were unexpected, due to the short link no bit errors should occur so it is anticipated that the physical layer interfaces of the PABX or TE are not fully up to the standard. Further investigations are therefore necessary to pinpoint the source of the fault.

On the tests between France and the Netherlands some problems occurred in call setup, but once a reset of the PBX at France was done this resulted in 100% success rate.

In general it therefore can be concluded that most of the call setup problems are caused by the equipment. Some shortcomings in the equipment stability (inverse multiplexers, PABX) have been identified, as well as some shortcomings in the D channel implementation of inverse multiplexers and other terminal equipment.

Some equipment is not very stable (PBX, inverse multiplexers) and it seems that some of the ISDN protocols (VN3) are not implemented properly.

Chapter 4 Cost comparison between ISDN and leased lines

4.1 Public transmission media for datacommunications

To connect geographic distributed LAN's with PC's, minicomputers and mainframes there is an enormous amount of equipment available. The most known examples are remote-routers and remote-bridges that can directly be connected to digital leased lines, but there are also many LAN/LAN connections possible that use modems and specific software like cc:mail and PC-Anywhere. LAN/LAN links that are used nowadays, use one of the following interlinks:

- PSTN
- analog leased line
- digital leased line

4.1.1 PSTN

With high-speed modems, the obtainable maximum dataspeed across the PSTN is about 9600 bit/s (without compression). Compared with the 64 kbit/s of a single B-channel, and taken into consideration that the connection costs are about the same for ISDN and the PSTN in most European countries, the PSTN is no real alternative for ISDN concerning datacommunications. However, the PSTN will still be the major network for datacommunications for private users during this decade. This is due both to the high installation costs of ISDN in some countries and the current lack of availability of ISDN services like Bulletin Boards.

4.1.2 Analog leased line

The costs of this line are based on the delivered quality and the distance. With the highest quality line, and high-speed modems, the obtainable speed is 19200 bits/s. PTTs discourage the use of analog leased lines because of the high operation costs. The Dutch PTT for example will increase the price of analog leased lines each year, while decreasing the prices of digital leased lines. In 1993 the prices of analog leased lines increased by 11%, whereas the prices of the digital lines decreased by 20%.

4.1.3 Digital leased line

The digital leased line has advantages over analog lines. The throughput is higher. However, digital lines are still expensive. Especially on very short connections (<10 km), the digital lines are much more expensive than the analog lines. However, this will change during the next few years since the prices of analog leased lines will increase and the prices of digital leased lines decrease.

4.1.4 Public packet switched networks

Datanet 1 is the public X.25 data network from PTT Telecom. It was started around 1980 and there were high expectations wrt this network: The principle of a public, packet-switched network has much potential advantages above digital point-to-point links, like bandwidth on demand, the possibility to reconfigure the network very quickly and the possibility to communicate to other countries. However, most of the companies kept their leased lines. The reason for this is that the X.25 equipment was not reliable in the first few years of the Datanet-existence [1]. Moreover, the costs of this network are much higher in comparison with the costs of leased lines in most communications scenarios. Datacommunications within the X.25 network are charged on a cost-per-byte base, whereas a leased line is charged with a fixed monthly amount. The Datanet is only considerable on terminal-traffic and comparable traffic since on that type of traffic there is little traffic per second.

4.1.5 ISDN

Considering the given options for data transmission, the ISDN is an interesting alternative. Since only the actual connection time of the ISDN is charged (like in the PSTN), it should be a good alternative for leased lines on connections that are not used constantly.

4.2 Cost comparison between ISDN and digital leased lines

A cost comparison has been made between the ISDN and digital leased lines. Since the analog lines are becoming more expensive each year, this is not a good alternative. Besides, companies like ESA only use digital lines because of their higher reliability. A comparison with the public X.25 data network is also not relevant. To send 100 kilobytes of information across ISDN within the Netherlands costs less than 20 cents, while it costs several guilders to send it across the X.25 network.

The tariffs on national leased lines only depend on the distance to cover. Therefore, for each specific distance a break-even point can be calculated between a leased line and ISDN. For international leased lines, there is only one tariff within each country. For example a leased line from Marseille to Hamburg has the same price as a leased line from Paris to Munchen. This means that given a link between 2 countries, a break-even point between ISDN and leased lines can easily be calculated. Table 4.1 shows the calculations for connections from the Netherlands to various European countries. The break-even point is the number of hours that ISDN can be used per day, based on a 20 days month, for the price of a 64 kbit/s leased line. In table 4.1, the initial installation costs of leased lines and ISDN are not considered. Table 4.2 shows the installation costs for ISDN and a 64 kbit/s leased line. Within companies, costs like these are usually written off in 60 months. Therefore, table 4.3 shows the break-even points with the initial installation costs considered.

It is also interesting to know how the break-even point changes if larger bandwidth is needed. Therefore, table 4.4 shows the break-even points for several connections from the Netherlands to Germany.

The tariffs for leased lines within a country are based on the distance to cover. Therefore, with each distance there is a different break-even point. Table 4.5 shows break-even points for several distances. Note that the 'local tariff is only used for the 5-km distance; for other distances, the so called 'interlocal' tariff is used.

Tariffs of international leased lines								
From:	To:		Dutch part:	Foreign part:				Total
Netherlands		Speed	Monthly:	Valuta:	Monthly:	Exch. rate	Month DFl	Month Dfl
	Germany	64 kbit/s	3600	DM	4220	1.122	4735	8335
	Belgium	64 kbit/s	3600	BEF	73500	0.055	4013	7613
	England	64 kbit/s	3800	Pound	1208	2.835	3425	7225
	Ireland	64 kbit/s	4500	Iri Pound	1583	2.735	4330	8830
	France	64 kbit/s	3800	FF	17400	0.328	5707	9507
	Spain	64 kbit/s	4500	Pta	362500	0.014	5024	9524
	Portugal	64 kbit/s	4500	Esc	811000	0.011	8799	13299
	Italy	64 kbit/s	4500	Lir	6680000	0.001	7856	12356
	Denmark	64 kbit/s	4500	Gfcs	9333	0.287	2678	7178
	Sweden	64 kbit/s	4500	Sek	15100	0.244	3689	8189
	Norway	64 kbit/s	4500	Nok	19500	0.259	5058	9558
	Austria	64 kbit/s	4500	Shilling	50000	0.160	7975	12475
	Switzerland	64 kbit/s	4500	Sfr.	2690	1.332	3583	8083
ISDN tariffs			ISDN call		**Break-even per month**		excl. installation costs	
ISDN per month			from NL,		number of hours ISDN use per day			
Neth.	Foreign	Foreign	Dfl./Hr.		1 month = 20 days			
	local curr	Dfl.						
85	64	72	57		7.17		Germany	
85	1050	57	57		6.55		Belgium	
85	28	79	66		5.35		England	
85	35	96	87		4.97		Ireland	
85	300	98	57		8.18		France	
85	8000	111	87		5.36		Spain	
85	4310	47	87		7.57		Portugal	
85	50000	59	87		7.02		Italy	
85	160	46	66		5.34		Denmark	
85	490	120	54		7.39		Sweden	
85	339	88	87		5.39		Norway	
85	400	64	87		7.08		Austria	
85	50	67	54		7.34		Switzerland	

Table 4.1 *Break-even point calculation ISDN versus 64 kbit/s leased line.*

From table 4.1, it can be concluded that the break-even point for international 64 kbit/s ISDN connections is between 5 and 8 hours of ISDN use per working day.

All the call tariffs used are the tariffs from the Dutch PTT as per 1 April, 1994. The tariffs valid for office hours (8:00-20:00) are used.

On the Dutch ISDN, there is a so-called answering tariff. This is an initial cost for each answered call that has been made. This tariff is Fl. 0.0085 for national calls and Fl. 0.15 for international calls. This tariff has not been taken into consideration with the cost break-even analysis, since it will only be of significant interest if many short calls are done.

	Leased line 64 kbit/s install costs					Install costs of leased line balanced out in 60 months
From NL	Dutch part	Foreign part			Total install.	
To:	Installation:	Valuta:	Installation:	Exch. rate	Dfl.	
Germany	2450	DM	2000	1.122	4694	78.2
Belgium	2450	BEF	100614	0.055	7944	132.4
England	2450	Pound	0	2.835	2450	40.8
Ireland	2450	Iri Pound	1600	2.735	6826	113.8
France	2450	FF	12000	0.328	6386	106.4
Spain	2450	Pta	500000	0.014	9380	156.3
Portugal	2450	Esc	200000	0.011	4620	77.0
Italy	2450	Lir	300000	0.001	2803	46.7
Denmark	2450	Gfcs	29500	0.287	10914	181.9
Sweden	2450	Sek	18500	0.244	6970	116.2
Norway	2450	Nok	6500	0.259	4136	68.9
Austria	2450	Shilling	3000	0.160	2929	48.8
Switzerland	2450	Sfr.	800	1.332	3516	58.6

ISDN Basic Access install costs				Installation costs of ISDN balanced out over 60 months		
Install Netherlands	Install foreign local curr	Install foreign Dfl.	Total install Dfl.	Dfl.		
600	130	0.0	600	10.0		Germany
600	3500	0.0	600	10.0		Belgium
600	400	0.0	600	10.0		England
600	420	0.0	600	10.0		Ireland
600	675	0.0	600	10.0		France
600	45000	0.0	600	10.0		Spain
600	30172	0.0	600	10.0		Portugal
600	400000	0.0	600	10.0		Italy
600	1520	0.0	600	10.0		Denmark
600	6200	0.0	600	10.0		Sweden
600	3689	0.0	600	10.0		Norway
600	1200	0.0	600	10.0		Austria
600	200	0.0	600	10.0		Switzerland

Table 4.2 *Installation costs of 64 kbit/s leased lines and ISDN.*

Tariffs international leased lines					Tariffs ISDN			
From:	To:	Speed:	Installation costs total Dfl	Per month total	ISDN per month Holland	Germany	ISDN call month	Break-even per month month = 20 days
Netherlands	Germany	64 kbit/s	4690	8326	85	71.81	57	7.15
		128 kbit/s	7140	17265	85	71.81	114	7.50
		192 kbit/s	8365	26499	170	143.62	171	7.66
		256 kbit/s	9590	32265	170	143.62	228	7.01
		384 kbit/s	12040	40216	255	215.43	342	5.81
		512 kbit/s	14490	50510	340	287.24	456	5.47
		768 kbit/s	16940	62757	510	430.86	684	4.52
		1024 kbit/s	21840	76154	850	558.76	912	4.10
		1536 kbit/s	25740	87321	850	558.76	1368	3.14
		2 Mbit/s	31490	100969	850	558.76	1824	2.78

Tariffs of international leased lines									
From: Netherlands	To:	Speed	Month Dfl	Install month 60 month write-off	Total month	ISDN cost month Dfl.	ISDN inst. 60 month write-off	ISDN, call from NL Dfl./hr.	Break even No. of hrs. a day
	Germany	64 kbit/s	8335	78.2	8413	157	12.4	57	7.23
	Belgium	64 kbit/s	7613	132.4	7745	142	14.2	57	6.66
	England	64 kbit/s	7225	40.8	7266	164	28.9	66	5.36
	Ireland	64 kbit/s	8830	113.8	8943	161	29.1	87	5.02
	France	64 kbit/s	9507	106.4	9614	183	13.7	57	8.26
	Spain	64 kbit/s	9524	156.3	9681	196	20.4	87	5.44
	Portugal	64 kbit/s	13299	77.0	13376	132	15.5	87	7.60
	Italy	64 kbit/s	12356	46.7	12402	144	17.8	87	7.03
	Denmark	64 kbit/s	7178	181.9	7360	131	17.3	66	5.46
	Sweden	64 kbit/s	81[illegible]	116.2	8305	205	33.2	54	7.47
	Norway	64 kbit/s	95[illegible]	68.9	9627	173	24.9	87	5.42
	Austria	64 kbit/s	12475	48.8	12524	149	13.2	87	7.10
	Switzerland	64 kbit/s	8083	58.6	8142	152	14.4	54	7.38

Table 4.4 *Break even point ISDN versus leased line on Netherlands-Germany connection for several transmssion speeds.*

Table 4.3 *Break even point calculation ISDN versus 64 kbit/s leased line including installation costs.*

It can be concluded that for bandwidths larger than 512 kbit/s, the break even point decreases in favour of leased lines. Note that for speeds from 1024 kbit/s and higher, an ISDN primary rate connection is cheaper than several basic rate interfaces.

Tariffs Dutch national		**leased lines**		**Tariff ISDN**	Call/hr		Break-even	hrs /day
Distance	Speed	Month cost		Month cost			Month=20 days	
5 km	64 kbit/s	956.25		85	2.25		19.4	
15 km	64 kbit/s	1440.5		85	11.49		5.9	
40 km	64 kbit/s	2181		85	11.49		9.1	
75 km	64 kbit/s	2605		85	11.49		11.0	
110 km	64 kbit/s	2870		85	11.49		12.1	
150 km	64 kbit/s	3150		85	11.49		13.3	
Tariffs Dutch national		**leased lines**		**Tariff ISDN**	Call/hr		Break-even	hrs/day
Distance	Speed	Month cost		Month cost			Month=20 days	
5 km	2 mbit/s	5019		850	67.5		3.1	
15 km	2 mbit/s	7570		850	344.68		1.0	
40 km	2 mbit/s	11446		850	344.68		1.5	
75 km	2 mbit/s	13678		850	344.68		1.9	
110 km	2 mbit/s	15068		850	344.68		2.1	
150 km	2 mbit/s	16538		850	344.68		2.3	

Table 4.5 *Break-even point calculation Dutch national leased line versus ISDN*

It can be concluded, that for 64 kbit/s, ISDN will be the most economic solution in most cases, but at 2 Mbit/s, ISDN is too expensive if the full bandwidth is needed for more than two hours a day. On connections that just cross the border of the local tariff zone, ISDN is relatively more expensive than on connections within the zone or well outside the zone (see the break-even point for the 15 km connection).

It is obvious that ISDN gives much more flexibility. This is not taken into account. Therefore, for each company the consideration to choose for either ISDN or a leased line shall include those considerations. Especially on 64 kbit/s connections, ISDN seems to be the most economic solution.

Chapter 5. Applications on the ISDN

5.1 Introduction

The integrated services digital network offers the possibility to introduce a range of new telematic services to the end-user. Especially on datacommunications, ISDN offers a new range of solutions because of the high speed of the circuits. Datacommunications on the PSTN is only possible with analog modems that reduce the maximum data throughput to ca 9600 bits/s, whereas ISDN offers digital circuits with a capacity of 64 kbits/s.

The ISDN testbed at ESTEC was and will be used for the evaluation of several of the possible telematic services.

5.2 Description of applications

5.2.1 Telephone applications

The ISDN offers many convenient services in the telephone area. Examples are Calling Line Identification and Closed User Group. However, many PABXs already offer these services on the PSTN on internal telephone traffic within companies. The extra convenience from the ISDN is that these services are delivered to each user, both on internal and on external lines. The calling line identification can also be used for security purposes on for example LAN-LAN connections.

The ISDN B channel capacity of 64 kbit/s is based on the 8-bits ADPCM (for a bandwidth of 4000 Hz, the Nyquist theorem states that 8000 samples per second are needed). However, with modern coding techniques it is possible to achieve compression ratios of up to 10:1 and more without recognized decrease of speech quality. Therefore to distribute one telephone call over a 64 kbit/s channel, is a waste of bandwidth. Several manufacturers therefore already produce compression equipment with intelligent switching capabilities, to distribute several telephone calls simultaneously over a 64 kbit/s channel. For companies like ESA, with a number of large sites in Europe and much telephone traffic between these sites, the implementation of such compression equipment can cause a large reduce in telephone costs.

5.2.2 Telefax applications

A new fax-standard has been specified for use with 64 kbit/s switched circuits. This group IV fax has a maximum resolution of 400 DPI. In comparison with the existing fax group III, both quality and speed improve much. Currently, only a limited number of group IV faxes are in use. This is mainly because of the price of a fax group IV device, which lies in the Fl. 8000,- range. Although a fax group IV can communicate with a fax group III, the advantages disappear when communicating with a fax group III: the quality and speed will be determined by the fax group III. Although the fax group IV can decrease transmission costs, it is not expected that the use will increase rapidly unless the prices are lowered distinctly.

5.2.3 File transfer applications

On the PSTN datacommunication is possible with the use of analog modems. However, the speed is limited to some 9600 bits/s. On the ISDN, higher speeds are possible. Also the connection time is shorter, at least on national connections. For PC-to-PC communications, currently a wide range of software and hardware is available, most from German and French manufacturers. A typical PC installation consists of an ISDN PC card with one or more software packages. There are two types of PC cards: passive ones and active ones. The active cards have a processor onboard to off-load the PC's CPU. Performance is therefore independent of the used PC. The passive cards do not have a processor and therefore fall back to the CPU of the PC. With the active PC cards, there is the option for an effective (fast) datacompression following the V42bis standard to come to a high net throughput over the ISDN. With textfiles, a net throughput of up to 200 kbit/s can be realised on a single B-channel. Another interesting option of ISDN cards is the ability to bundle the two B-channels to one 128 kbit/s datachannel.

The most important French and German ISDN-equipment manufacturers have specified the Common ISDN application program interface (CAPI). Software compatible with this standard will run on any ISDN PC card that is driven by a CAPI-driver. All the ISDN PC-cards sold in the Netherlands and Germany come with the CAPI.

ITU-TSS has specified a uniform datatransfer standard, called Eurospeed. Different programs that use this standard to exchange data, have to be compatible. Eurospeed has options for channel bundling and data compression.

5.2.4 LAN-LAN links

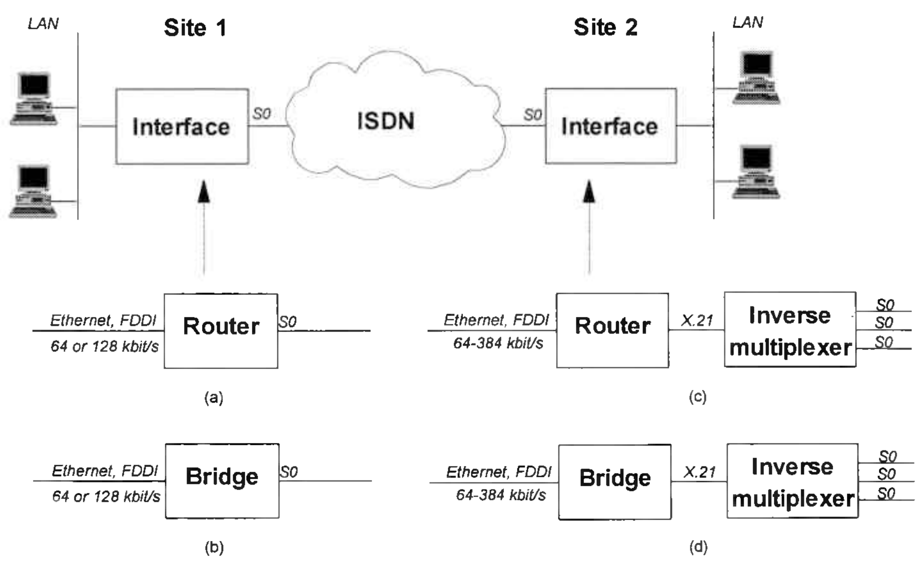

Figure 5.1 *Different options for LAN-LAN links over the ISDN. (a) An ISDN router, (b) ISDN bridge, (c) non-ISDN router, (d) non-ISDN bridge.*

From chapters three and four, it was concluded that the ISDN can be an economic alternative for leased lines to interconnect LANs over long distances. Since most of these links are idle for a considerable amount of time, the leased lines are relatively expensive. There are several ways to implement a LAN-to-LAN interconnection. Figure 5.1 shows some solutions.

There are a several routers and bridges available with a native ISDN interface. They can usually interconnect two LANs that use the Ethernet, Token Ring or FDDI topology. For routers and bridges without ISDN interface, an alternative is to use an ISDN terminal adapter or an inverse multiplexer to provide leased line emulation with X.21/V.35/RS-449/G.7xx interface.

With all the routing and bridging equipment there is one major problem that needs special attention:

Some routing protocols regularly exchange updates over the network to make sure that every other router has a current routing table (the only way to ensure that all nodes can communicate). RIP, for instance, sends an update every 30 seconds, which means that an ISDN router at a remote site would be setting up a call twice a minute, hardly cost-effective. To prevent this from happening, ISDN routers offer extensive spoofing and filtering. The former fools the stations on the remote LAN into thinking that they are permanently connected to the backbone (thus keeping sessions alive). At the same time, the ISDN router filters all packets coming from the LAN and headed for the wide area, storing them in cache memory and examining them to see if they are carrying data or network information (updates). If data is detected, a call is set up and it is passed along. If network information is being carried, the router checks to see if there has been any change since the last update was sent. (RIP updates go off even if nothing has changed since the previous update). When changes are detected, this information is held in a buffer until the router dials up the ISDN link. At this point, the network information is shipped with the data.

5.2.5 Dial-in services

Services like bulletin boards and CompuServe will become available on the ISDN soon. The ISDN offers higher data rates than the PSTN, and therefore is more suitable for interactive and multimedia applications. Interactive services like the World-Wide-Web can be used with reasonable speed over the ISDN. Therefore it is stated that ISDN is the first step towards the 'multimedia superhighway'.

At this moment, analog modems are being used to perform remote login on the ESA networks. This method has proven to be both slow and unreliable. The ISDN is a fast, reliable and cost-effective medium to use for dial-up services. To evaluate these services, a router configuration[1] has been set up at the ESTEC ISDN testbed. The router has access to the Ethernet on one side and a maximum of 6 B-channels on the other side. These B-channels can be accessed via the ISDN by remote PC's with build-in ISDN cards and appropriate software. Reliability and performance were measured and evaluated.

[1] Schneider & Koch Multi Protocol Router

5.2.6 Multimedia applications

There is a growing need for interactive applications on PC's and workstations that can be used simultaneously by multiple users. Two of these applications are (desktop) videoconferencing and collaborative working. Using compression techniques it is possible to videoconference over a 128 kbit/s or even a 64 kbit/s link. ITU-TSS has specified compression methods in their standards H.261 and H.320. The compression algorithms are based on the principle not to send all the picture information continuously, but to only send the differences that have occurred in comparison with the last picture. With an application like videoconferencing, the picture will remain fairly constant since the participants usually sit and do not move too much. With 128 kbit/s, a fair quality can already be achieved unless the participants in the videoconference move quickly.

Some manufacturers have developed interactive workstations with built-in videoconferencing and collaborative working hardware and software. It is possible to call other users over the ISDN, then open a window with the video capture of the other user and the option to work at the same moment on the same application. Since the video service on these systems usually is implemented conforming to the H.320 and H.261 standard, also large videoconference systems (e.g. GPT (GE-Plessey Technology) /BT (British Telecom)) can be reached from the desktop.

5.3 Terminal adaptation and inverse multiplexing

To run specific applications over the ISDN, the available equipment does not always support ISDN or it needs more than 64 kbit/s bandwidth. For these problems, special equipment is available.

5.3.1 The terminal adapter

Apparently ISDN equipment will be largely incompatible with today's telecommunications devices. Since the ISDN is expected to evolve from the existing telecommunications networks, during the transitional period solutions to the problem of interworking between ISDN and non-ISDN equipment need to be developed.

From a technical point of view there are two major compatibility issues that must be addressed. The first relates to the interworking between the interexchange facilities of an ISDN and those of an existing non-ISDN such as the public switched telephone network. Such interworking will likely be required for an interim period, either to provide connections between pairs of TEIs attached to disjointed ISDNs or between a TEI attached to an ISDN and a conventional terminal attached to an existing network. The second issue concerns the adaption of existing non-ISDN telecommunications terminals to the user-to-network interface of an ISDN, to allow their continued use during a period of transition to the ISDN.

In recent years considerable effort has gone into the development of specifications for the conversion of the most important types of non-ISDN data terminal equipment (DTE) in use today-namely, those whose communications characteristics have previously been standardized by the ITU-TSS. The status of this work is summarized in six ITU-TSS recommendations.

One of the most common terminal adapters is the X.21 terminal adapter. Its function is to adapt a DTE conforming to the ITU-TSS X.21 specifications to the requirements of the S or T reference points. On the physical layer, apart from electrical differences there is also a rate adaptation problem since the X.21 signal may have a bit rate below 64 kbit/s. The rate adaptation problem is solved by adding bits to the X.21 signal on one side and distract them on the other side of the ISDN link. Obviously the biggest problem is the signalling adaptation. A detailed description of this problem is however outside the scope of this report.

Trials with existing terminal adapters at ESTEC have resulted in some specific problems that may occur while using X.21 terminal adapters. They are listed below.

- Terminal adapters from the manufacturer Controlware have been used in the Netherlands (ESTEC), Germany (ESOC),and France (ESA HQ and ESA/CNES Toulouse). ESTEC and ESOC were using the German 1TR6 signalling protocol, ESA HQ and ESA Toulouse the French VN3 protocol. Between ESTEC and ESOC there were no problems, but between ESTEC and France call setup was only possible in one direction. The problem occurs because of the service indicator that the terminal adapter sends. It can be set either to X.21 services or 64 kbit/s unrestricted. However, in both modes errors occurred. This is due to an implementation problem of the D-channel protocol.

5.3.2 Inverse multiplexing equipment

Since the ISDN is a circuit switched network, a connection between two places can be made over several different routes. Therefore, if a datacommunications connections over 64 kbit/s is requested, and thus several B-channels are used, the data arrives desynchronized at the destination. To solve this problem, many different inverse multiplexers exist. In principle, they consist of one or more ISDN BRI or PRI on one side, usually an X.21 interface on the other side, and in between some delay lines and protocol conversion and synchronization handlers. Until recently, many proprietary and incompatible inverse multiplexers were available. As soon as a company had decided to choose for brand A, all the inverse multiplexing equipment had to be brand A. Recently, the Bandwidth On Demand Interoperability Group (BONDING) has specified a standard for multiplexing equipment. In September, 1993, this standard has been proposed to the Telecommunications Industry Association. The title of the standard is "Aggregation of Multiple Independent 56 kbit/s or 64 kbit/s Channels into a Synchronized Wideband Connection" [15].

5.3.2.1 The BONDING standard

The purpose of the standard is to define a frame structure and procedures for establishing a wideband communications connection by combining multiple switched 56/64 kbit/s channels through the use of a Channel Aggregation Unit. The draft standard provides:

- A method for automatically synchronizing and aligning multiple bearer channels to support application data rates including N x 8 kbit/s, N x 56 kbit/s and N x 64 kbit/s, to be transmitted over 56 kbit/s and 64 kbit/s bearer channels.

- A method for providing end-to-end negotiation of such items as operating mode, bearer channel rate, phone numbers, etc.

- A method for monitoring data transfer throughout the call to determine if and when a channel fails. A failed channel is one that loses synchronization due to line errors or phase changes or is disconnected while a session is in progress.

- A call failure recovery procedure that supports various options. These include call disconnect, rate reduction, bandwidth replacement or rate reduction in conjunction with bandwidth replacement.

- A method for supporting a dynamically variable transmission rate-during a data call.

- A method for providing a remote loopback function.

At the transmitting end, user data is placed in a framing structure in each bearer channel and transmitted over multiple, independent channels. At the receiving end, all channels are phase aligned and synchronized (using the framing structure) to recreate the original data stream. This framing and synchronization is transparent to the attached application.

Due to the network routing the channels used for the wideband connection are independent of each other, and thus the data in each channel might be individually delayed relative to the data in other channels. The specified frame structure for each 56/64 kbit/s bearer channel provides for the alignment of data octets from the individual channels to their original sequence before reforming the individual channels into a composite serial data stream at the terminating end. Overall transit delay for the end-to-end connection is equal to the longest transit delay from all the channels plus a constant delay due to the frame alignment that depends on implementation.

Once aligned, data transfer can be constantly monitored throughout the call. The failure of a channel, for reasons such as call disconnection, phase slip or high error rate, can be automatically detected. Various fault isolation and recovery procedures are defined in response to these scenarios.

Up to 63 individual switched 56/64 calls can be combined to form a single transparent N x 56/64 channel.

Modes of operation

Five modes of operation are supported. The following two modes are required:

Mode 1:

This mode supports user data rates that are multiples of the bearer rate. It provides the user data rate with the full available bandwidth, but does not provide an in-band monitoring function. The overhead octets are removed after the call is phase aligned. Error conditions on one or more channels that disturb overall system synchronization are not recognized automatically after the call is in active state.. Recovery from these error conditions during the active state requires manual or external intervention.

Transparent Mode:

In this mode, incoming (outgoing) channels are "cut-through" to the applications served by the Channel Aggregation Unit. The Channel Aggregation Unit performs no delay equalization or parameter negotiation. This mode is useful when delay equalization is performed by some other means (or is not required) and one endpoint is not a Channel Aggregation Unit.

The following three modes of operation are optional.

Mode 0:

This mode provides initial parameter negotiation and Directory Number exchange over the master channel, then reverts to data transmission without delay equalization. It does not provide an inband monitoring function. This mode is useful when the calling endpoint requires Directory Numbers, but the delay equalization is performed by some other means (e.g., attached video codec).

Mode 2:

This mode supports user data rates that are multiples of 63/64 of the bearer rate. An in-band monitor function provides a continuous check for delay equalization and an end-to-end bit error rate test. (Error rate testing is accomplished by performing a cyclic redundancy check calculation on an octet sequence before transmission and testing the same sequence for errors on the receive end.) The user data rate is the bandwidth remaining after the insertion of overhead octets (i.e., 98.4375% or 63/64 of the total network bandwidth).

Mode 3:

This mode supports user data rates that are integral multiples of 8 kbit/s, including N x 56 and N x 64 kbit/s. All channels use the same bearer channel rate. An in-band monitor function provides a continuous check for delay equalization and an end-to-end bit error rate test. The overhead octets required for monitoring are provided by adding bandwidth (most likely an additional bearer channel), thereby preserving the full user data rate. The overhead octets are included in each bearer channel.

5.3.2.2 The inverse multiplexers used at ESA

The currently available inverse multiplexers provide either no BONDING or BONDING Mode 1. Most of the tests that have been done were done with Controlware TAXI inverse multiplexers. The Controlware TAXI equipment does not (yet) provide BONDING. To set up a call over 6 B channels, all ISDN numbers had to be specified. If because of an error one of the channels failed, there was no error recovery and only the connected channels were aligned. Currently, the Controlware TAXI inverse multiplexers are not in operational use at ESA, but are only used to support temporary experiments.

A few tests were performed with inverse multiplexers from Ascend. These multiplexers provide BONDING mode 1 compliance. They showed more flexible operation in comparison with the Controlware TAXI equipment. If a connection of 384 kbit/s was requested, the calling Ascend called the receiving Ascend over 1 B-channel. The receiving Ascend then specified which other ISDN numbers to dial for the other five B channels. If one of the channels fails, the Ascend keeps trying to call the failed channel and to establish the preferred connection. In the mean time, it will align the successful channels and provide a synchronous channel. The Ascend inverse multiplexers have been implemented at various ESA sites to serve ESA's operational videoconference service.

Unfortunately, no other BONDING supporting equipment has been evaluated. Therefore it is uncertain if such equipment would be 100% compatible with the Ascend equipment.

5.4 Results of experiments

5.4.1 Telephone services

Since no ISDN telephones were available, no experiments have been done. From the experiments with Controlware TAXI inverse multiplexers, it can be concluded that at least the Calling line identification works between the Netherlands and Germany.

5.4.2 Telefax services

No experiments with fax group 4 have been done. However, during the EURIE '93 showcase in December 1993, demonstrations were given with fax group 4. The quality of the fax is good, although not as good as the quality obtained with a high quality laser printer (this document).

5.4.3 File transfer

File transfer between PC's was performed with different combinations of software and hardware. Two PCs were equipped with ITK ISDN cards, one PC with a Diehl ISDN card and one PC with an AVM PC card. All cards contained a separate processor. The software packages were all using CAPI as an interface platform between software and ISDN card.
Figure 5.2 shows the configuration.

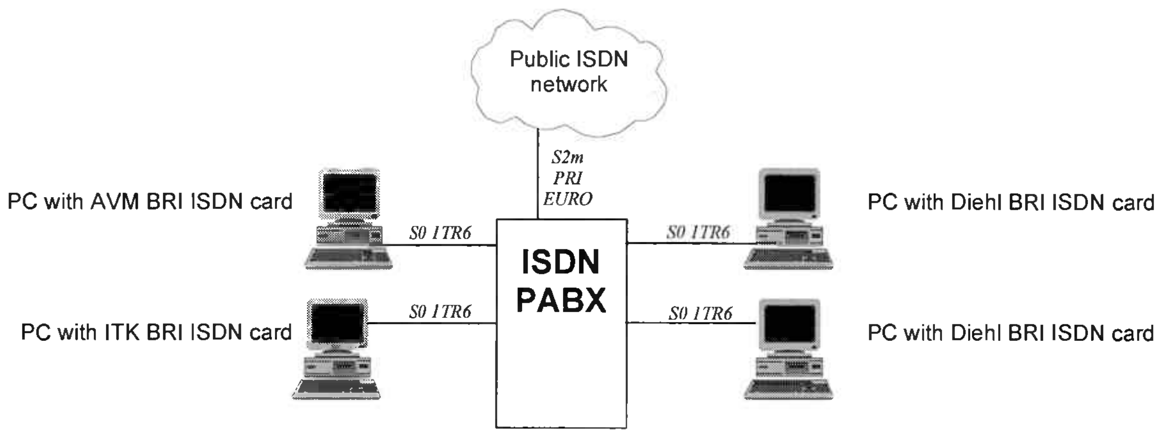

Figure 5.2 *PC-to-PC file transfer.*

The following combinations were tested:

PC 1	PC 2	Software	Average throughput [kbit/s]
Diehl	Diehl	ITK	62
Diehl	AVM	ITK	62
Diehl	ITK	ITK	62
ITK	ITK	ITK+compr.	85
AVM	ITK	ITK	62
Diehl	Diehl	AVM	63

The ITK software had an option for datacompression. However, this option only works if both PCs have an ITK PC card.

5.4.4 LAN-LAN links

A LAN-LAN link was set up locally at ESTEC. Two Ethernet networks consisting of a single PC were connected via CISCO routers. Figure 5.3 shows the configuration. The CISCOs were configured to perform IP bridging. With IP bridging, only 1 B channel can be used and therefore the maximum throughput is 64 kbit/s. Figure 5.3 shows the configuration. File transfer was established using FTP. The net throughput was measured at an average of 57 kbit/s.

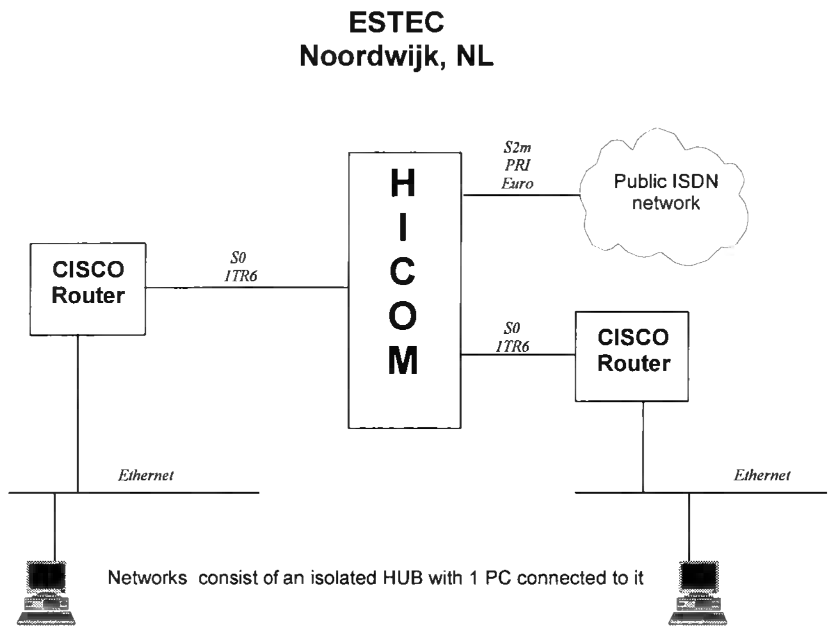

Figure 5.3 *64 kbit/s LAN-to-LAN link*

To obtain a throughput of 128 kbit/s, the configuration of figure 5.4 was established. The Controlware TAXI inverse multiplexers provided the 128 kbit/s X.21 channels over ISDN. The net throughput increased to an average of 112 kbit/s.

The functionality of the routers showed no errors. A call is established automatically as there is traffic to be send. However, with the configuration of figure 5.4, the functionality depends on the Controlware TAXI inverse multiplexers.

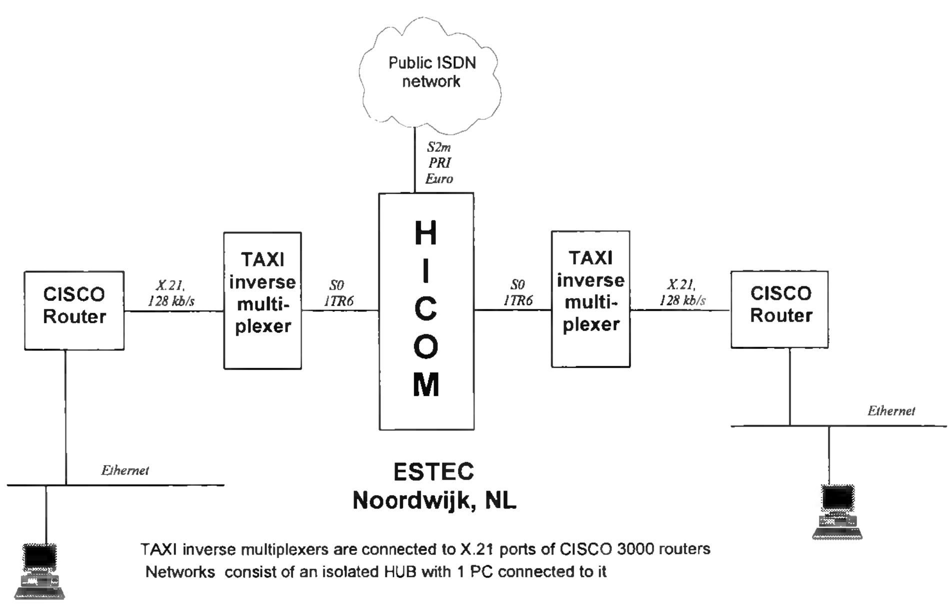

Figure 5.4 *128 kbit/s LAN-to-LAN link.*

5.4.5 Dial-in service

The router configuration of figure 5.5 was established. To measure performance, 1 megabyte files were sent with FTP.

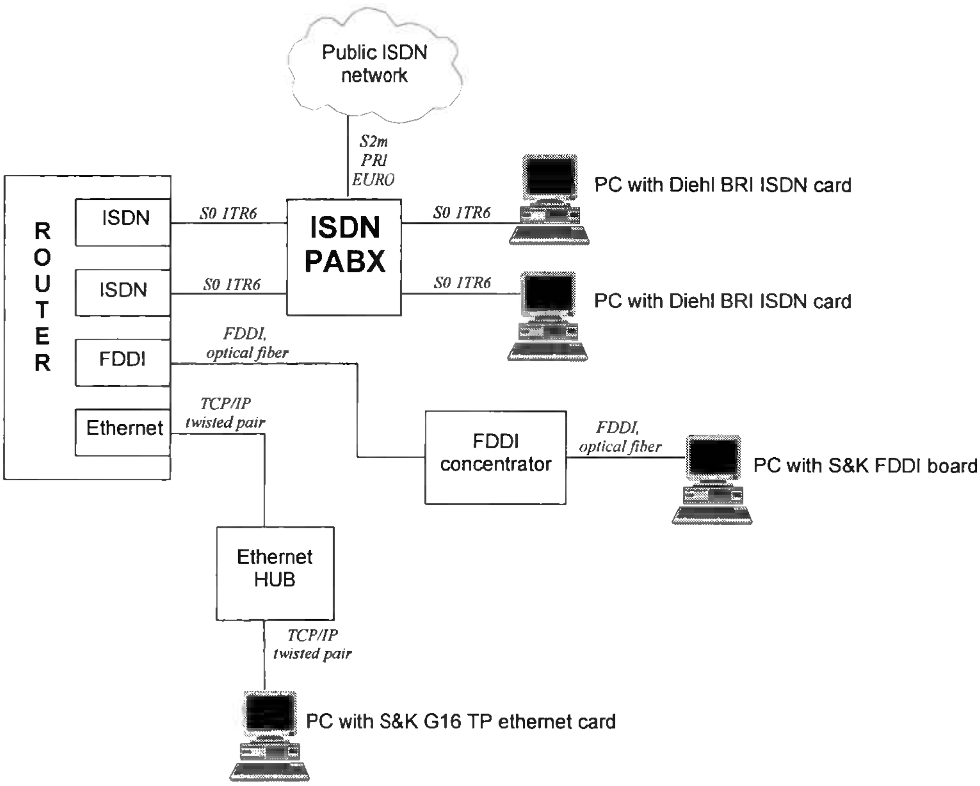

Figure 5.5 *ISDN remote dial-in configuration*

The following combinations were tested:

From:	To:	Average throughput [kbit/s]:
FDDI	Ethernet	815
Ethernet	ISDN PC1	56
Ethernet	ISDN PC2	54
FDDI	ISDN PC1	56
FDDI	ISDN PC2	56
ISDN PC1	ISDN PC2	49

Simultaneous sending from Ethernet to ISDN PC1 and from FDDI to ISDN PC2 resulted in 55 and 56 kbits/s throughput respectively. It is not possible to use two B-channels on the PC to provide 128 kbit/s bandwidth. ISDN-to-ISDN routing was considerably slower than the other combinations. However, this type of transfer will not be used in a practical setup since it is a dial-up service where the ISDN users are temporarily users.

Chapter 6 provides a theoretical analysis on the throughput that can be achieved on LAN-to-LAN connections.

5.4.6 Multimedia applications

Since equipment was not yet available, this subject is pending. A demonstration of one vendor's equipment displayed a good video quality over 128 kbit/s and a number of functional options to provide collaborative working/editing.

Chapter 6. Theoretical analysis of user data throughput on 64 kbit/s links, using Ethernet and HDLC framing and TCP/IP and SPX/IPX protocols

See Appendix B for a detailed overview of all different network protocols and frame formats.

6.1 Introduction

As described in chapter 5, practical experiments were performed using ISDN links as LAN interconnections. In this part, an analysis on the theoretical throughput that can be achieved is done. Besides the analysis of the configurations tested, also other protocols and media will be investigated.
As a first assumption, the ISDN B-channel will be considered error-free. The practical experiments have proven that the channel is almost error-free, especially on local connections.

The following variables are defined [14]:

A = number of bits in an ACK frame
C = channel capacity (bit/s)
D = number of data bits per frame
F = D+H (total frame length)
H = number of bits in the frame header
L = probability that a frame or its ACK is lost or damaged
I = interrupt and service time + propagation delay
U = efficiency of channel
W = window size

The efficiency of a specific protocol combination is defined as the quotient of real user data within a frame and total amount of data within the frame. The real user data is defined as the number of data bits within the layer 2 (datalink layer) packet of the network protocol.

For the error-free channel with large window size obviously the efficiency is given by

$$U=\frac{D}{H+D} \tag{6.1}$$

The efficiency of several different configurations will be calculated using this formula.
If we multiply the efficiency with the channel capacity (64 kbit/s for a single ISDN line) then the real user throughput on the ISDN line remains.

With all protocols the assumption is made that higher-layer packets are not fragmented across lower layer packets and that lower layer packets always carry a maximum of one higher-layer packet.

6.2 Internetwork connections with the TCP/IP protocol suite

The TCP protocol provides a reliable end-to-end datastream using positive acknowledgement with retransmission and using a sliding window. Although the size of the sliding window changes continuously[1] during transmission, it will be assumed constant.

6.2.1 Ethernet-to-Ethernet bridging via an ISDN link

The Ethernet frame contains 208 bits of overhead and 368-12000 bits of data (see appendix A). The data-part of the frame contains higher-level packets.
On an Ethernet, there is a delay between packets of about 9.6 μs. On the ISDN link, the Ethernet frames are encapsulated within HDLC frames. Since the bandwidth of an ISDN line is only 64 kbit/s, the Ethernet packets can be delivered fast enough to create a continuous stream of HDLC packets on the ISDN line. Figure 6.1 shows how the Ethernet frames are bridged through the ISDN link. The HDLC overhead is assumed to be 64 bits per packet, although in some implementations this may be 48 bits. This difference occurs because the frame check sequence (see Appendix B) may either consist of two or four bytes.

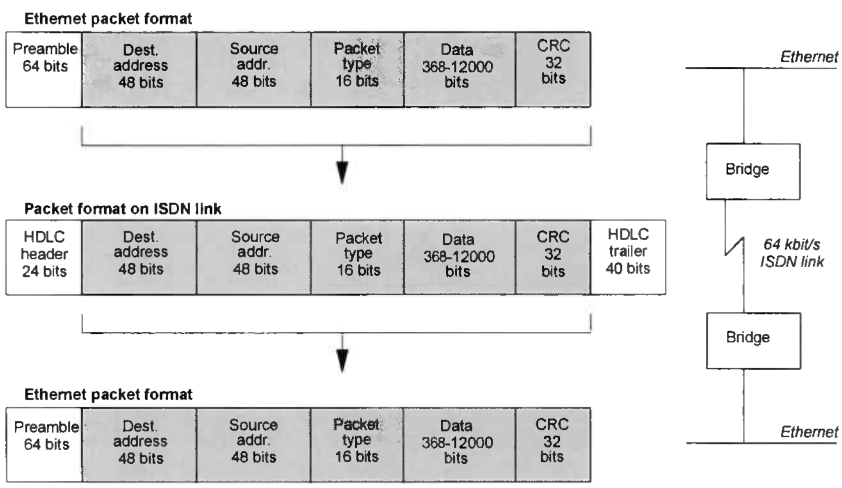

Figure 6.1 *Bridging of Ethernet frames through an ISDN circuit.*

Both the IP packet and the TCP segment contain 192 bits of overhead and an undefined amount of data bits. During actual file transfer with FTP, the data part of the TCP segment contains the user data. Therefore, the total overhead on a (HDLC) packet travelling across the ISDN network is:

HDLC overhead (64 bits) + Ethernet overhead (144 bits) + IP overhead (192 bits) + TCP overhead (192 bits) = 592 bits.

Apparently the efficiency of the above described protocol stack depends on the number of data bits within the Ethernet frame. Figure 6.2 shows the efficiency as a function of the size of the HDLC frames travelling across the ISDN network.

[1] It is adapted to line quality and delay

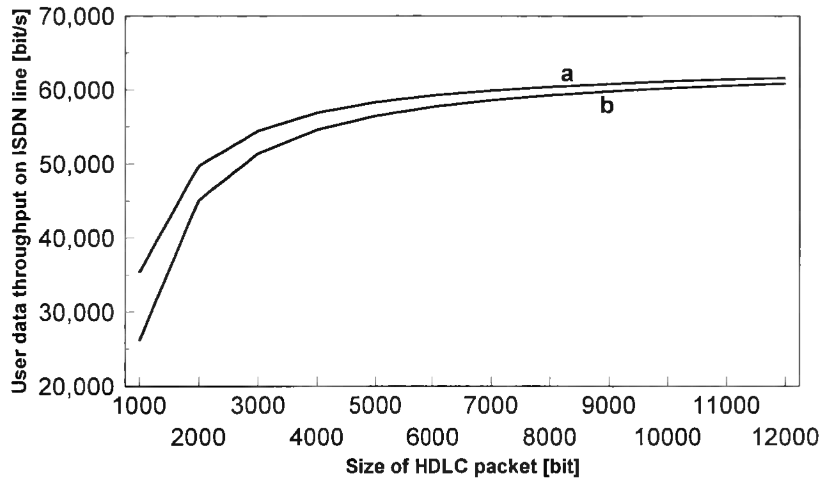

Figure 6.2 *Net user data throughput on a 64 kbit/s ISDN link using Ethernet bridging and routing and TCP/IP protocols as a function of the size of the HDLC packet. (a) routing; (b) bridging.*

It can be concluded that the maximum theoretical user data throughput on an ISDN 64 kbit/s line using Ethernet bridging and TCP/IP protocols exceeds 60 kbits/s.

6.2.2 Ethernet-to-Ethernet routing via an ISDN link

Figure 6.3 shows how Ethernet frames are routed through an ISDN link:

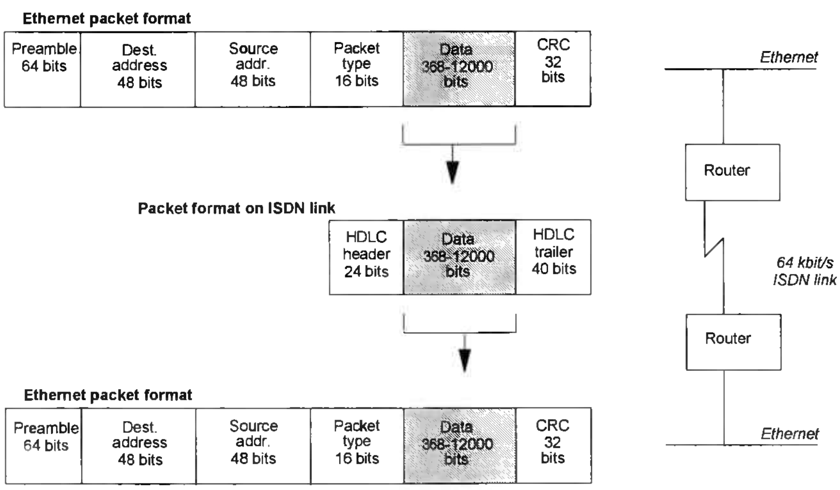

Figure 6.3 *Routing of Ethernet frames through an ISDN line.*

It can be concluded that the overhead on the ISDN link is smaller than the above described bridging method. Therefore, in theory, the efficiency should be slightly better. Figure 6.2 shows the efficiency. With the dial-in router showed in figure 5.5, there is Ethernet or FDDI on one side, but on the side of the standalone PC there is no network media. What will happen is that the Data part of the Ethernet or FDDI frame will be encapsulated within an HDLC packet as shown in figure 6.3, and at the standalone PC the HDLC header and trailer will be stripped off and the data part will be processed without encapsulating it in for example an Ethernet frame. The theoretical throughput on the ISDN link remains the same as shown in figure 6.2 (a), since the overhead on the ISDN line is the same.

6.2.3 ISDN PC-to-FDDI routing via an ISDN link (see figure 5.5)

The above outlined principle also holds for this configuration: The PC encapsulates the TCP/IP datagram in an HDLC packet and the router puts the data part of the HDLC packet in an FDDI frame. In the other direction, the router puts the data part of the FDDI frame in an HDLC packet, and the standalone PC distracts the data part from the HDLC packet and processes the TCP/IP datagram. The efficiency on the ISDN link is the same as the above described situation.

6.2.4 Throughput on non-ideal ISDN connections using sliding window protocols

In the paragraphs 6.2.1-6.2.3 the ISDN channel was assumed to be ideal. However, in practical situations, there may be errors and delays on the line. In theory, the delay should be unimportant: if we consider the traffic to travel with a speed of 2/3 c, and the framelength 1600 bits, and a length of the ISDN line of 5000 km, the number of frames on the line is 1. However, the actual frame length will be larger and the size of the sliding window is 8 with the HDLC protocol. Therefore, it is not expected that the sender has ever to wait for an acknowledgement before he can continue sending data.

On national ISDN connections, British Telecom specifies a maximum delay of 13 ms from NT to NT.

However, on international ISDN connections there may be large propagation delays. British Telecom specifies a maximum of 260 ms from NT to NT.This is the delay due to a satellite hop (Satellite at at distance of 39000 km to both sender and receiver, propagation speed of $3*10^8$). Also, router manufacturers implement proprietary versions of the HDLC protocol with variable window sizes and packet lengths. These routers often automatically tune the parameters based upon the transmission speed and a zero delay line. Still, they often keep a very low window size because of a 'fear of errors' and the amount of retransmissions should an error occur.

The error performance of ISDN can in the worst case degrade to some 10^{-7} bits/s. With a frame size of 10000 bits this would mean that one out of every 1000 frames will have an error. The expected number of transmissions is 1/(1-L) [14], so to receive W frames without error, W/(1-L) of them must be transmitted. With a large window size, formula (6.1) becomes:

$$U=\frac{D}{H+D}*(1-L) \qquad (6.2)$$

With the given worst case error performance, the efficiency would only be decreased by 1 ‰. Mostly, the error performance is much better, thus the efficiency suffers hardly from bit errors. Therefore, there is no reason to keep the window size low.

If the window size is small (W< 1 + 2CI/F), and there are no errors, the following formula for the efficiency holds [14]:

$$U = \frac{D}{H+D} * \frac{W}{1 + \frac{2CI}{H+D}} \tag{6.3}$$

In the table below, some values for the second factor in (6.3) are calculated using different values for I and the framesize.

H+D (framesize)	W (windowsize)	I [s] (prop. delay)	decreasing factor
1000	8	0.1	0.58
1000	8	0.26	0.23
1000	2	0.013	0.75
1000	2	0.26	0.06
5000	3	0.1	0.84
5000	3	0.26	0.39
5000	1	0.013	0.75
5000	1	0.26	0.13
10000	1	0.013	0.86
10000	1	0.26	0.23

It can be concluded that with small frame sizes and large transmission delays the efficiency of sliding window protocols degrades severely, even with window sizes of 8. When using frame sizes of ten thousand bits, the 'delay problem' can be solved easily by using any window size larger than 1.

6.3 Internetwork connections with the IPX/SPX protocol suite

Whilst the TCP protocol is using a sliding window, the SPX protocol is usually implemented only providing positive acknowledgement with retransmission, which means it has to wait for the acknowledgement before it will send the next packet. In some implementations, the sender sends a burst of packets and then waits for an acknowledgement for these packets. The whole burst of packets is then acknowledged by a single acknowledgement.

The first case is a stop-and-wait protocol. Apparently the performance of such a protocol is very dependent on the transmission delay on the line. If we assume the transmission line to be error free, then the following holds:

The sender starts sending a frame at t=0. At time F/C the last bit has been sent. At time F/C+I the last bit has arrived at the receiver and the receiver is ready to start sending the acknowledgement frame. At time (F/C) + I + (A/C) + I the sender has processed the acknowledgement and is ready to send the next data frame. The bandwidth occupied by this frame is C multiplied by the time taken, which is F + A + 2CI. The number of data bits actually transferred is D, so the channel efficiency is

$$U = \frac{D}{H + D + A + 2CI} \tag{6.4}$$

If the efficiency is multiplied by the channel capacity then the result is the amount of transmitted real user data bits per second out of the channel capacity.

6.3.1 Ethernet-to-Ethernet bridging via an ISDN link

See figure 6.1 for the principle of bridging.

The IPX datagram has 240 bits overhead and 0 to 4368 bits of data. The SPX datagram contains 96 bits overhead and 0 to 4272 bits of data. However, if the drivers are optimized by Novell, the size of the IPX and SPX datagrams may be larger. Therefore calculations have been made with HDLC packet sizes up to 12000 bits.

The total overhead within an HDLC frame is:

> HDLC overhead (64 bits) + Ethernet overhead (144 bits) + IPX overhead (240 bits) + SPX overhead (96 bits) = 544 bits.

The length of an SPX acknowledgement is 96 bits (a typical file transfer situation has only one-way traffic, so the acknowledgements travel in data-less SPX packets). Encapsulated in an HDLC packet this HDLC packet will have a size of 544 bits.

6.3.2 Ethernet-to-Ethernet routing via an ISDN link

Figure 6.3 gives the principle for routing. As described in paragraph 6.2.2, routing should give a slightly better performance over bridging since the total overhead per HDLC packet is smaller. The total overhead within an HDLC frame is now:

> HDLC overhead (64 bits) + IPX overhead (240 bits) + SPX overhead (96 bits) = 400 bits.

6.3.3 Ethernet-to-FDDI routing via an ISDN link

See paragraph 6.2.3.

6.3.4 Throughput on non-ideal ISDN connections

As described in paragraph 6.2.4, ISDN links and for example satellite links may cause serious transmission delays. Equation 6.3 shows that this will have significant influence on the performance while using the stop-and-wait protocol. The influence of errors is much less. The expected number of retransmissions is L/(1-L). The efficiency then becomes

$$U = \frac{D}{\left(\frac{L}{1-L}\right)(F+CT)+(F+A+2CI)} \quad (6.5)$$

If the receiver's service time is neglected, the sender can set its timeout interval just above the time required for the acknowledgement to arrive: T = A/C + 2I. But even with a transmission delay of 0.3s, and a 'L' of 0.001, the first factor in the denominator of (6.5) is neglectible in comparison with the second factor and therefore errors will be neglected.
Table 6.2 shows what the influence of transmission delays is when using the stop-and-wait protocol. Ethernet bridging and errorfree transmission are assumed.

Size of HDLC frame	Transmission delay [s]	Throughput of user data [kbit/s]
1000	0.013	9.1
1000	0.1	2.0
1000	0.26	0.8
4000	0.013	35.6
4000	0.1	12.8
4000	0.26	5.8
12000	0.013	51.6
12000	0.1	28.9
12000	0.26	16.0

Table 6.2 *Influence of transmission delays on network performance*

6.3.5 Performance with a packet burst-and-wait protocol

As mentioned before, in some implementations of SPX there is not an acknowledgement for each SPX packet sent, but one acknowledgement after a burst of packets. This method can improve the performance. If the number of packets per acknowledgement is n, the channel efficiency becomes

$$U = \frac{nD}{n(H+D)+A+2CI} \quad (6.6)$$

Figure 6.4 shows the performance for all the above described situations for Ethernet bridging. Figure 6.5 shows the performance for all the above described situations for Ethernet routing. In both figures, bit errors are not taken into account.

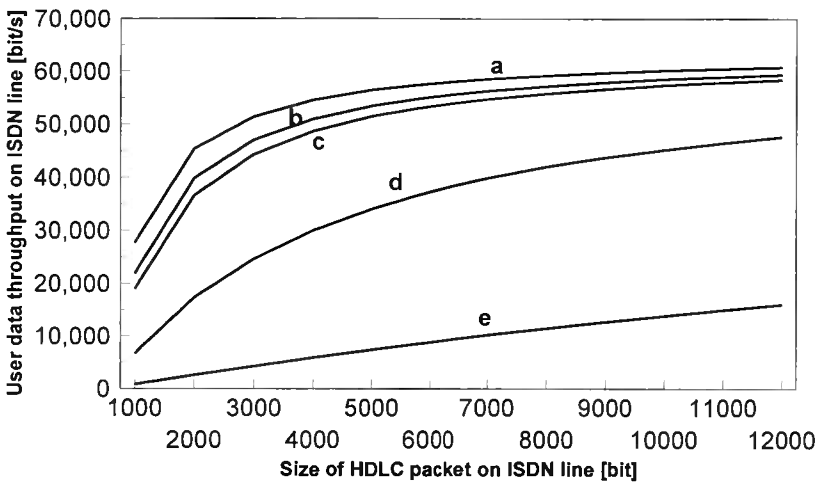

Figure 6.4 *Net user data throughput on a 64 kbit/s ISDN link using Ethernet bridging and IPX/SPX protocols as a function of the size of the HDLC packet. (a) burst mode, 1 ACK per 10 SPX packets, no delay; (b) burst mode, 1 ACK per 100 frames, 0.26s delay; (c) stop-and-wait mode, 1 ACK per SPX packet, no delay; (d) burst mode, 1 ACK per 10 SPX packets, 0.26s transmission delay; (e) stop-and-wait mode, 1 ACK per SPX packet, 0.26s transmission delay.*

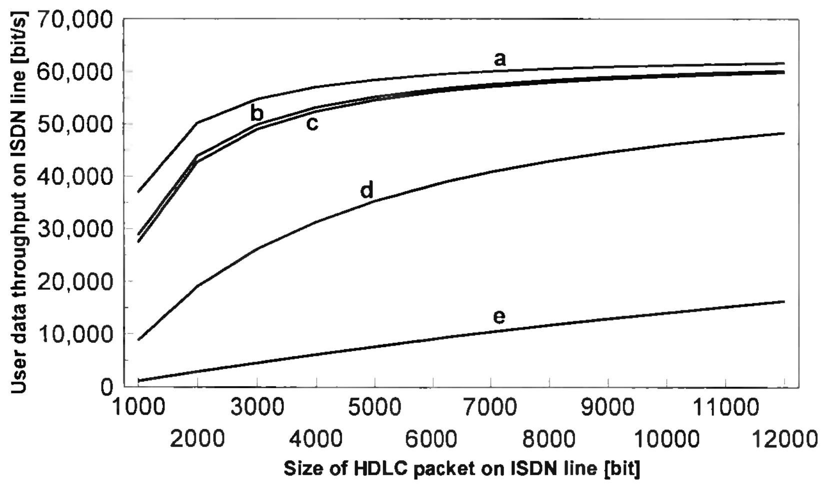

Figure 6.5 *Net user data throughput on a 64 kbit/s ISDN link using Ethernet routing and IPX/SPX protocols as a function of the size of the HDLC packet. (a) burst mode, 1 ACK per 10 SPX packets, no delay; (b) burst mode, 1 ACK per 100 frames, 0.26s delay; (c) stop-and-wait mode, 1 ACK per SPX packet, no delay; (d) burst mode, 1 ACK per 10 SPX packets, 0.26s transmission delay; (e) stop-and-wait mode, 1 ACK per SPX packet, 0.26s transmission delay.*

If errors are considered, the following formula is valid:

$$U = \frac{nD}{\frac{L}{1-L}(n(H+D)+CT)+n(H+D)+A+2CI} \qquad \textbf{(6.5)}$$

With a burst of 100 frames, a framesize of 10000 bits, and a random bit error rate of 10^{-7} the value for L is 0.1. It is obvious that the bit errors will have a significant influence on the performance. Table 6.3 shows the influence of bit errors and transmission delays on the performance. Bit errors are assumed to occur randomly; ethernet bridging and IPX/SPX protocols are used.

HDLC packet size [bit]	Burst size [HDLC packets]	Transm. delay [s]	Throughput on error free 64 kbit/s line [kbit/s]	Throughput with ber of 10^{-7} [kbit/s]	Throughput with ber of 10^{-9} [kbit/s]
4000	10	0	54.6	52.4	54.6
4000	10	0.013	52.4	52.2	52.4
4000	10	0.26	30.0	29.8	30.0
4000	100	0	55.2	53.0	55.2
4000	100	0.013	55.0	52.8	55.0
4000	100	0.26	51.0	48.9	51.0
10000	10	0	60.2	59.6	60.2
10000	10	0.013	59.2	58.6	59.2
10000	10	0.26	45.2	44.8	45.2
10000	100	0	60.5	54.4	60.4
10000	100	0.013	60.4	54.3	60.3
10000	100	0.26	58.5	52.7	58.5

Table 6.3 *Influence of bit errors and transmission delays on network performance.*

It can be concluded that transmission delays have more influence on the network performance than bit errors. In the table, errors are assumed to occur randomly, whilst in reality the errors will occur in bursts.

6.4 PC-to-PC file transfer

The PC-to-PC file transfer software packages use the HDLC encapsulation of data. There is no layer three information, so the real user data is put into an HDLC frame. The maximum packet size used is 2048 bytes, or 16384 bits. Figure 6.6 gives the theoretical throughput.

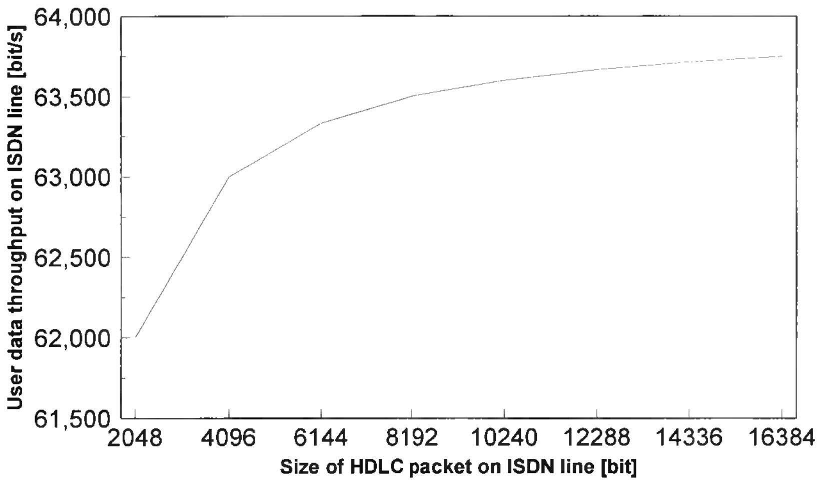

Figure 6.6 *PC-to-PC file transfer throughput on an ISDN line with HDLC framing.*

6.5 Comparison between theoretical findings and measured results

6.5.1. Network performance

The experiments performed were all done with the TCP/IP protocol stack. With the CISCO routers, a maximum user data throughput of 57000 bit/s was achieved. The maximum theoretical user data throughput is 60816 bit/s with IP bridging and 61623 bit/s with IP routing (see figure 6.1, 12000 data bits per ethernet frame). In all the calculations, the processing time of the equipment was neglected. From the difference between the theoretical and practical results the estimated processing time can be calculated. When the processing time is taken into account, formula 6.1 changes into:

$$U=\frac{D}{H+D+CI} \qquad (6.6)$$

With the given theoretical and practical values, the processing time per packet equals 13ms within the CISCO 3000 routers and 19 ms with the S&K multiprotocol router. Real values for processing time lie in the 10μs-10ms range. Since the S&K MPR is based on a PC, the high value for the processing time lies in the fact that delays occur within the PC because of the architecture of the PC and other tasks it is performing. The CISCO 3000 router is a dedicated router, and the high value of the processing time lies in the fact that the router is sending routing information besides user information.

6.5.2 Performance of PC-to-PC file transfer

The experiments done showed a maximum throughput of 63 kbit/s. The maximum theoretical throughput is 63750 bit/s (HDLC packet size of 16384 bits). This leads to a estimated processing time of 3 ms per HDLC packet. This value is realistic, since the filetransfer is performed between PCs.

6.6 Conclusions from the theoretical findings

The maximum measured throughput of user data on a network-to-network link over an ISDN link was 57 kbit/s. The maximum theoretical throughput exceeds 60 kbit/s. This is due to the following factors:

- The processing time in the router or PC is in reality not zero, but takes a considerable amount of time[1]. This is the most significant factor that influences the actual throughput.

- Besides HDLC dataframes, there are also some supervisory frames transmitted, even within a data transfer session.

- The HDLC protocol uses bit stuffing for data transparency. This means that bits are added to the data. However, this does not occur frequently, since one additional bit is stuffed when the pattern 0111111 appears within the data.

With the TCP/IP protocols good performance can be achieved as long as the window size and packet length are not too small. Under normal conditions, the maximum allowed packet and frame sizes will be used. The TCP/IP protocol usually has a window size of 8.

On satellite links, delays can occur[2]. However this will have small impact on the performance of the TCP/IP link as long as the packet and window sizes are not too small.

Since the IPX/SPX protocol software usually uses the stop-and-wait mechanism, the performance may be very poor, especially when the packet size is small. When delays occur on the line, the performance degrades severely because the sender is idle most of the time waiting for acknowledgements. The performance can be increased using a 'burst-and-wait' protocol, where multiple data packets are acknowledged by a single acknowledgement packet. Using this technique in combination with a large packet size results in a reasonable performance.

ISDN connections have proven to be almost error free. Therefore possible errors will have neglectible influence on the performance. Since errors usually occur in bursts, they will only affect a single frame. Because of this low occurrance of errors, there is no reason to keep low window sizes and/or small packet sizes. Increasing both values increases the performance of

---

[1] 10 μs - 10 ms per packet

[2] The expected delay on a satellite link is 260 ms per hop, i.e. 520 ms round trip delay.

network protocols over ISDN lines. Calculations have shown that even with a random bit error rate of 10^{-7} bit/s and a burst length of 100, the errors have only small influence on the performance.

With PC-to-PC file transfer, the maximum measured throughput was 63 kbit/s[1] (without datacompression and channel bundling). The maximum theoretical throughput exceeds 63.5 kbit/s. Again the most significant reason for this difference is the processing time needed by the PCs.

[1] The file transfer software used a proprietary data transmission protocol

Chapter 7. Conclusions and recommendations

7.1 Conclusions

The availability of international 64kbit/s ISDN circuits has proven to be satisfactory. Most of the call-attempts carried out during the testing were successful. Within the Netherlands, 100% of the call attempts were successful. Between the Netherlands and Germany, 99.1% of the call attempts were successful.

The quality of 64kbit/s ISDN circuits is good. The performed bit error rate tests show that the network is error free for more than 99.9% of the time. Errors occur often in bursts. The error performance is well within the limits given by the ITU-TSS G.821 specification, even if the large bursts of errors that sometimes occurred are considered to be 'normal' errors instead of controlled slips that do not count to the error performance objectives.

The time needed to establish a call differred between national and international calls. On national calls, call establishment took around 2 seconds. On international connections, the time needed was around 10 seconds. These times could not be measured in an accurate way, since the calls were established with inverse multiplexing equipment. This equipment needs some 10 seconds after call establishments to synchronize.

Some established calls were dropping during bit error rate tests. The reason why this happened is unknown. Since monitoring by the German PTT did not show any errors, it is assumed that the call drops were caused by the Controlware TAXI inverse multiplexers used for the tests. One reason for the call-drop may be that the Controlware TAXI inverse multiplexers cannot handle short bursts of errors, which appeared on the international ISDN connections that were used for the tests. A test with Controlware CITAM terminal adapters between ESTEC (Noordwijk, NL) and DLR (Köln, D) showed no call drops, which leads to the conclusion that the Controlware TAXI inverse multiplexers may not always be capable of coping with bit error bursts and consequently sometimes fail.

The Controlware TAXI inverse multiplexers tested still show some shortcomings in terms of reliability. After intensive use, they sometimes need a hardware reset to function properly again. Investigations on the reason for call drops have to be done.

The tests with Ascend inverse multiplexers showed that these were much more flexible than the Controlware TAXI units. Their BONDING mode 1 compliance guarantees compatibility, although this has not been tested. The call setup is more convenient, and the possibility to add bandwidth to the established connection by dialling more B-channels is useful. Also the error recovery is better than the error recovery of the Controlware TAXI units. However, one ASCEND unit used for tests at ESA/CNES in Toulouse caused a total ... of the PABX.

The bit error rate tests over an internal PABX link showed that some the bit errors (appr. 1 in 10^9) are caused by the combination of the terminal equipment and the PABX. These bit errors were not expected on this kind of short interconnection and point to the conclusion that the actual physical layer interfaces are not yet trouble-free.

On the tests between France and The Netherlands some problems occurred in call setups, but once a reset of the PBX at France was done this resulted in 100% success rate.

In general it can therefore be concluded that most of the connectivity problems are caused by the equipment. Some equipment is not very stable (PBX, inverse multiplexers) and it seems that some of the ISDN protocols (VN3) are not implemented properly within the terminal equipment and inverse multiplexing equipment.

The cost break-even analysis between ISDN and leased lines shows that on international connections, an ISDN B channel can be used from 5 to 8 hours per working day for the price of a leased line. Since most leased lines are used only a few hours a day, the ISDN can be an economic alternative.

However, for 2 Mbit/s connections leased lines are still the most favourable solution. If the full bandwidth is used for more than 3 hours per working day, a leased line is cheaper than an ISDN primary rate connection.

It is obvious that the fact that ISDN gives much more flexibility is not taken into account with the cost analysis. Therefore, for each company the consideration to choose for either ISDN or a leased line will be different. Especially for 64 kbit/s connections, ISDN seems to be an economic solution.

The most important application on the ISDN will be telephony. Although this service will have some additional features in comparison with telephony on the PSTN like calling line identification and probably a better quality in the future when 7 KHz telephony is available, the average user will probably not recognize this as a big improvement. The most interesting applications are in the data communications sector.

The new fax group IV standard offers high speed and good quality in comparison with the fax group III. Since many people complain about the quality and the thermal paper often used with fax group III, it could be expected that laser printer fax group IVs will become very popular. However, since the equipment is still very expensive and since most of the faxes in use are fax group III, it is expected that the fax group IV equipment will not be used on a large scale during the next few years.

The ISDN offers many possibilities for internetworking of computers. Various equipment is available to connect LANs over long distances, and to provide remote connectivity to single users. When connecting LANs via ISDN links, care has to be taken for the data filtering of the routers. If all traffic on a network causes the router to call the remote LAN, the ISDN lines will be online for a large percentage of the time and this will void the saving aspects.

LAN-to-LAN connections have been tested using CISCO routers. An internal test in ESTEC showed good performance (net user data throughput up to 57 kbit/s over a single B channel). The network of NLR in Amsterdam was linked with an ESTEC network to provide data connectivity for telescience experiments. These experiments were concluded successful.

To provide remote connectivity to the ESTEC LANs for remote users, analog modems are used. This method is slow and unreliable. Therefore a dial-in router[1] was built and configured to provide remote access over the ISDN. The configuration functioned without severe errors, although the software was still in beta-status and sometimes halted on the remote PCs. The multiprotocol router provided access to both an FDDI and an Ethernet simultaneously. The achieved net user throughput on the ISDN lines was 56 kbit/s.

On satellite links, severe delays can occur. However this will have small impact on the performance of the TCP/IP connection as long as the packet and window sizes are not too small. With packet sizes of more than 4000 bits and a window size of 8, good performance can be achieved.

No experiments were done with the IPX/SPX protocol suite. The theoretical analysis showed that also with this protocol suite good performance can be achieved (net user data throughput over 60 kbit/s) if the so-called burst method is used.

The errors that occur on ISDN lines do not have much impact on the performance. Since errors occur usually in bursts, they will affect only a single packet. Because of the low error rate, there is no reason to use small window and/or packet sizes, as only few retransmission have to take place.

7.2 Recommendations

Experiments have shown problems with both terminal adapters and inverse multiplexers of the units under tests. Especially within France problems occurred. Therefore it is advised to perform tests with equipment from other manufacturers and to analyse the D channel signalling in order to trace the reason for the problems.

On local connections via internal calls, some bit errors occurred. Some more testing is required in order to detect which equipment is responsible for these errors (i.e. ISPBX or terminal equipment).

The multiprotocol router that was built and tested used proprietary software. Novell has released its "Multiprotocol Router for ISDN" software. Manufacturers like AVM, Diehl and ITK have implemented this MPR software for their ISDN PC cards. These MPR solutions will be tested in future.

With compression techniques, up to 10 voice calls can be put on a single 64 kbit/s connection. It is advised that ESA evaluates this compression equipment in order to reduce telephone costs between the various ESA sites.

[1] Schneider & Koch Multi Protocol Router

In order to get an impression about the time needed for international call establishments, it is advised to procure some ISDN telephones and measure the times needed to establish calls between various countries.

References

[1] H. Ekkelenkamp, I. Verstraaten, R. Wijbrands,
Het ISDN boek.
Amsterdam: Tutein Nolthenius 1992.

[2] H.J. Helgert,
Integrated Services Digital Networks : architectures, protocols, standards.
New York: Addison-Wesley publishing company, Inc., 1991.

[3] CCITT,
"Recommendation X.1: International user classes of service in public data networks and integrated services digital networks (ISDNs)",
Blue Book, Vol.VIII.2, Geneva: ITU, 1989.

[4] CCITT,
"Recommendation I.430: Basic user-network interface: layer 1 specification",
Blue Book, Vol.III.8, Geneva: ITU, 1989.

[5] CCITT,
"Recommendation G.703: Physical/Electrical characteristics of Hierarchical Digital Interfaces",
Blue Book, Vol. III.4, Geneva: ITU, 1989.

[6] CCITT,
"Recommendation Q.930: ISDN user-network interface layer 3 - general aspects",
Blue Book, Vol.VI.11, Geneva: ITU, 1989.

[7] CCITT,
"Recommendation Q.920: ISDN user-network interface data link layer -general aspects",
Blue Book, Vol.VI.10, Geneva: ITU, 1989.

[8] CCITT,
"Recommendation X.25: Interface between data terminal equipment (DTE) and data circuit-terminating equipment (DCE) for terminals operating in the packet mode and connected to public data networks by dedicated circuit",
Blue Book, Vol. VIII.3, Geneva: ITU, 1989.

[9] CCITT,
"Recommendatio Q.931: ISDN user-network interface layer 3 specification for basic call control",
Blue Book, Vol. VI.11, Geneva: ITU, 1989.

[10] CEPT,
"Memorandum of understanding for the implementation of an European ISDN service by 1992",
Computer Networks and ISDN systems, no. 21, pp. 69-74, 1991.

[11] Eurie '93 Organization,
Eurie '93 Handbook, a User's guide to Euro-ISDN,
London, Ovum Ltd., 1993.

[12] CCITT,
"Recommendation G.821: Error performance of an international digital connection forming part of an integrated services digital network",
Blue Book, Vol. III.5, Geneva: ITU, 1989.

[13] CCITT,
"Recommendation G.822: Controlled slip rate objectives on an international digital connection",
Blue Book, Vol. III.5, Geneva: ITU, 1989.

[14] A.S. Tanenbaum,
Computer Networks, 2nd edition.
New Jersey: Prentice Hall, Inc., 1988.

[15] Telecommunications Industries Association (TIA),
"Aggregation of multiple independent 56 kbit/s or 64 kbit/s channels into a synchronized wideband connection",
Standards proposal No. 3014 (if approved, to be published as TIA/EIA-619), December 1993.

[16] D. Comer,
Internetworking with TCP/IP : principles, protocols and architecture.
New Jersey: Prentice Hall, Inc., 1988.

Appendix A. Results of tests performed on international ISDN connections

A.1 Explanation of tests and test equipment

Tests:

1) Availability of ISDN circuits.

ISDN call made:

This test checked availability of ISDN channels. Calls were defined as successful if connectivity was established within 1 minute.

2) Functionality of Controlware TAXI inverse multiplexers.

ISDN call timed out:

Check the call is dropped after a disconnection request by the inverse multiplexers, activated by defined conditions on the X.21 data port.

ISDN call re-establish after time-out:

Checks the call is automatically re-established when a re-connection request is received by the Controlware TAXI inverse multiplexer, again via the X.21 data port.

ISDN circuit down on failure:

Check the effects of a specified ISDN failure caused by physically breaking the circuit at the S_0 reference point.

ISDN circuit recovers after failure:

Checks the operation capability following reconnection of the S_0 circuit break.

3) Bit error rate test of synchronous channel aggregated circuits.

Bit error rate test:

A check on the quality of the 64-384 kbit/s synchronous channel, using a Network Probe 7100 as a data generator and error detector, with a physical loopback provided on the remote Controlware TAXI X.21 interface.

Test equipment:

At ESTEC, an ISDN PABX[1] was used to provide basic rate interfaces. Terminal equipment used were inverse multiplexers[2], terminal adapters[3], and a test device[4] to perform bit error rate tests and D-channel decode. For the tests with France, inverse multiplexers from Ascend were used.

A.2 Test results

A.2.1 Local ESTEC connection via internal call (see figure A.1).

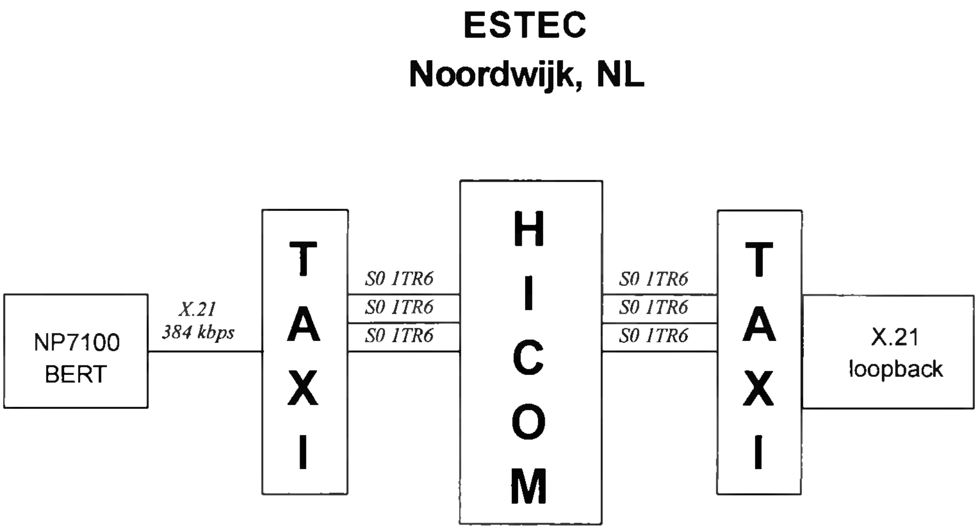

Figure A.1 *Local ESTEC connection via internal call*

Availability of ISDN lines

No.	Availability tests	Nr. of B-channels	Nr. of attempts	Start time	Nr. of successes	% of successes
01.01	ISDN call made	6	48	n.a.	6*48=288	100

[1] Siemens HICOM 340 PABX

[2] Controlware TAXI with 3 BRIs and one X.21 interface

[3] Controlware CITAM terminal adapter with 1 BRI and two X.21 ports

[4] NCC Network Probe 7100

Functionality of Controlware TAXI inverse multiplexers

No.	TAXI inverse multiplexers functionality tests	Nr. of B-channels	Nr. of attempts	Start time	Nr. of successes	% of successes
01.02	ISDN call timed out	6	12	n.a.	12*6=72	100
01.03	Call re-established after time-out	6	12	n.a.	12*6=72	100
01.04	ISDN circuit down on failure	6	12	n.a.	12*6=72	100
01.05	ISDN circuit recovers after failure	6	12	n.a.	12*6=72	100

Remarks:
On the internal connection all the tests were 100% successful. Since the ISDN PABX was dedicated to the tests, the times at which the tests took place are not relevant.

Quality tests (Bit error rate tests)

Test description	ESTEC-ESTEC ISDN connection via internal calls, 6 B-channels aligned with TAXI inverse multiplexers		
Nr.: T311011			
Start time	30-11-93 17:25:17	Error free seconds	65242
End time	01-12-93 11:33:20	% Error free secs	100.00
Elapsed time	0 Days 18:07:22	Errored seconds	0
Bit errors	0	% Errored secs	00.00
Block errors	0	Sev. errored secs	0
Bit count	25067963662	% Sev. errored secs	00.00
Block count	397904185	Degraded minutes	0
Out of sync	0	% Degraded mins	00.00
Out of sync secs	0	Unavailable secs	0
Bit error rate	0	% Unavailable secs	00.00
BPS clock	384000	Available seconds	65242
		% Available secs	100.00

Test description	ESTEC-ESTEC ISDN connection via internal call,		
Nr.: T312061	6 B-channels aligned with TAXI inverse multiplexers		
Start time	03-12-93 11:59:18	Error free seconds	267844
End time	06-12-93 14:27:59	% Error free secs	100.00
Elapsed time	3 Days 02:24:14	Errored seconds	10
Bit errors	20	% Errored secs	00.00
Block errors	12	Sev. errored secs	0
Bit count	102911329816	% Sev. errored secs	00.00
Block count	3140700	Degraded minutes	0
Out of sync	0	% Degraded mins	00.00
Out of sync secs	0	Unavailable secs	0
Bit error rate	1.94e-10	% Unavailable secs	00.00
BPS clock	384000	Available seconds	267854
		% Available secs	100.00

Test description	ESTEC-ESTEC ISDN connection via internal call,		
Nr.: T312131	6 B-channels aligned with TAXI inverse multiplexers		
Start time	10-12-93 16:19:18	Error free seconds	254148
End time	13-12-93 14:58:20	% Error free secs	99.98
Elapsed time	2 Days 22:36:34	Errored seconds	46
Bit errors	304	% Errored secs	00.02
Block errors	92	Sev. errored secs	0
Bit count	97663056484	% Sev. errored secs	00.00
Block count	47710335	Degraded minutes	0
Out of sync	0	% Degraded mins	00.00
Out of sync secs	0	Unavailable secs	0
Bit error rate	3.11e-09	% Unavailable secs	00.00
BPS clock	384000	Available seconds	254194
		% Available secs	100.00

Remarks: As can be seen, there are some bit errors on the internal connection. They are caused by the Controlware TAXI inverse multiplexers and/or the Siemens HICOM 340 PABX. It was expected that this connection would be error free.

A.2.2 Local ESTEC connection via Leiden exchange (see figure A.2).

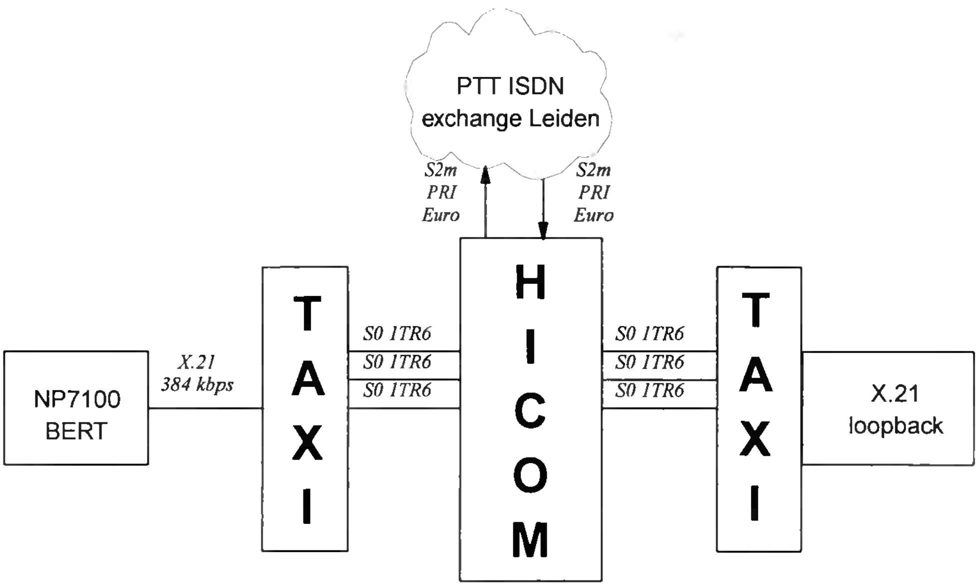

Figure A.2 *Local ESTEC connection via external call to Leiden*

Availability of ISDN lines

No.	Availability tests	Nr. of B-channels	Nr. of attempts	Start date & time	Nr. of successes	% of successes
02.01	ISDN call made	6	12	22-12-93 10:00	12*6=72	100
02.02	ISDN call made	6	12	22-12-93 12:00	12*6=72	100
02.03	ISDN call made	6	12	22-12-93 14:00	12*6=72	100
02.04	ISDN call made	6	12	22-12-93 16:00	12*6=72	100

Functionality of Controlware TAXI inverse multiplexers

No.	TAXI inverse multiplexers functionality tests	Nr. of B-channels	Nr. of attempts	Start date & time	Nr. of successes	% of successes
02.05	ISDN call timed out	6	12	22-12-93 11:00	12*6=72	100
02.06	Call re-established after time-out	6	12	22-12-93 11:00	12*6=72	100
02.07	ISDN circuit down on failure	6	12	22-12-93 11:00	12*6=72	100
02.08	ISDN circuit recovers after failure	6	12	22-12-93 11:00	12*6=72	100

Remarks: On the ESTEC-to-ESTEC connection via an external call to Leiden all the tests were 100% successful.

Quality tests (Bit error rate tests)

Test description	ESTEC-ESTEC ISDN connection via Leiden exchange, 6 B-channels aligned with TAXI inverse multiplexers		
Nr.: T312231			
Start time	23-12-93 9:46:17	Error free seconds	7200
End time	23-12-93 11:46:24	% Error free secs	100.00
Elapsed time	02:00:00	Errored seconds	0
Bit errors	0	% Errored secs	00.00
Block errors	0	Sev. errored secs	0
Bit count	2766280595	% Sev. errored secs	00.00
Block count	84422	Degraded minutes	0
Out of sync	0	% Degraded mins	00.00
Out of sync secs	0	Unavailable secs	0
Bit error rate	0	% Unavailable secs	00.00
BPS clock	384000	Available seconds	7200
		% Available secs	100.00

Test description	ESTEC-ESTEC ISDN connection via Leiden exchange, 1 B-channel aligned with TAXI inverse multiplexers		
T401131			
Start time	12-01-94 12:03:11	Error free seconds	91294
End time	13-01-94 13:25:45	% Error free secs	99.99
Elapsed time	1 Day 01:21:38	Errored seconds	4
Bit errors	5	% Errored secs	00.01
Block errors	4	Sev. errored secs	0
Bit count	5846215421	% Sev. errored secs	00.00
Block count	178417	Degraded minutes	0
Out of sync	0	% Degraded mins	00.00
Out of sync secs	0	Unavailable secs	0
Bit error rate	8.62e-10	% Unavailable secs	00.00
BPS clock	64000	Available seconds	91298
		% Available secs	100.00

A.2.3 Test of ESTEC - ESOC connection (see figure A.3)

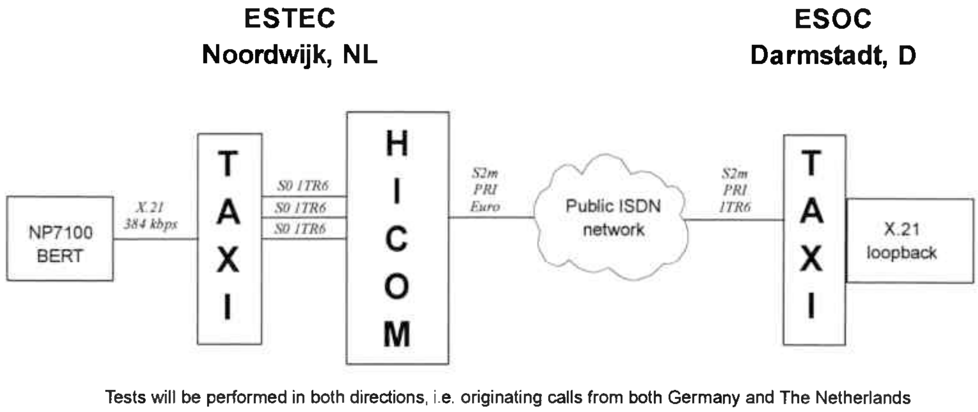

Figure A.3 *ESTEC-ESOC connection*

The functionality tests for the Controlware TAXI inverse multiplexers were performed four times on each international connection and showed no errors.

Availability of ISDN lines

No.	Availability tests	Nr. of B-channels	Nr. of attempts	Start date & time	Nr. of successes	% of successes
03.01	ISDN call made from ESTEC	6	6	05-01-94 12:00	6*6=36	100
03.02	ISDN call made from ESTEC	2	3	05-01-94 12:10	3*2=6	100

No.	Availability tests	Nr. of B-channels	Nr. of attempts	Start date & time	Nr. of successes	% of successes
03.03	ISDN call made from ESOC	6	12	05-01-94 11:50	12*6=72	100

Quality tests (Bit error rate tests)

Test description	ESTEC - ESOC Darmstadt ISDN connection,		
	6 B-channels aligned with TAXI inverse multiplexers		
Start time	16-12-93 12:15:00	Error free seconds	7548
End time	16-12-93 2:21:00	% Error free secs	99.99
Elapsed time	2:05:49	Errored seconds	1
Bit errors	1	% Errored secs	00.01
Block errors	1	Sev. errored secs	0
Bit count	2898816000	% Sev. errored secs	00.00
Block count	88467	Degraded minutes	0
Out of sync	0	% Degraded mins	00.00
Out of sync secs	0	Unavailable secs	0
Bit error rate	3.45e-10	% Unavailable secs	00.00
BPS clock	384000	Available seconds	7549
		% Available secs	100.00

A.2.4 **Test of ESTEC - DLR connection** (see figure A.4).

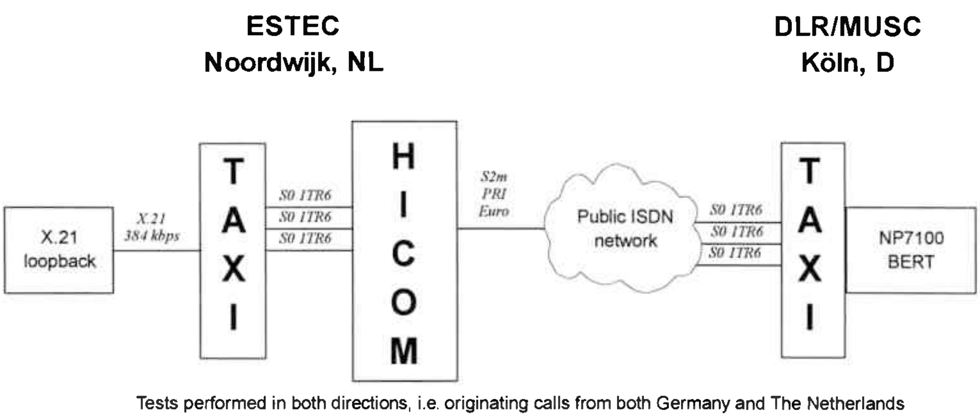

Figure A.4 *ESTEC-DLR connection*

Availability of ISDN lines

No.	Availability tests	Nr. of B-channels	Nr. of attempts	Start date & time	Nr. of successes	% of successes
04.01	ISDN call made from ESTEC	6	12	06-01-94 15:15	11*6+5=71	98.6
04.02	ISDN call made from ESTEC	6	12	06-01-94 16:20	12*6=72	100

No.	Availability tests	Nr. of B-channels	Nr. of attempts	Start date & time	Nr. of successes	% of successes
04.03	ISDN call made from DLR	6	12	05-01-94 8:30	12*6=72	100
04.04	ISDN call made from DLR	6	12	04-01-94 16:00	12*6=72	100
04.05	ISDN call made from DLR	2	5	04-01-94 16:10	5*2=10	100
04.06	ISDN call made from DLR	1	5	04-01-94 16:15	5*1=5	100
04.07	ISDN call made from DLR	6	12	05-01-94 16:25	10*6 + 2*5 = 70	97.2
04.08	ISDN call made from DLR	6	12	06-01-94 15:30	12*6=72	100
04.09	ISDN call made from DLR	6	12	06-01-94 8:45	9*6 + 3*5 = 69	95.8

Remarks: Test 04.09 shows 2 unsuccessful calls; the messages concerning the reasons were: 'Reason unspecified' and 'Reason: network congested'. This last message was only seen once during the whole test program.

Quality tests (Bit error rate tests)

Test description	DLR Köln-ESTEC ISDN connection		
Nr.: T401053	6 B-channels aligned with TAXI inverse multiplexers		
Start time	05-01-94 9:14:09	Error free seconds	4734
End time	05-01-94 10:33:08	% Error free secs	100.00
Elapsed time	01:18:54	Errored seconds	0
Bit errors	0	% Errored secs	00.00
Block errors	0	Sev. errored secs	0
Bit count	1818834854	% Sev. errored secs	00.00
Block count	55508	Degraded minutes	0
Out of sync	0	% Degraded mins	00.00
Out of sync secs	0	Unavailable secs	0
Bit error rate	00.00	% Unavailable secs	00.00
BPS clock	384000	Available seconds	4734
		% Available secs	100.00

Remarks: 100% error free.

Test description	DLR Köln-ESTEC ISDN connection,		
Nr.: T401061	1 B-channel aligned with TAXI inverse multiplexer		
Start time	05-01-94 16:43:36	Error free seconds	18406
End time	06-01-94 08:27:18	% Error free secs	99.99
Elapsed time	05:06:48	Errored seconds	1
Bit errors	8662	% Errored secs	00.01
Block errors	276	Sev. errored secs	1
Bit count	1178747132	% Sev. errored secs	00.00
Block count	18709255	Degraded minutes	0
Out of sync	1	% Degraded mins	00.00
Out of sync secs	38179	Unavailable secs	0
Bit error rate	7.35e-06	% Unavailable secs	00.00
BPS clock	64000	Available seconds	18407
		% Available secs	100.00

The call was dropped after 5 hours. The reason for this drop was a synchronisation problem with the Siemens HICOM at ESTEC. It is possible that the large burst of errors was caused because of this desynchronisation just before the call was dropped.

Test description	DLR Köln-ESTEC ISDN connection,		
Nr.: T401063	6 B-channels aligned with TAXI inverse multiplexer		
Start time	06-01-94 10:18:33	Error free seconds	14383
End time	06-01-94 14:18:33	% Error free secs	99.95
Elapsed time	03:59:49	Errored seconds	6
Bit errors	17	% Errored secs	00.05
Block errors	7	Sev. errored secs	0
Bit count	5528352927	% Sev. errored secs	00.00
Block count	87751633	Degraded minutes	0
Out of sync	0	% Degraded mins	00.00
Out of sync secs	0	Unavailable secs	0
Bit error rate	3.07e-09	% Unavailable secs	00.00
BPS clock	384000	Available seconds	14389
		% Available secs	100.00

This test shows a result comparable to the results of the local ESTEC-ESTEC connections. It is therefore assumed that the bit errors are caused by the Controlware TAXI inverse multiplexers.

Test description	DLR Köln-ESTEC ISDN connection,		
Nr.: T401071	1 B-channel aligned with TAXI inverse multiplexers		
Start time	06-01-94 16:06:28	Error free seconds	63998
End time	07-01-94 09:55:30	% Error free secs	99.99
Elapsed time	17:46:41	Errored seconds	2
Bit errors	2406	% Errored secs	00.01
Block errors	2	Sev. errored secs	1
Bit count	4098270664	% Sev. errored secs	00.00
Block count	125073	Degraded minutes	0
Out of sync	1	% Degraded mins	00.00
Out of sync secs	82	Unavailable secs	0
Bit error rate	5.88e-07	% Unavailable secs	00.00
BPS clock	64000	Available seconds	64000
		% Available secs	100.00

Call was dropped on 7th of January, at approximately 8:40. Controlware TAXI devices were programmed to re-establish connections if they were broken. The reason for the disconnection is unknown, but is likely to be due to the Controlware TAXI-equipment. The German PTT monitored the ISDN lines during this test, and they did not report outages on the ISDN lines.

To evaluate the reason for the call-drops that occurred, a similar test was performed with two Controlware CITAM X.21 terminal adapters. These units provide two X.21 ports and one BRI. Test was performed over 1 B-channel.

Test description	DLR Köln - ESTEC ISDN connection,		
T401181	1 B-channel aligned with CITAM X.21 terminal adapter		
Start time	17-01-94 12:00:44	Error free seconds	86397
End time	18-01-94 12:01:37	% Error free secs	99.99
Elapsed time	1 Days 00:00:00	Errored seconds	3
Bit errors	23	% Errored secs	00.01
Block errors	3	Sev. errored secs	0
Bit count	5532575321	% Sev. errored secs	00.00
Block count	168845	Degraded minutes	1
Out of sync	0	% Degraded mins	00.06
Out of sync secs	0	Unavailable secs	0
Bit error rate	4.16e-09	% Unavailable secs	00.00
BPS clock	64000	Available seconds	86400
		% Available secs	100.00

No problems occurred with the Controlware CITAM adapters.

A.2.5 **Test of ESOC - DLR connection** (see figure A.5).

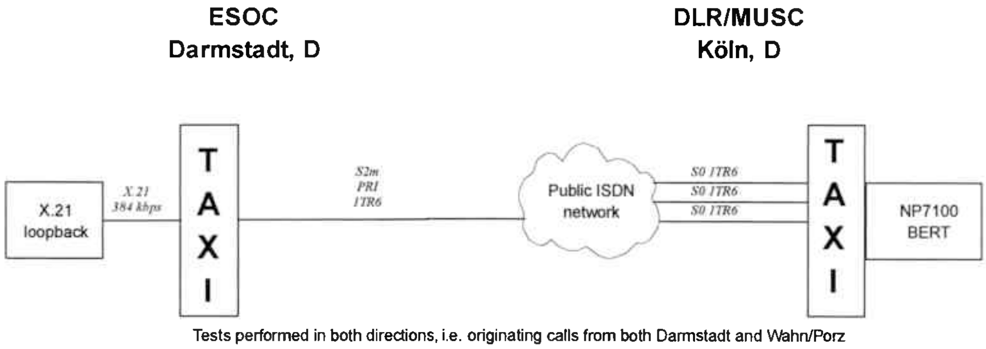

Figure A.5 *ESOC-DLR connection*

Availability of ISDN lines

No.	Availability tests	Nr. of B-channels	Nr. of attempts	Start date & time	Nr. of successes	% of successes
05.01	ISDN call made from ESOC	6	12	05-01-94 11:30	12*6=72	100

No.	Availability tests	Nr. of B-channels	Nr. of attempts	Start date & time	Nr. of successes	% of succ.
05.02	ISDN call made from DLR	6	15	05-01-94 11:00	12*6 + 2*5 + 1*4 = 86	95.6

Remark: Not all of the calls were successful; decoding of the D channel proved that the unsuccessful calls did not reach the destination address.

Quality tests (Bit error rate tests)

A bit error rate test was performed on 5 January 1994, from 12:40 to 13:40. 6 B-channels were aligned but at 13:35 the calls were disconnected because of an unknown reason. No record of error quality was recorded.

Test description	DLR Köln-ESOC Darmstadt ISDN connection		
Nr.: T401056	6 B-channels aligned with TAXI inverse multiplexers		
Start time	05-01-94 13:34:32	Error free seconds	4793
End time	05-01-94 14:54:32	% Error free secs	99.97
Elapsed time	01:19:54	Errored seconds	1
Bit errors	133	% Errored secs	00.03
Block errors	1	Sev. errored secs	0
Bit count	1841887841	% Sev. errored secs	00.00
Block count	56211	Degraded minutes	1
Out of sync	0	% Degraded mins	01.25
Out of sync secs	0	Unavailable secs	0
Bit error rate	7.24e-08	% Unavailable secs	00.00
BPS clock	384000	Available seconds	4794
		% Available secs	100.00

Remarks: The test shows that a single burst of bit errors occurred.

A.2.6 Test of ESTEC - NLR connection (see figure A.6)

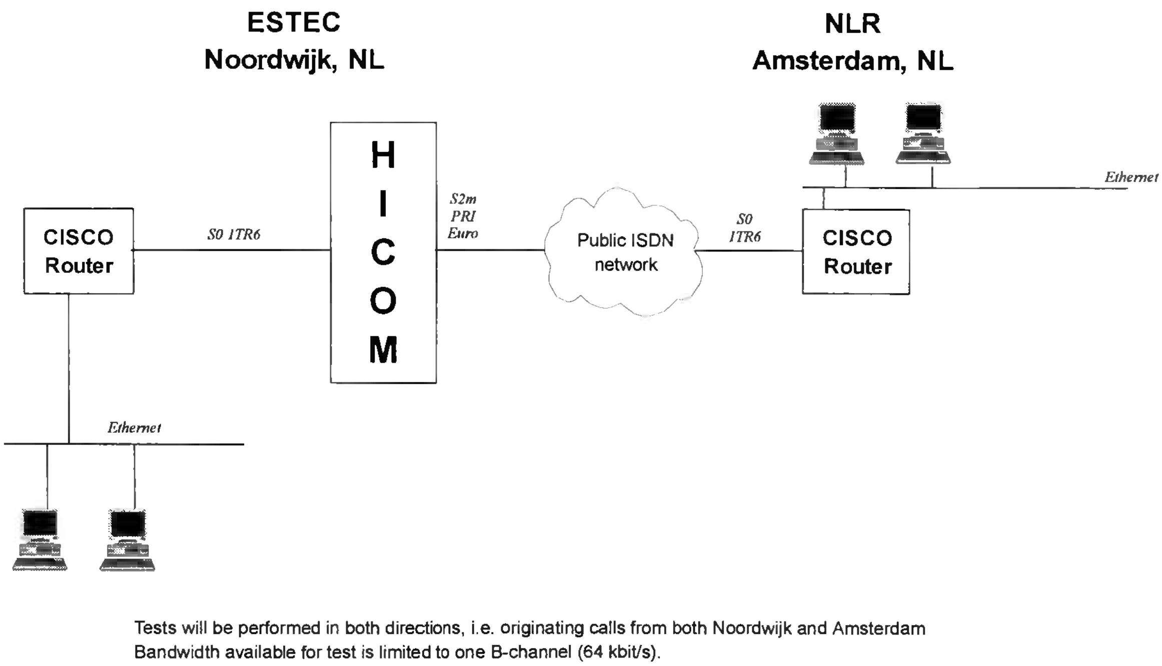

Figure A.6 *ESTEC - NLR Amsterdam connection*

Tests with Controlware TAXI inverse multiplexers were not performed between ESTEC and NLR Amsterdam, due to unavailability of suitable BRI connection. At the time of testing at NLR, only a single B-channel was being used for the test. Instead of the Controlware TAXI inverse multiplexers, two CISCO 3000 ISDN routers were used. The tests between ESTEC and NLR mainly consist of network connectivity tests and evaluate the performance of the network over ISDN. See chapter 5.x for results on these WAN-tests. The only relevant test concerning availability of ISDN bearer channels were the call attempts that were performed with the CISCO routers. These attempts were all successful.

ESTEC-NLR Amsterdam LAN-LAN interconnection

No.	General tests	Nr. of B-channels	Nr. of attempts	Start time	Nr. of successes
06.01	ISDN call made (ESTEC-NLR A'dam)	1	6	15:00	6
06.02	ISDN call made (ESTEC-NLR A'dam)	1	6	16:00	6

A.2.7 Test of ESA Paris - various sites connections

These connections were only tested during a short period of time.
Instead of Controlware TAXI inverse multiplexers, multiplexers from ASCEND were used.
Tests at Paris were performed on two differeny days.

Paris test 1. March 16, 1994.

The call setup success rate to ESOC & ESTEC was 100%. A bit error rate test between ESA Paris and ESTEC, measured for a period of 1361 seconds, resulted in 20 bit errors, which means a main BER of $3.8*10^{-8}$. No channels were lost.

Paris test 2. March 24, 1994.

The call setup success rate degraded in comparison with earlier tests. Generally, only 4 or 5 channels were connected to ESTEC and ESOC out of 6 attempts. Call attempts from Paris to ESA Toulouse were even less successful: typical only 2 to 4 channels were connected out of 6 attempts.

A.2.8 Test of ESA Toulouse - various sites connections

The test method was equal to the Paris situation.

Toulouse test . March 17, 1994.

The call setup success rate to ESTEC and ESOC was variable (3-6 channels for 18*6 channel tests), degrading to the point where no connections were obtainable, but rectified following a PBX reset to give 100% connectivity to ESOC & ESTEC. However, the ASCEND inverse multiplexer stated 'Network Problems' during the connection period, but this appears due to use of PBX provided S_0 connections, instead of France Telecom direct S_0's. A bit error rate test between Toulouse and ESOC dropped after 26 minutes (cause unknown), a second attempt gave 0 bit errors for a period of 1200 seconds. Two BERTs were performed between Toulouse and ESTEC. The first resulted in 0 errors for a 1200 seconds period; the second gave 186 errors in 1200 secs, which means a BER of $4.04*10^{-7}$. All BERTs were performed over a 384 kbit/s connection.

A.3 Discussion of test results

The Controlware TAXI inverse multiplexers showed some problems during the test cycle. A few times they needed to be reset after a test had been completed. Function was only regained after that reset.

The time that call-establishments took could not be measured in an accurate way. All the calls were established by the Controlware TAXI inverse multiplexers, and the messages generated by these devices do not contain specific information about call-setup times. After one or more calls have been answered, the Controlware TAXI unit needs an additional number of seconds (appr. 10) to align and synchronize the B-channels into an aggregated channel. However, the time from initiating the calls on a Controlware TAXI-unit until alignment of the connected B-channels was roughly measured. The time to call and align 6 B-channels using Controlware TAXI-devices was found to be between 10 seconds for local calls up to 14 to 20 seconds on international (Netherlands-Germany) calls.

On the local ESTEC connection via an internal call all the availability and Controlware TAXI functionality tests were successful. The measured bit error rate was in the range of 10^{-9}. The bit errors were caused by the combination of the Controlware TAXI inverse multiplexers and the HICOM 340 ISDN PABX.

On the local ESTEC connection via an external call to Leiden all the availability and Controlware TAXI functionality tests were also successful. The measured bit-rate was in the order of 10^{-10}. Comparing this to the results of the local ESTEC connection via internal call shows that the Noordwijk to Leiden circuit does not contribute to the number of errors.

On the international connections the Controlware TAXI functionality tests were performed four times and were all successful.

On the ESTEC - ESOC connection all the availability tests were successful. Only one bit error rate test was performed which was in the range of 10^{-10}.

On the ESTEC - DLR connection not all of the availability tests were successful; 98.84% of the call attempts were successful. On this connection once the message 'network congestion' appeared. During the bit error tests, several random call drops were experienced. These occurred between 1 and 17 hours after initial call connection time. It appears from the recorded Controlware TAXI messages that a normal disconnect from either the network or the Controlware TAXI system was occurring. Further tests using Controlware CITAM terminal adapters didn't show these call-dropouts which indicates a high likelihood for the call drops being caused by a malfunction of the Controlware TAXI inverse multiplexers.

On the ESOC - DLR connection a similar call drop problem was experienced.

On the ESTEC and ESOC to Paris and Toulouse connections several problems were experienced. Since the connectivity problem was solved in Paris by resetting the PBX, it is assumed that some of the problems were caused by equipment rather than by the network. The French D channel protocol implementation -first they used VN2 and later upgraded it to VN3- seems to cause problems for equipment manufacturers, since an X.21 TA was able to receive calls at France from the Netherlands, but was not able to call itself from France.

Appendix B. Overview of network protocols

B.1 DARPA Internet

Most of the information in this paragraph is extracted from [16].

B.1.1 DARPA Preface

The U.S. Defense Advanced Research Projects Agency's (DARPA) early research of computer interconnection led to the development of the Internet suite of protocols. This suite of protocols consists of a set of network standards that specify the details of how computers communicate, and conventions for interconnecting networks and routing traffic. The Internet suite of protocols is commonly referred to as TCP/IP after its two main standards.

The DARPA protocol stack contains four protocols: Internet Protocol (IP), Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and an Interior Gateway Protocol called the Routing Information Protocol (RIP).

B.1.2 Internet Protocol Operation

IP (Internet Protocol) is located at the Network Layer (3) of the OSI model. It provides a specification for packet formatting, packet delivery, and routing. IP also describes how packets are processed and errors handled. IP's most fundamental service is the unreliable, connectionless, best effort delivery of packets. IP service is called unreliable because delivery of packets is not guaranteed. If packets are lost, duplicated, or delivered out of order, the sender and receiver are not informed. The service is called connectionless because each packet is treated independently of the others. The service is called best effort because IP does its best to deliver packets. It only discards packets when resources are exhausted or underlying networks fail.

The basic unit of transfer in IP is the datagram. Datagrams are handled by software rather than hardware so they can be any size software designers choose. But datagrams must ultimately travel across networks in physical network frames. If a datagram is too large for a network frame to accommodate, it is fragmented, sent across the network, and reassembled at the destination. Each fragment travels independently of the other fragments and has the same format as the original datagram. Each contains a header that duplicates most of the original datagram header with some extra bits that indicate it is a fragment.

ICMP is a special case among all the protocols using IP. ICMP is an integral part of IP that handles error and control messages. Gateways and hosts use ICMP to send reports of problems about datagrams back to the original source that sent the datagram and to send new routing information. The message type determines the format of the message as well as its meaning.

0	4	8	16		31
Version	Length	Type of service	Total Length		
Ident			Flags	Fragment Offset	
Time		Proto	Header Checksum		
Source IP Address					
Destination IP Address					
Options					Padding
Data					

Table B.1 *Format of the Internet datagram.*

Version specifies the IP protocol version of the datagram. It is used to verify that the sender, the receiver, and the gateways in between them agree on the format of the datagram. If standards change, machines will reject datagrams with protocol versions different from theirs, preventing them from misinterpreting datagram contents according to an outdated format. This is a 4-bit field with a range of values from 0 - 15 decimal.

Header Length indicates the size of the datagram header measured in octets. The header has a 20-byte fixed portion and an options portion of variable length.

Type of Service is an 8-bit field, divided into 5 subfields: Precedence, Delay, Throughput, Reliability, and Unused.

Precedence	D	T	R	Unused

Precedence (Bits 0 - 2). The 3 most significant bits contain the Precedence field, which indicates the importance of a datagram. Values range from 0 - 7.

Delay (Bit 3) The delay bit requests low delay when it is set. Otherwise Normal delay is requested.

Throughput (Bit 4) The throughput bit requests high throughput when it is set. Otherwise Normal throughput is requested.

Reliability (Bit 5) The reliability bit request high reliability when it is set. Otherwise normal reliability is requested.

Unused (Bits 6- 7) The 2 least significant bits are unused.

Total Length is a 16-bit field that specifies the total datagram length, including header and data, in octets. The maximum length is 65,536 octets.

Fragment fields contains information regarding datagram fragmentation. It is subdivided to Original Datagram ID and Fragment Control.

Original datagram ID. This 16 bit field contains a unique integer that identifies the datagram. Its purpose is to allow the destination machine to collect all the fragments from a datagram. The destination uses the ID along with the source address to identify the datagram to which the fragment belongs.

Fragment Control. The 2 low order bits of the 3 bit *Flags* field control fragmentation.

May Fragment.	The first fragment control bit indicates whether or not the datagram may be fragmented.1 = Do Not Fragment. 0 = May Fragment.
Last/Only Fragment	This bit specifies whether this is the last (or only) fragment. The last/only fragment is the fragment with the highest offset. This bit is used to mark the end of the original datagram. 1 = More Fragments. 0 = Last/Only Fragment.

Fragment Offset. This 13-bit field specifies the position of a fragment in the original datagram. The offset is measured in 8-octet units, starting at offset zero.

The *time-to-live* field is 8-bits long. Time-to-live is measured in hops, or seconds. The maximum value of TTL is 255.

The 8-bit *Protocol* field indicates the transport level (Layer 4) protocol.

Header Checksum is a 16-bit field used for error detection.

Source address. This 32-bit field contains the sender's address.

Destination address. This 32-bit field contains the receiver's address.

Options, when present, are used for security, source routing, error reporting, debugging, time stamping, and other control information. This is an optional field of variable length. The length depends on the options selected.

B.1.3 Transmission Control Protocol Operation

TCP is located at the Transport Layer (Layer 4) of the OSI model. It provides reliable, end-to-end transmission that many application processes require, such as file transfer and electronic mail. TCP is connection oriented. It establishes a virtual circuit between two applications and maintains it until the transmission is complete. With TCP, applications talk to each other, rather than to intermediate machines. TCP distinguishes among multiple destinations, called ports, on a machine. Because hosts can run several application programs simultaneously, the application programs must be able to send and receive datagrams independently of the address of the host.

TCP delivers a stream of bytes, called a segment, from an application program. The destination machine receives the same sequence of bytes that the sender passes to it. Application programs see a TCP connection as a virtual circuit. To establish a connection, application programs at both ends of the connection must agree that the connection is desired. Once a connection is established, machines can begin transferring data. To make data transfer efficient, TCP software usually collects enough data from a stream to fill a large datagram before transmitting. However, protocol software can "push" data to force the transfer before the buffer is full.

TCP guarantees reliable transmission of data using positive acknowledgement with retransmission. A recipient sends an acknowledgement every time it receives a packet, but TCP allows groups of packets to be sent before an acknowledgement arrives. A sliding window keeps track of the packets that the sender has transmitted but the receiver has not yet acknowledged. After the sender receives an acknowledgement for the first outstanding packet, it "slides" the window along the outstanding sequence numbers and sends another. TCP also sets a timer when it sends a packet, and retransmits the packet if an acknowledgement has not been received when the timer expires.

TCP is a full-duplex connection allowing for simultaneous transfer of data in both directions. TCP's ability to piggyback outgoing acknowledgements with outgoing data, rather than sending separate control frames, makes efficient use of network bandwidth. A single header format is used for all TCP segments, whether the segment is carrying user data, control messages, or both. The purpose and contents of a segment are indicated by the control bits contained in the Code field. Receivers use the Window field to advertise the number of octets they are willing to accept at one time. This end-to-end flow control prevents the receiver from becoming flooded with data, and allows TCP to operate over a wide variety of physical networks.

0	4	8	16 31
Source Port			Destination Port
Sequence Number			
Acknowledgement Number			
Off	Reserv	Code	Window
Checksum			Urgent Pointer
Variable Length Options			Padding
Data			

Table B.2 *Format of the TCP segment*

The *source port* and *destination port* fields in the TCP header contain the TCP port numbers that identify the application programs at the ends of the connection.

Sequence number indicates the position in the sender's byte stream of the data in the segment.

The *acknowledgement number* field identifies the position of the highest byte that the source has received..

Data offset is a 4-bit field containing an integer that specifies the number of octets in the header. The data offset is necessary because the options field varies in length.

The field marked *reserv.* is reserved for future use.

Some segments carry only acknowledgement while some carry data. Others carry requests to establish or close a connection. TCP software uses the 6-bit field labeled *Code* to determine the purpose and contents of the segment.

Window size advertises the amount of data the sender is willing to accept, measured in data octets.

Checksum is used to verify the integrity of the TCP segment (both header and data). Software on the sending machine adds a pseudo-header to the segment to prevent misdelivery from IP, and computes the checksum over the entire result. The pseudo-header contains the source and destination IP addresses, protocol (TCP), and the total length of the segment. To compute the checksum TCP uses 1-bit arithmetic and takes the ones complement of the ones complement sum. The receiving machine performs the same computation to verify that the segment arrived intact at the correct destination (host and port).

Options The sender uses the Options field to specify a maximum segment size it is willing to receive. The Options field varies in length, depending on whether or not options are included.

The *urgent pointer* field is used to mark urgent data.

B.2 Novell IPX/SPX

B.2.1 Novell IPX/SPX preface

Novell's IPX/SPX is one of the most popular network protocol suites in use today. It is based on the Xerox Network Systems (XNS) protocols. Using XNS as a basis for local area networking protocols is logical because of its straightforward approach to addressing networking and internetworking issues.

Novell's components correspond directly to the seven layers of the OSI model. The network architectures supported at the Physical and Data Link layers, include Ethernet, IEEE 802.3, IEEE 802.5, ARCNET and a variety of others. Network layer functions are handled by the Internetwork Packet Exchange (IPX) protocol and Transport layer functions by the Sequenced Packet Exchange (SPX) protocol.

At the Network Layer, IPX provides the functions of addressing, routing, and switching packets from the source node to the destination node. It provides a connectionless communications link between nodes. Since IPX does not assure either packet delivery or packet sequence, it relies on the higher level protocols to provide this service.

At layer 4, SPX provides connection-oriented, guaranteed packet delivery services between two nodes; flow control and duplicate packet suppression.

B.2.2 Internetwork Packet Exchange (IPX) Operation

IPX (Intemetwork Packet Exchange) is located at the Network Layer (3) of the OSI model. It provides the Network Layer functions of addressing, routing, and packet switching from the source node to the destination node. IPX's most fundamental service is the unreliable connectionless, best effort delivery of packets. IPX service is called:

- unreliable because delivery of packets is not guaranteed. Only reasonable assurance of delivery is implied by IPX; it is the responsibility of the higher level protocols to ensure the delivery of the packet.

- connectionless because each packet is treated independently of the others.

- best effort because IPX does its best to deliver packets. It only discards packets when resources are exhausted or underlying networks fail.

The basic unit of transfer in IPX is the datagram. The maximum datagram packet size is 576 octets, which includes the header information (30 octets) and the data (up to 546 octets). However, in actual implementation this limit may be ignored. This is due to Novell's optimization of the NetWare drivers for each type of Network Interface Card (NIC) for higher throughput.

The format of the IPX datagram is given in table B.3..

0 7 15

Checksum	
Packet Length	
Transport Control	Packet Type
Destination Network	
Destination Node	
Destination Socket	
Source Network	
Source Node	
Source Socket	
Data	

Table B.3 *Format of the IPX datagram*

The *checksum* field is always set to a special hexadecimal value, FFFF, by IPX to indicate that the checksum is not calculated. This field is included in the datagram header for conformance with the original XNS header, but is not necessary since the LAN frame provides a cyclic redundancy check (CRC) at the Data Link Layer.

Length. An IPX datagram is always an integral number of 16-bit words. If the data is an odd number of octets, the trash octet is added to make the length a multiple of 1-bit words. The allowable length for the datagram, including both the header and the data, is between 30 and 576 octets. However, if the NetWare driver for the node's NIC has been optimized by Novell, this limit may be ignored.

The transport control field, used only by internetwork bridges and gateways, indicates the hop count for packets routed between bridged networks. It is set to 0 by IPX prior to sending the packet. The first four bits are reserved and set to 0; the last four bits are used for the hop count. The hop count is incremented by each router the packet passes through on its way to the destination node. A packet is dropped by an internetwork router on the sixteenth hop.

The *packet type* field is a one-octet field that identifies the type of service offered or desired by the packet.

The *destination address* is composed of the three subsections described below:

- *Network.* This four-octet hexadecimal address uniquely identifies the packet's destination network. This field is used by internetwork routers for routing packets between networks. The network address is displayed in hexadecimal, networkbyte order.

- *Node.* This six-octet field specifies the destination node's physical hardware address. Destination node addresses for IEEE 802.3 and 802.5 protocols use all six octets. However, other protocols, such as ARCNET, require fewer octets. If the node address is fewer than six octets, the address is placed in the least significant bit positions of the node address field and the most significant bits are filled with zeros. The node address is displayed in hexadecimal, network-byte order.

- *Socket.* This two-octet address uniquely identifies the destination node process. Sockets can be well known (WKS) or dynamically assigned on an as-needed basis.

The *source address* is composed of the three subsections described below:

- *Network.* This four-octet hexadecimal address uniquely identifies the packet's source network. This field is used by internetwork routers for routing packets between networks. The network address is displayed in hexadecimal, network-byte order.

- *Node.* This six-octet field specifies the source node's physical hardware address. Source node addresses for IEEE 802.3 and 802.5 protocols use all six octets. However, other protocols, such as ARCNET, require fewer octets. If the node address is fewer than six octets, the address is placed in the least significant bit positions of the node address field and the most significant bits are filled with zeros. The node address is displayed in hexadecimal, network-byte order.

- *Socket.* This two-octet address uniquely identifies the source node process. Sockets can be well known (WKS) or dynamically assigned on an as-needed basis.

B.2.3 Sequenced Packet Exchange (SPX) Operation

The Sequenced Packet Exchange (SPX) protocol is located at the Transport Layer (4) of the OSI model. It provides a connection-oriented path between two network nodes. This connection must be established prior to any data transmission. SPX provides guaranteed delivery of packets, in sequence, and without errors or duplication. It also provides flow control.

The basic unit of transfer in SPX is the datagram. The maximum datagram packet size is 546 octets, which includes the header information (12 octets) and the data (up to 534 octets).

The format of the SPX datagram is shown in table B.4.

0	7	15
Connection Control	Data Stream Type	
Source Connection ID		
Destination Connection ID		
Sequence Number		
Acknowledge Number		
Allocation Number		
Data		

Table B.4 *Format of the SPX datagram*

The *connection control* field is the first octet of the SPX header. It contains Control Flags eight, one-bit flags; the first four flags control the data flow between endpoints and the other four are undefined.

The *data stream type* field is a one-octet field that identifies the data within the packet. It is used by higher-layer protocols as a convenient way to send control inforrnation. The values are defined as follows:

Hexadecimal Value	Meaning
00-FD	Client defined
FE	End of connection
FF	End of connection acknowledgement

The *source connection id* field is a two-octet field that contains a unique connection identification number assigned by the source.

The *destination connection id* field is a two-octet field that contains a unique connection identification number assigned by the destination.

The *sequence number* field is a two-octet field that keeps count of the packets exchanged in one direction on the connection. Both the ends of the connection maintain their own sequence numbers.

The *acknowledgement number* field is a two-octet field that indicates the sequence number of the next packet the connection expects to receive.

The *allocation number* field is a two-octet field that, in conjunction with the acknowledge number, indicates the number of packets that may be outstanding. The allocation number specifies the sequence number up to and including the packet numbers from the other end which can be accepted.

Note: On some networks the actual number of data octets may be greater than the maximum of 534 octets. This is due to Novell's optimizing the NetWare drivers to maximize performance of network interface cards.

B3. HDLC

High-level Data Link Control is the ISO standard for Data Link layer operations in Wide Area Networks. It takes responsibility for managing synchronous, code-transparent, serial-by-bit information transfer between nodes that are joined by data links. HDLC manages line discipline, interprets commands, generates responses, controls data flow and manages error control and recovery functions.

In Wide Area Network topologies, data links are either point-to-point (two stations on a link) or multi-point (more than two station per link). Line discipline is simple in point-to-point environments; the two nodes establish which is to transmit, which is to receive, and that the receiver to ready to receive. Next, data is transmitted in one or more acknowledged blocks. When the transmission is complete, the transmitter-receiver relationship is terminated. In multi-point environments, access to the data link must be managed. This is accomplished by the choice of one of three response modes which establish whether the stations using the link share a peer-equal relationship, or are in a master-slave configuration:

Normal Response Mode--stations are master and slave. The primary station controls the access of the secondary stations to the transmission carrier by a scheme known as polling; each secondary node must wait until it is polled before it can transmit. This mode is usually used in multi-point configurations.

Asynchronous Balanced Mode--stations are peer-equals and either may initiate transmission without receiving permission from the other. Usually used only in point-to-point configurations.

Asynchronous Response Mode--the primary station has responsibility for line initialization, error recovery and disconnection, but secondary stations can transmit without explicit permission from the primary. This mode is rarely used.

In each of these response modes the network stations identify themselves by an address, which is one portion of the HDLC frame. Depending on the particular HDLC installation, the address may be one or two Bytes in length. Other aspects of line discipline managed at this layer are determined by whether the link is half- or full-duplex, that is, whether data transmission can occur only in one direction at a time (half-duplex) or in both directions simultaneously (full-duplex).

There are three formats of HDLC frames:

Supervisory frames which indicate the state of the transmitter-receiver. Supervisory frames include Receive Ready, Reject, Receive Not Ready and Selective Reject.

Information--frames which transmit data.

Unnumbered--control frames which are used for initialization, polling and status reporting. Some of these frames contain a variable-size Information field for the control information.

Each frame begins with a one-Byte flag, which alerts the receiver to the frame's presence. The address field (one or two Bytes) is the next frame component, followed by a control field with optional Information field (one or more Bytes). After the higher-layer-protocol data, a two-Byte Frame Check Sequence (using a cyclic redundancy checksum) allows the receiving station to verify the accuracy of bit transmission. The frame ends with another one-Byte flag.

HDLC provides end-to-end verification and error recovery by sequencing and acknowledgement procedures. Each HDLC user maintains a Send Counter and Receive Counter, called Ns and Nr, for frames carrying information. Values for both Ns and Nr are transmitted as part of Information frames; Supervisory frames transmit a value for Nr; there is no Ns or Nr field in Unnumbered frames. The HDLC user checks the Nr and Ns values in the relevant frames against its internal counter values. If the values maintain proper sequencing and none are missed, the internal counters are updated and acknowledgements are provided in the various Information and Supervisory format frames. If the Nr and/or Ns values in the frame do not remain sequenced, the receiving HDLC user returns a negative acknowledgement and request for retransmission by means of one of the Supervisory frames.

A REJ (Reject) frame identifies the first frame in the sequence not received correctly and requires the sending node to retransmit that frame and all frames sent after it. A SREJ (Selective Reject) identifies a single frame to be retransmitted; this takes advantage of the HDLC entity's ability to buffer out-of-sequence frames for future use, and reduces the number of frames being retransmitted on the network. The RNR (Receive Not Ready) frame acknowledges the last frame received, but tells the sending station to stop sending. This is usually an indication of a temporary problem.

Regular exchanges of these acknowledgements are important to network flow control, as they prevent the data link bandwidth from being blocked by large numbers of frame retransmissions. The protocol enforces a frequent acknowledgement pattern by limiting the number of frames that the sender can send prior to an acknowledgement from the receiver. The limit is established in a buffer called a window; with each frame received, the pointer within the window advances. The size of the window varies with what is called the modulo of the HDLC version in use. Modulo 8 operation allows a maximum window size of 7; frames numbered 0 through 7 may be transmitted before the sending station expects a confirmation from the receiving station. Modulo 128 allows a window of 127 frames. When the maximum window value (127) is reached, the transmitting user sends no further frames until an acknowledgement is received.

Address 1 or 2 bytes	Control 1 or 2 bytes	Optional information Var. # of bytes	Data	Frame Check Sequence 2 or 4 bytes

Table B5. *Format of the HDLC frame.*

The address field is given in hexadecimal, which the normal format for HDLC addresses.

The Control field identifies the frame as one of the three frame formats (Information, Supervisory and Unnumbered).

Two flags on each end of the HDLC packet (not shown) mark the beginning and the end of the packet.