CS4218 Project Paper

Oliver Gerler
9954465
CS4218: Telecommunication Networks Architectures
University of Limerick, Ireland

December 10, 1999

To Anarwenya

Chapter 1

Project Description

Consider the interconnection of several, mutual remote LANs via an ISDN wide area network. For simplicity, assume the LANs are all of the same type, i.e. Ethernet. Each LAN will be interfaced to the ISDN network through an appropriate Inter-Working Unit (IWU), in this case an NT2. The NT2 will control the access of the TE2s (i.e. PCs and Workstations on each LAN) to the ISDN network. The objective of this project is to analyse the functions of the NT2, to arrive at a good understanding of its role as an IWU, and to consider its protocol architecture.

Assume end user stations will run applications which use the TCP/IP protocols. The interconnection of the LANs via the NT2 must not require any modification to end user software applications. User applications will not be aware of the presence of the ISDN, they will have the illusion of a uniform IP addressing space. This means that user application messages will be packetised at the IP layer on the source user machine within a given LAN. Each IP packet (datagram) will have, inter alia, the IP address of the destination machine in the packet header. The NT2 must monitor the MAC frames on every LAN to decide which frames need to be transported over the ISDN network.

Analyse the functions of the NT2 under the following general headings:

1. Suggest how the set of NT2s may be interconnected. Read up on the use of the PPP protocol in this kind of configuration.

2. Discuss what terminal adaptation, if any, is required and how it may be implemented. Would it be feasible to individually adapt every connected user station or to have a common adaptation function integrated in the NT2?

3. Discuss possible protocol stack configurations appropriate to the NT2. Draw a diagram to clarify the configurations suggested.

4. Discuss various user requirements which may necessitate the bandwidth of either B, H_0 or H_1 channels to interconnect the NT2s and under what

circumstances permanent connections versus dial-up connections should be used. In the case of dial-up connections, analyse how the D channel signalling is triggered in response to the arrival of an IP packet. Also show how the 'bearer capability' parameter in the SETUP message is assembled.

5. Discuss the bufferring requirements, if any, of the NT2 for traffic in both directions, i.e. LAN to ISDN and vice versa.

6. Discuss address translation issues. Source and destination machines will be identified via IP addresses. However, ISDN does not work in terms of IP addresses. Analyse the passage of an IP-packet from source to destination machine and discuss the transformations which take place at every NT2.

7. If the number of interconnected LANs becomes very large, what problems does this cause?

8. Consider non-functional issues relating to the NT2 operation, e.g. performance, security, authentication, access control, etc.

Chapter 2

Project Development

2.1 Interconnection

2.1.1 Structure

The network defined in the project description will have a structure as shown in fig. 2.1. The goal is to connect the remote LANs via an ISDN structure to form a network with a uniform IP address range, which means that all hosts will have the same network address regardless of their location.

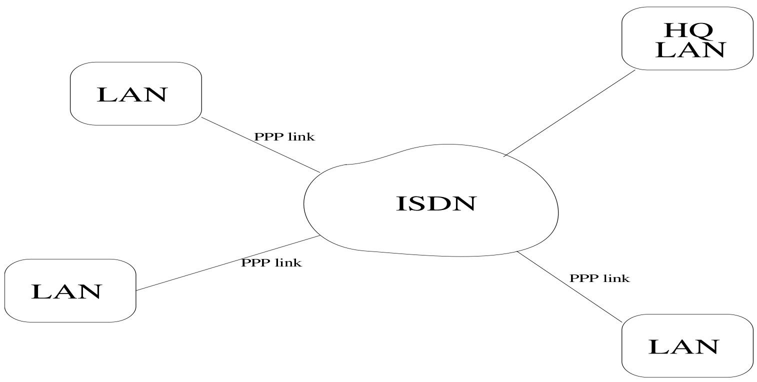

Figure 2.1: Network according to project description

One possibility to achieve this uniform address range is to divide the given address space into several subnets and set the netmasks according to the subnet. Each NT2 will then compare the destination address contained in the IP packets which it receives from its local LAN against the subnet-mask of this local LAN and decide if this particular packet has to spawn a connection to the ISDN.

Another possibility is to have the NT2 listen to all IP packets as before but then decide on each IP address separately, whether to route it into the ISDN. This method would need a table with entries of all IP addresses that can occur in all connected LANs. This table can become very big and the lookup time and requirement for storing the table would be too large for this system to be applicable.

The subnet mask has to be defined according to the number of hosts in each LAN and the available address space. For 8 subnets there are only 3 subnet bits necessary and with an address range from, for example, a class B net, 8190 hosts in each subnet are possible. This would be sufficient for many of the companies which have need for remote subnets.

Each of the LANs will be connected to the ISDN using an interface consisting of an NT2 and an NT1 as shown in fig. 2.2, where the NT2 works as the Inter Working Unit (IWU).

Figure 2.2: Interface between LAN and ISDN

According to the specifications given in ch. 1, the NT2s will be connected using a network based on the ISDN standard. These NT2s will communicate with each other by use of the Point-to-Point-Protocol (PPP).

2.1.2 PPP Configuration

To quote from the RFC[1]:

> The Point-to-Point Protocol is designed for simple links which transport packets between two peers. These links provide full-duplex simultaneous bi-directional operation, and are assumed to deliver packets in order. It is intended that PPP provide a common solution for easy connection of a wide variety of hosts, bridges and routers.

[1] See [3, ch.1].

This protocol is usually used in connecting two peers per modem or serial lines. It can be used in more bandwidth consuming lines, too. In either case it is providing connectivity for higher level protocols such as TCP/IP.

One link is for example a dial-up connection from a host to a provider where a user is connected to the network with his modem.

To handle these links, which usually are low speed ($\leq$ 1.5Mbps[2]), PPP consists of 3 main components:

- A method for encapsulating multi-protocol datagrams.
- A Link Control Protocol (LCP) for establishing, configuring, and testing the data-link connection.
- A family of Network Control Protocols (NCPs) for establishing and configuring different network-layer protocols.

The PPP encapsulation is used to disambiguate multiprotocol datagrams.

The LCP establishes and maintains each point-to-point link. This protocol sets up the link, negotiates configurations for the link's operation and finally ensures that the termination of the link is graceful.

There are several different NCPs within PPP, which handle the operation of the network protocols (like IP) over the link. For each network protocol there is an corresponding NCP.

PPP is intended to be used for dial-up links or leased lines from hosts to access nodes. It is also possible to interconnect routers with this protocol[3]. There are additional standards on how to use PPP on Frame Relay (FR) links, this method has not proven very popular, though[4]. In this project the NT2 forms the host which acts as a gateway between the local LAN and the ISDN.

The usual terminology is client for the host which starts the connection by dialling up the access node, which then performs as the server. Here, however, different settings will be necessary as there will be connections coming from the network side, too, and the access node should be able to call the host to establish a connection. As PPP is a client-server model it will be possible to have a listening server running on the LE (which forms the access node) and on the NT2 as well as a client waiting for traffic from its side from the network on the host and the access node.

Depending on the necessities of bandwidth (see forward Sec. 2.4), the link from the NT2 to the LE will differ: This includes low-bitrate links such as 14.4, 33.6, and 56 kbit/s modems, 56 and 64 kbit/s ISDN B-channels, or even sub-T1 links. On all these links PPP can be used. The configuration however will differ.

For low bitrate connection such as transferring packets containing mail a modem link or one B channel will be enough. This offers bitrates of up to 64 kbps. As we are concerned with an ISDN a digital connection using at least a B

[2]See [7, ch. 5]. The data rate is not restricted by PPP, though, and this protocol can be used in interconnecting routers, too.

[3]See [10, p.95].

[4]See footnote in [10, p.82].

channel will be assumed. In principle there is no difference in protocol if other speeds and bandwidths are used.

On ISDN links LAPD is used and so PPP is carried within these frames as can be seen in fig. 2.3. PPP introduces a 2 byte protocol field. This field determines if the data builds up to the encapsulated protocol, the LCP or the NCP. These 2 bytes are all overhead that PPP is generating. And even these two bytes can be compressed to one if both stations negotiate to do so.

Figure 2.3: Format of a PPP frame within an ISDN LAPD frame

PPP also supports various mechanisms to compress the overhead[5] that can help to improve the throughput over the link.

In overall PPP will provide a bandwidth on the link that is only slightly less than the provided bandwidth by the LAPD beneath. To further increase the possible data throughput an additional protocol[6] is defined to combine several pyhysical links from a client to a server to a logical bundle. The drawback is that this only works if all links originate from the same client and link to the same server, which can be a problem in large access nodes, since each physical link in a bundle is a full new PPP link (with its own LCPs and NCPs) and each of these links is setup by a new call to the access node which can result in a connection to a different server. To circumvent this there is a new experimental protocol to support connections to different nodes instead of simply one node at the server side[7].

The actual connection between the NT2 and the LE can be, according to the needs of bandwidth, dial-up or permanent as a switched or leased line. In both cases PPP is useable and provides the required connectivity. If permanent connections are used it is not necessary that a server and a client is up and running on both peers as there will not be any further setup of the link once it is up and running.

2.1.3 Connections of the NT2s across the ISDN

In principle there are two possibilities to interconnect two of the NT2s:

- direct
- routed

[5] Namely Protocol-Field-Compression (PFC) [3, ch. 6.5] and Address-and-Control-Field-Compression (ACFC) [3, ch. 6.6].

[6] The PPP Multilink Protocol (PPP ML) [4].

[7] The Multi-link Multi-node PPP Bundle Discovery Protocol [9].

In the direct connection, one NT2 is making a phone call to the other NT2 and connects as if it is using a common POTS network. These two NT2s can then connect per PPP and exchange their data with the help of a higher protocol which is running on that PPP link. For that each NT2 must know the numbers of all other NT2s and must make the routing decision for each outgoing packet. The connections between these two NT2s act like a circuit switched connection and need no more managment. In this configuration, however, the functions provided by ISDN and PPP are not used at all and depending on the available B channels only that many connections to other NT2s can be build at one time. This solution is not feasible if newer network technologies are available.

As there is only an ISDN as connecting network it cannot be assumed that this network knows how to route these packets. It is therefore not possible to connect only to the local next hop router (as there is none) and let the network do the routing to the other NT2. This would, however, be a usual setup and would use the features of ISDN and PPP. On the contrary, it is necessary to have one link to every other NT2 as one link cannot provide connectivity to many destinations and therefore cannot contain packets for more than one destination.

If one of these NT2s has a Primary Rate Interface (PRI) while the others only have a Basic Rate Interface (BRI), all of them can connect to the NT2 connected by the PRI, which for example can be the HQ. This NT at the PRI acts as router to all other remote LANs and has a PPP connection to all of them.

Most of these connections will be leased lines as there is traffic flow possible in both directions, which is necessary to initiate a link from both ends (example: ftp). But if one remote LAN only has client functions like users who want to fetch their mail from the server in the HQ or browsing in the net, there is no need for a permanent connection and dial-up on demand would be the solution.

2.2 Terminal Adaptation

The TCP packets from the local LAN with destinations outside of it are transported to the NT2 by IP. At that point they will be wrapped into PPP and transmitted on ISDN lines. The router in the HQ will look at the destination address of the packet and resend it over the required line to get the packet one hop nearer to its destination. At the destination NT2 the reverse operation will be applied and the TCP-Packet will be put onto the local network which operates with IP.

To be able to transport these packets across the ISDN terminal adaptation will be necessary. This adaptation mainly performs following functions[8]:

Rate Adaptation: A data stream not equal to multiples of 64 kbps has to be padded and filled into multiples of 64 kbps data streams. The data rate in Ethernet is not constant so this will have to be done by the NT2.

[8] See [12, p. 196].

Signalling Conversion: The signalling protocol of the connected device has to be converted to the ISDN signalling protocol. As Ethernet does not have such a protocol no adaptation will be necessary.

Physical Interface Conversion: The Ethernet uses coaxial cables or twisted pair connections (see forward ch. 2.4.1). These cables have to be interfaced to the twisted pair interface of ISDN.

Digitization: If there are analogue devices to be connected a conversion has to be done. Such a converter is usually called A/B-adapter. As Ethernet transmits digital signals, no digitization will be necessary.

The most problematic issue to interface Ethernet to ISDN will be the rate adaptation. In order to convert the unsynchronized and irregular occurrances of IP packets on the Ethernet to a steady 64 kbps data stream for the ISDN, buffering will be necessary. If the average data rate of the Ethernet packets to be routed into the ISDN is smaller than 64 kbps, the additional bits in the B channel data streams will remain unused and therefore set to a specified value. The allocation of the bits from the packet to the octets in the ISDN data stream will not be byte by byte, but will rather a fraction of an octet according to the required speed. For example an 8 kbps stream occupies only the first bit in every 64 kbps ISDN octet.

This adaptation will take place in the NT2s and not in every work station as there is no need to adapt the rates in the local network. Furthermore, the concept of CSMA/CD would not allow a synchronised and constant data flow from one sender.

2.3 Protocol Stack Configurations

There will be no change necessary in the configuration of the user workstations. They just have to set up their TCP stack to be able to connect to the local LAN. As the NT2 does the monitoring of the IP packets the workstations need not to be aware of the ISDN that lies between them and other machines in the same IP address space.

Information going from one application layer, like Mail or FTP data, to the application layer on a second machine will see a structure as shown in fig. 2.4.

Figure 2.4: Stack Configuration

IP isolates upper layers from specific network technologies but provides enough information for the router to specify the outgoing path for the packet to reach its destination. Higher layers like TCP or user applications are not affected by the structure of the network or the way the packet goes. Their connection is end-to-end.

So the local NT2s only have to look into the IP packet that they get from the local net and initiate the ppp device which adds the two protocol bytes to the data and sends it on their established link.

The NT2 at the destination site does the reverse operation and builds a packet that can be transmitted over the Ethernet.

2.4 Bandwidth Consideration

2.4.1 Bandwidth on Ethernet

The Ethernet which is used for the remote LANs can be build in several ways:

10Base2 describes the kind of network which is mostly used in existing nets. It is built of strings of coaxial cables which form an end-terminated bus. As the name suggests it supports data rates of (theoretically) 10 Mbps. The problem is that if the cable fails at one point the whole string is dead.

10BaseT is the most common way to build a local net nowadays. It supports easy and cheap cabling (twisted pair). The user stations are physically connected as a bus but by using a hub the net will continue to work if one cable or one user station is faulty. It also provides speeds of up to 10Mbps.

100BaseT is a newer version of 10BaseT and provides 100Mbps with the same cabling as 10BaseT if only used for short distances. For full length special cables must be used which then provide the full speed of this kind of connection.

10Base5, also called "ThickCable" or "YellowCable" is the oldest version of this kind of physical connections. It used a thicker version of coaxial cables than 10Base2. The handling of this (quite unflexible) cable and the taps to connect the user stations was quite complex and expensive. It is not very widespread nowadays.

Usual hubs for 100BaseT are backward compatible to 10BaseT (as network cards are, too) and both hubs for 100BaseT and 10BaseT mostly provide a BNC connector for connecting an existing 10Base2 network.

There are other ways to build a local LAN to obtain a higher speed but as this will not anymore use Ethernet technology they will not be discussed in this paper.

2.4.2 Bandwidth on ISDN

There are different specifications for the traffic channel type in ISDN:

Type	Data Rate	# B Channels	Channels Free
B	64 kbps	1	1
H_0	384 kbps	6	6
H_{11} (US)	1536 kbps	24	24
H_{12} (EU)	1920 kbps	30	30
E1	2048 kbps	32	30
E2	8448 kbps	132	120
E3	34.368 Mbps	537	480
E4	139.264 Mbps	2176	1920
E5	565.248 Mbps	8832	7680

Table 2.1: Traffic Channel Types

The required bandwidth in a network is heavily dependent on the purpose of the connections. As the basic B channels, each providing a connection with 64 kbps, can be bundled or multiplexed to much higher rates at higher costs, the choice on the type of connection has to be made to fit the usage.

It would be of no need to buy a 2.048 Mbps link if it is only needed to transfer some mail or files. On the other hand, one B channel is not enough to carry realtime data or video streams.

Which traffic channel type is appropriate for the purpose of the connection depends on its usage and the type of physical connection between the terminal equipment and the local exchange. If this connection is only copper then not more than 2048 kbps will be possible.

On the other hand a bandwidth of 155 Mbps or 620 Mbps (as defined by channel type E4 and E5) are physically available if fiber optics are used and can then provide access to multiple videoconferencing or video streaming[9]. As the local LANs will only be Ethernet it is of no use to have such high capacity links as they cannot be fully utilised.

2.4.3 Dial-Up vs. Permanent Connections

One B channel will be enough for users in remote LANs who only occasionally connect to other servers for fetching their mails or transferring some files. As these connections will only be established a few times in a period of time there is no need for a permanent connection and dial-up can and will be used.

If some of the users want to use the facilities provided by the HQ or other LANs such as for example some redirection of an X application window or interactive applications which require small response times the use of an H_0 will

[9] See [12, pp. 390 f.].

be appropriate. These applications mostly occupy the link for longer continuous times, so permanent connections seem to be suitable.

The links with higher data rates such as H_{11} or H_{12} can be used for videoconferencing but they will only be enough for one conference at any time, but that also depends on the desired video and audio quality. These are also connections that will be established for longer times and so seem to require a permanent connection. This, however, depends on the frequency of such sessions and it will require a thorough calculation of costs to determine which option - a dial-up or a permanent connection - will be cheaper for the end user.

2.4.4 D Channel Triggering and Bearer Capability

In the case of dial-up connections the link has to be established before it can be used. In order to do so the NT2 has to execute certain operations to setup a link.

On the receipt of an IP packet the NT2 will first look (with help of a lookup table or by masking the subnet bits) at the destination IP address to determine if this packet needs to be sent over the ISDN at all. For outbound packets it will further check if a link to the necessary destination is already up and running and will use this link if there is one.

In the case of the arrival of an outbound message without a fitting link up and running a SETUP message has to be generated and sent to the LE in order to establish a link to another NT2. The SETUP message[10] contains the mandatory fields (the Protocol Discriminator, the Call Reference Value and the Message Type, here "SETUP") and several optional Information Elements.

One of these Elements is the Bearer Capability Parameter (BCP). This Element is used in the SETUP messages to request a bearer service as specified in I.231. Unlike many other message parameters, which are passed through the net from the source to the destination, this element is used by the network itself to establish the connection and decide upon its features. The contents of the BCP are the requested bearer service and information about the terminal capabilities, like encoding or flow control mechanisms.

Up to now, the ITU-T has defined 10 circuit-mode and 3 packet-mode bearer services[11] as shown in table 2.2 and table 2.3.

More bearer services, especially at higher data rates, will be defined in the future. However, this mix of services so far defined provides the capability of meeting a wide variety of user requirements.

Table 2.4 shows the fields for transfer capabilities.

2.5 Buffering

Buffering will be necessary at the NT2s to ensure that no packets will be lost during the call and link setup and establishment phases. These phases need

[10] See fig. 8.14 in [12, p. 213].
[11] See [12, p.243].

64 kbps unrestricted
64 kbps, speech
64 kbps, 3.1kHz audio
alternate speech/unrestricted 64kbps
64kbps multiuse
2 × 64 kbps unrestricted
384 kbps unrestricted
1536 kbps unrestricted
1920 kbps unrestricted
Multiple-rate unrestricted

Table 2.2: Circuit-Mode Bearer Services (8 kHz structured)

Virtual call and permanent virtual call circuits
Connectionless
User signalling

Table 2.3: Packet-Mode Bearer Services

much less time than they did with previous old analogue telephone systems, nevertheless the speed of networks has increased, too and there will be need to store these packets.

This buffering will be done locally in the NT2 and the amount of memory required depends on the actual speed of the local network and the time it needs to set up the call. If the buffer is full packets will be discarded but that is no problem since the first packets are still in the buffers and can initiate a connection between two computers.

These buffers are also necessary to perform the rate adaptation into the ISDN as there is no constant data flow in the Ethernet.

There is no or little need for buffers in the other direction as the traffic flow into an Ethernet needs not to be synchronized. If the packet load on the lcoal Ethernet is not too high the NT2 will be able to forward the packets as they arrive. Nevertheless, for robustness and use in heavily used Ethernets buffering is at any time a good method to avoid packet loss.

2.6 Address Translation

ISDN uses a different addressing scheme as IP does, therefore some address translation has to take place. The topic of address translation has to be divided in two areas. One is the configuration where the NT2s route the packets into a ISDN that is capable of routing data itself. The other one covers direct connects of two NT2s, which is the simpler way.

In a routed environment (as explained in ch. 2.1.3), the packets are going

Intermediate rate
Network independent clock on transmission and reception
Flow control on transmission and reception
Number of data and stop bits, parity
Duplex mode
Layer 2 and Layer 3 informations

Table 2.4: Transfer Capabilities

through the Internet and therefore have to have assigned (unique) source and destination IP addresses. This can be achieved by assigning a unique IP address to every single workstation in every remote LAN or by using a technique called Network Address Translation (NAT). In a setup using NAT only the NT2s need to have assigned IP addresses and they are the sender and receiver of the packets that are crossing the Internet. At the destination NT2 the packet has to be reconverted to get its original destination address for the local LAN as it was converted by the source NT2 in the first place. In this configuration the IP addresses of the local LAN do not matter, they only have to be unique in these remote LANs.

As we here only have an ISDN that connects two NT2 on the direct way the translation becomes even simpler. All an NT2 has to do is to look up in some table which other NT2 it has to call if a certain IP address is in the destination field of the IP packet.

In both configurations the NT2 will probably call some default router (the HQ in this case) and therefore does not have to divert outgoing calls to different destinations. The IP packets will travel along a path as shown in fig. 2.4. At the remote NT2s the packets will be converted to packets on a PPP line, which was set up by LAPD and routed on a B channel. The HQ router will have a look on the destination address in the embedded IP packet and forward the packet in the direction of its destination. The NT2 at the further end of the path then will put the IP packet back on the local Ethernet.

2.7 Expansion Problems

If the network has to be expanded problems occur. It depends on the design of the existing network on how big this problems become.

If every NT2 calls every other NT2 without the help of a central router, each of the NT2 has to know the addresses of all other NT2s. If there is a change in the number of active NT2s all other NT2s have to be reconfigured. On the other hand, if there is a central router which diverts the packets in different directions, only this router has to be reconfigured and all remote LANs will have access to the new LAN without reconfiguration needs in their local environment.

For very large configurations the requirements for the central router become very high and expensive. It is in principle no problem to increase the perfomance of a router by cascading them as needed, but other, decentralizing solutions seem to be more robust. In an anlaysis of the network traffic the most used routes can be determined and direct routes can be built between these LANs without the necessity to use the router. For this, the local NT2 acts more as a router than a gateway and will divert the packets to different destination NT2s as the IP address in the packet will specify.

But this introduces restrictions. An NT2 is connected to the ISDN for example by an BRI. This interface only provides 2 B channels and so only two different destination NT2s can be connected at any time.

2.8 Operating Considerations

2.8.1 Security and Authentication

The most critical step in a security relevant environment is the authentication of the identity of the client in a secure fashion. Several technologies[12] have been developed to support this phase of creating a remote-access session, most notably TACACS[13], PAP, CHAP and RADIUS.

PPP offers two of these authentication protocols[14] for verifiying identities:

- Password Authentication Protocol (PAP).
- Challenge Handshake Authentication Protocol (CHAP).

PAP is the simpler authentication protocol. Using PAP, one system can send identification (a username or the login) and a password to the other. The peer validates the password and either accepts or rejects the connection. This approach usually is too weak as the password will be sent over the link and any third party that monitors that link can discover the password.

CHAP provides stronger security than PAP as it uses challenging to verify the correct password. For that the server is generating a challenge value which is sent to the peer that wants to login. Both are then generating a response value out of the challenge by use of the locally stored password. The client sends back its generated response value and the server compares the two responses. It will only accept the connection if both response values are the same. The algorithm to find the response value must be unidirectional to prevent reverse application of the algorithm by a third party to find the password. The most common algorithm is MD5[15].

Both RADIUS[16] and TACACS[17] provide a remote access-control mechanism. RADIUS as the newer protocol implements more functionality and security into the access node: it acts as a client/server architecture in which the

[12] See [11, pp. 167ff.]

[13] An updated variation of the TACACS protocol, called TACACS+, was introdcued by Cisco Systems Inc., a leading manufacturer of Internet routers.

[14] See [10, pp. 73f.].

[15] See [10, p. 121].

[16] See [5] and [11, pp. 168ff.].

[17] See [2].

Network Access Server (NAS) is the client and the RADIUS server can hold and deliver the authentication and profile information directly or act as a proxy to other information servers. So the authentification information can be stored centrally and has not to be spread all over the net. If the RADIUS server returns an *Access-Accept* message it can contain a user profile which can be loaded into the NAS which is most relevant for Quality of Service (QoS) for example.

Once the link is authenticated and set up there can still occur theft of data from the lines. To prevent this the communication can be encrypted using technologies like Pretty Good Privacy (PGP). This offers a unidirectional encrypting using private and public keys. PGP can also be used as a digital signature and will in Austria (as the first country in the EU) be legally acknowledged as such from January 1, 2000.

2.8.2 Performance

The speed of a network is the most important part in its conception and design. In the end a network has to work. If the user once recognizes that a network actually exists it is too slow or its configuration is not well done.

To increase the data rates several approaches can be taken. One, already mentioned at the end of ch. 2.1.2, is PPP Multilink (ML). It combines several low speed physcial PPP links to one logical link. To provide more and better services on these low speed links, which will help to improve the overall service of the network, different additional implementations[18] have been developed.

Another way to introduce a concept of priority into the network is Quality of Service (QoS)[19]. QoS basically defines different types of traffic and services which have different priorities in being transported over the network.

It is for example not important if packets containing some mail messages are delivered some seconds earlier or later. It will, however, matter for audio or video data which has to arrive at the destination not only continously and without different delays between the packets but also in the correct order.

QoS now introduces mechanisms that try to ensure that higher priority data will be transported over the net at the expense of lower priority class services. This also increases the predicability of the behaviour of the network which helps to optimize its configuration and usage.

One out of many policies that QoS introduces is the concept of a Committed Information Rate[20] (CIR). This is a data rate which is guaranteed to the user and paid by him that his connection to the network will provide at any time. The user can excess this rate but has to be aware that these additional packets will only be transported if there is no other data to be transported.

In a configuration where RADIUS is used as an access control mechanism and the user connects with the help of the PPP ML protocol an upper Port Limit can be specified according to the CIR. This Port Limit value can be part of the profile that can be loaded into the NAS at a successful login to place

[18] See [7] and [8].
[19] See [11].
[20] See [11, pp. 89f.].

an upper limit on the number of ports (or physcial links which are combined to one logical bundle) that can be configured into a PPP ML session. Within the context of QoS, this attribute can be used to define an upper bound on the aggregate access bandwidth available to a single user. As the traffic rate reaches a point of saturation of the links for a determined interval, the user equipment can decide to activate additional circuits into the PPP ML session, but has to be aware of the excess of the CIR and its consequences.

Appendix A

Acronyms

ACFC Address and Control Field Compression

ATM Asynchronous Transfer Mode

BCP Bearer Capability Parameter

BRI Basic Rate Interface

CHAP Challenge Handshake Authentication Protocol

CIR Committed Information Rate

CSMA/CD Carrier Sense Multiple Access / Collission Detect

FR Frame Relay

FTP File Transfer Protocol

IP Internet Protocol

ISDN Integrated Services Digital Network

ITU International Telecommunication Union

ITU-T ITU Telecommunication Standardization Sector

IWU Inter Working Unit

kbit/s kbps kilo bits per second

LAN Local Area Network

LAPD Link Access Protocol on the D channel

LE Local Exchange

LCP Link Control Protocol

Mbit/s Mbps Mega bits per second

MD5 Message Digest 5

ML Multilink

NAS Network Access Server

NAT Network Address Translation

NCP Network Control Protocol

NT1 Network Termination 1

NT2 Network Termination 2

PAP Password Authentication Protocol

PFC Protocol Field Compression

PGP Pretty Good Privacy

POTS Plain Old Telephone System

PPP Point-to-Point Protocol

QoS Quality of Service

PRI Primary Rate Interface

RADIUS Remote Authentication Dial In User Service

RFC Request For Comment

TACACS Terminal Access Controller Access Control System

TCP Transmission Control Protocol

TE1 Terminal Equipment type 1

TE2 Terminal Equipment type 2

Bibliography

[1] Gerler Oliver, *Lecture Notes to CS4218,* University of Limerick, 1999

[2] Finseth C., *RFC 1492: An Access Control Protocol, Sometimes Called TACACS,* Network Working Group, July 1993, [http://www.faqs.org/rfcs/rfc1492.html]

[3] Simpson W., *RFC 1661: The Point-To-Point Protocol (PPP),* Network Working Group, July 1994, [http://www.faqs.org/rfcs/rfc1661.html]

[4] Sklower K., Lloyd B., McGregor G., Carr D., Coradettim T., *RFC 1990: The PPP Multilink Protocol (MP),* Network Working Group, August 1996, [http://www.faqs.org/rfcs/rfc1990.html]

[5] Rigney C., Rubens A., Simpson W., Willens S., *RFC 2138: Remote Authentication Dial In User Service (RADIUS),* Network Working Group, April 1997, [http://www.faqs.org/rfcs/rfc2138.html]

[6] Shenker S., Partridge C., Guerin R., *RFC 2212: Specification of Guaranteed Quality of Service,* Network Working Group, September 1997, [http://www.faqs.org./rfcs/rfc2212.html]

[7] Jackowski S., Putzolu D., Crawley E., Davie B., *RFC 2688: Integrated Services Mappings for Low Speed Networks,* Network Working Group, September 1999, [http://www.faqs.org/rfcs/rfc2688.html]

[8] Bormann C., *RFC 2689: Integrated Services Mappings for Low Speed Networks, Part II,* Network Working Group, September 1999, [http://www.faqs.org/rfcs/rfc2689.html]

[9] Malkin G., *RFC 2701: Nortel Networks Multi-link Multi-node PPP Bundle Discovery Protocol,* Network Working Group, September 1999, [http://www.faqs.org/rfcs/rfc2701.html]

[10] Thomas Stephen A., *IPng and the TCP/IP Protocols,* Wiley Computer Publishing, 1996, ISDN 0-471-13088-5

[11] Ferguson Paul, Huston Geoff, *Quality of Service,* Wiley Computer Publishing, 1998, ISBN 0-471-24358-2

[12] Stallings William, *ISDN And Broadband ISDN With Frame Relay And ATM,* 4th ed., Prentice Hall, 1999, ISBN 0-13-973744-8

[13] Muller Nathan J., Davidson Robert P., *LANs to WANs,* Artech House Inc, 1990, ISBN 0-89006-410-5