Information Technology: Research and Development (1982), *1* (197–216)

THE EURONET DIANE NETWORK FOR INFORMATION RETRIEVAL

M. PURSER

Trinity College, Dublin, Ireland

(*Received 12 February 1982, revised 14 March 1982*)

ABSTRACT

The Euronet Diane network was developed to support information retrieval. This paper reviews the design of the network itself, its utilization, its performance and its tariff structure. The problems of interfacing to the network and the higher-level facilities available on the network are then discussed. Finally the future of the network is reviewed.

1. INTRODUCTION

Euronet Diane is the information retrieval service covering the countries of the European Economic Community and some others, which became operational in 1980. The service is composed of the network (Euronet), which links together user terminals and host computers, and the facilities for information retrieval (Diane), largely consisting of file management and searching software on the hosts.

The origins of Euronet date back to the early 1970s, when Directorate General XIII of the Commission of the European Communities, aided by the Committee for Information and Documentation in Science and Technology (CIDST), began a series of studies on how to improve access to scientific and technical information in Europe. The motivation lay partly in the realization that Europe, short on many other natural resources, had a major resource in knowledge and information which was not easily and universally available; and partly in response to transatlantic developments, which were seen to be a potential threat to the European information industry. A guiding spirit behind the whole Euronet project has been M. Georges Anderla, lately of D.G. XIII, and it is fitting that his name should be recorded here.

A key concept in the Euronet project has been that information should be equally available throughout the Community, with no geographical or linguistic barriers. Geographical barriers were overcome by providing a network linking together all member states, and providing essentially standard and distance-independent tariffs. This network was ordered from the PTTs who placed a contract with a Franco-British consortium, SESA-Logica, in 1977. The linguistic barriers have been tackled in various ways, for example: by issuing Euronet publicity and support documentation in all Community languages, and by providing 'neutral' standard command procedures for accessing and searching the databases on the host. However, the fact remains that the databases themselves are in their language of origin, predominantly English.

During the course of the development of Euronet, a large number of studies were undertaken on aspects as varied as tariff policy, the potential market for Euronet services, the design of communications protocols, and the provision of support services for users. Some of these studies led to the development of hardware and software products which are now in use; others provided useful policy guidance; and some studies were abortive. All this work has created Euronet Diane as it is now, with over a year's running experience behind it: a network covering the ten Community member states, with some 30 hosts, carrying over 100 databases, in eight countries, supporting some 30000 searches per month from 2000 subscribers.

This article aims to review the network aspects of Euronet Diane. The review is organized under four headings:

1. The network itself; its design, use and performance.
2. Interfacing to the network.
3. Higher level facilities for the user.
4. The future of Euronet.

2. THE NETWORK

Figure 1 illustrates the geographical and topological layout of the network. The original four nodes at Frankfurt, London, Paris and Rome have been supplemented by a fifth at Zürich. The nodes are switches, which switch traffic from a source along the correct route towards its destination. Additionally, there are multiplexors at Athens, Brussels, Copenhagen, Dublin and Luxembourg. These multiplexors do

Table 1. Approximate number of ports available on Euronet (mid-1981)

Country	*Asynchronous ports (X28)*	*Synchronous ports (X25)*
Belgium	16	4
Denmark	16	4
France	32	30
Germany	64	30
Greece	16	4
Ireland	16	4
Italy	48	30
Luxembourg	16	4
Netherlands	32	4
UK	64	28

not switch traffic; they concentrate traffic from different sources into the network onto a single line to the appropriate switch, and do the reverse on output. Three types of device are attached to the network: host computers, user terminals and special devices.

Host computers provide the information retrieval services. They are, in PTT parlance, 'Packet-Mode Data Terminal Equipment (DTE)' and are interfaced via special ports which use the X25 protocol (see below). Hosts are permanently connected to these X25 ports, which operate at 2400 and 4800 bits/s.

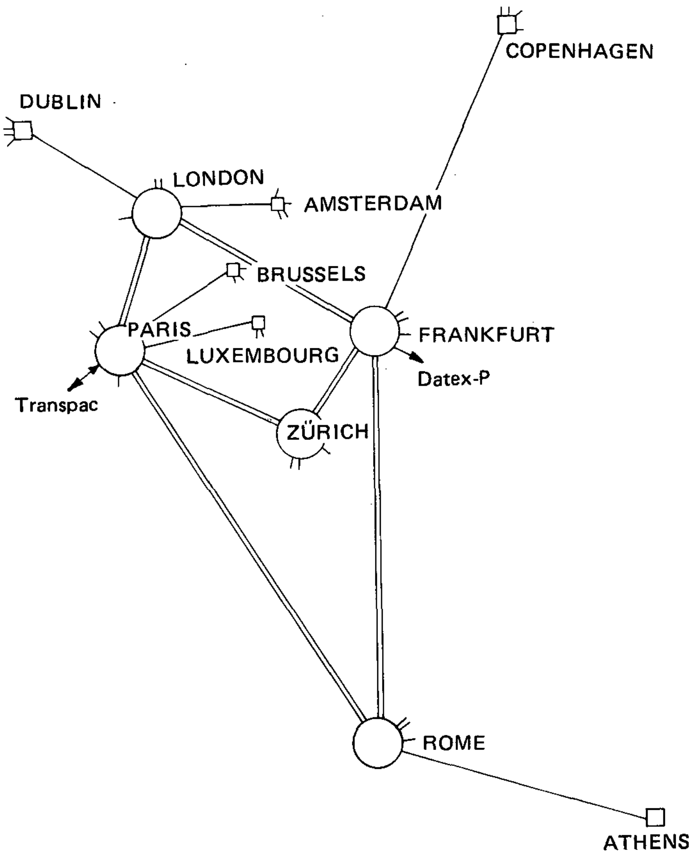

FIG. 1. Euronet topology (mid-1981); ◯ = node, □ = multiplexor

User terminals are, in general, 'Character-Mode DTEs' and are either permanently connected to the network at an X28 port, via 1200 bps lines; or can be connected and disconnected at will by the user, using a dial-up connection to an X28 port. These dial-up connections are typically via the telephone network using 300 bps modems, although in Germany and France access via other networks is also supported. An exceptional terminal is the Euronet Virtual Terminal which must be connected (permanently) to an X25 port.

Euronet has also connected to it devices (computers) which are neither hosts nor user terminals. These are essentially interfaced using X25 ports; examples being the Network Monitoring Centre in London, the NBST network in Dublin, and the Transpac gateway in Paris.

Table 1 lists the number of different types of ports provided for each node and multiplexor.

2.1 The packet-switching technique

The packet-switching technique itself is well known (Roberts, 1978; Sanders, 1978). Data are sent between subscribers' DTEs in the form of packets of characters. In Euronet, the maximum packet size is 128 octets (bytes). Packets are received at a node completely before being forwarded to another node. This means that many independent 'conversations' can share the trunk links between nodes, by interleaving packets—a form of dynamic Time Division Multiplexing (TDM). The trunk links operate at 48K bps. The technique is to be contrasted with circuit-switching, in which capacity on trunk links is permanently allocated to conversations (calls) whether or not there is currently data flowing on that call. The technique is obviously well-suited to traffic which is 'bursty', i.e., has relatively long pauses between bursts of data on each conversation. Interactive traffic, as for information retrieval, is 'bursty'.

The technique of interleaving packets also applies on links between nodes and packet-mode DTEs. A single packet-mode DTE can, in principle, handle many simultaneous conversations. The maximum number it can handle is determined when the operator of the packet-mode DTE subscribes to a connection to the network. If, for example, he subscribes to 10 'channels', then the network will be configured to support 10 simultaneous conversations on his connection, and his software must be similarly configured.

Simple character-mode DTEs, such as an asynchronous VDU or teletype, clearly cannot handle packets directly, with their headers, etc. They are supported via a PAD—a Packet Assembly/Disassembly service usually implemented by software at the node. The incoming character-by-character stream from a character-mode DTE is assembled into a packet by the PAD for forwarding through the network; the outgoing packet is disassembled into a character-by-character stream for sending to the character-mode DTE.

There are many ways in which a packet-switched network can be built. Euronet uses the 'Virtual Call' technique. A user establishes a call (e.g., between a terminal and a host) and is apparently allocated an end-to-end channel for the duration of the call. This channel is often called a 'Virtual Circuit' (VC). All subsequent traffic associated with this Virtual Call is routed along the VC that has been established. In the case of Euronet the VC is a route, determined at call establishment and fixed for the duration of the call. The nodes along the VC contain information about that VC as long as it exists. This technique means that (after call establishment) individual packets need not contain, for example, the destination address—all they need is to identify the VC to which they belong. The technique is to be contrasted with the 'datagramme' technique, in which each packet is similar to a letter in the post, carrying all the information it needs to reach its destination. Because Euronet uses fixed routing (once the call is established), and because packets are queued first-in-first-out at nodes for forwarding, packets arrive in the sequence in which they were sent, ignoring errors. This relieves the destination (node or DTE) from the need to resequence packets on reception.

The VC technique is standardized by CCITT according to the well-known X25 protocol defining the packet-mode DTE's interface to the network (CCITT Seventh Plenary Assembly). X25 defines how VCs are established and cleared, with special Call Request and Call Clear packets; how data are transferred, in Data packets; how certain error conditions are handled, with Reset packets; and some special facilities, such as the Interrupt packet.

The PAD is defined by the CCITT X3 recommendation (CCITT Seventh Plenary Assembly), which includes provision for parameters to control the PAD's behaviour. For example: when does a PAD forward an assembled packet: when the user types Carriage Return, when 128 octets have been accumulated, or when a timeout expires? Another example: does the terminal work in full duplex mode with the PAD echoing characters to it; or does the host echo characters and the PAD not; or does the terminal have local (half-duplex) echo and the PAD not? The PAD parameters determine how the PAD behaves in such cases.

Two further CCITT recommendations are associated with the PAD: X28 and X29 (Figure 2). X28 (CCITT Seventh Plenary Assembly) defines the interface between the character-mode DTE and the PAD, including the procedures for all set-up and clearing, for sending and receiving data (character set employed, etc.) and for reading and setting the PAD parameters to suit the terminal and/or the user. X29 (CCITT Seventh Plenary Assembly) essentially defines the rules whereby a remote host may control the PAD, in particular read and set its parameters. It is a weakness of the specifications that the X28 and X29 procedures for controlling PAD parameters are uncoordinated: a host can change the parameters without the user being aware of it, unless they are explicitly read by the user (and vice versa).

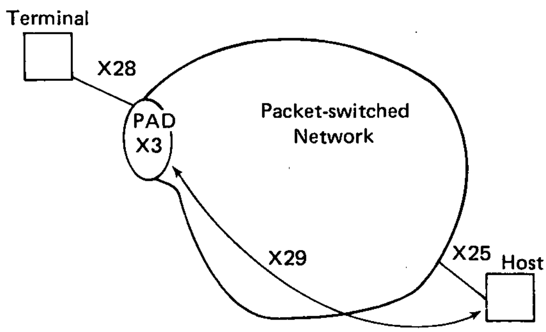

FIG. 2. X-series protocols

Two important aspects of the packet-switching technique, in particular the VC technique, should be mentioned: Flow Control, and Error Control.

Flow Control is a not very apposite name for an important requirement: that a source of data should not be allowed to flood a data sink with more than it can handle. In networks three types of flow control can be necessary:

1. Between the DTE and the network.
2. Between the nodes of the network.
3. End-to-end between communicating DTEs.

Euronet provides flow control as follows: X25 defines a 'window' mechanism whereby a DTE may not send more than a certain number of data packets (the 'window size') on a given VC before receiving a confirmation that at least some have arrived. This confirmation effectively updates the 'window edge'. Euronet employs a similar procedure between nodes. Euronet has no explicit end-to-end flow control but relies on the step-by-step flow control to ensure it. In effect, this means that if a sink DTE cannot accept data, a queue of packets builds up back across the network to the source, with the 'window size' number of packets queued at each node along the route. In theory, a subscriber can select his window size on subscribing. However, Euronet does not actually perform accordingly! It appears to assume that a receiver's window size is the same as the sender's.

Error Control is at various levels. Firstly, X25 Level 2 (effectively HDLC) contains error detection and correction-by-retransmission procedures, whereby packets are supposed to be guaranteed error-free transmission between nodes, or between packet-mode DTEs and nodes. Error-free means bit error rates of 10^{-10}. Of course errors can escape this mechanism, and could occur in the nodes themselves rather than on the lines. All packets are checked for the validity of their headers (is it a defined packet type? and if so, is it valid at this moment?), and if an error is found a Reset is generated or, in extreme cases, the call is cleared. Data packets have sequence numbers, used by the window mechanism, and a Reset effectively reinitializes the sequence numbers after scrapping all data packets in transit in the network. It is left to the end-to-end users to recover from this situation of lost packets. A particular type of error which could cause a Reset is an error in a sequence number itself. The sequence numbers also allow a receiver to check that a data packet has not been lost, or its sequence changed. In practice the use of the Reset and sequence numbers is obscure. Sequence numbers are only valid on a given 'hop', and the same packet may have a new sequence number on the next 'hop'. Resets can be generated by a DTE or within the network and their consequences are unpredictable—generally packets are lost and the network provides no help in resynchronizing the end-to-end transmission. Fortunately Resets are, in practice, very rare unless a DTE is seriously misbehaving, e.g., infringing the window mechanism.

2.2 Use of Euronet

Euronet supports some 35 host computers providing information retrieval services in all participating countries, except Greece and Ireland. Euronet has some 2000 'users', that is, issued NUIs, Network User Identifiers. (Of course many people may share an NUI, and some owners of NUIs may scarcely use Euronet.) There are over 300 asynchronous ports to support users, the majority of which are dial-in. Additionally, users of national packet-switched networks such as Transpac (France), Datex–P (Germany) can gain access to Euronet. As more national networks become available, Euronet itself will have fewer explicit users since they will be diverted by their PTT Administration to become subscribers to their national networks. Such a subscriber would access a host in his own country via the national network, and only international calls would be routed through Euronet via a 'gateway' linking a national and a Euronet node (Figure 3). As gateways are provided directly between national networks, Euronet will be further by-passed, until eventually the only traffic to be carried by it will be that to/from DTEs connected only to Euronet (Figure 4). Thus the use of Euronet is in a very fluid state, since most of the major participating countries are already operating, commissioning, or planning national networks. However, statistics from Euronet up to mid-1981 (when only France had a national network) give interesting information on user characteristics. These statistics were gathered by the Network Monitoring Centre (NMC) in London, and collated by the Euronet Launch Team in Luxembourg

(Euronet Launch Team, 1981a).

Since the NMC gathers stations for the purpose of billing users, rather than analysing their behaviour, a degree of uncertainty applies. For example: the NMC measures 'segments' transmitted rather than packets. A 'segment' is the charging unit of 64 octets and billing is per packet, reckoned in segments rounded up to the nearest whole number. Thus a data packet of 1 octet and one of 64 octets are both reckoned as a segment. Similarly a data packet of 65 or 128 octets is counted as two segments. Despite these uncertainties some broad conclusions can be drawn.

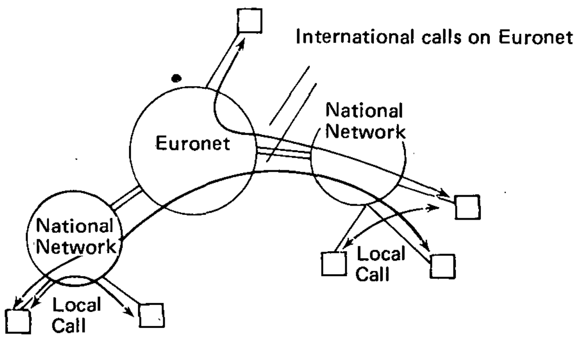

FIG. 3. National networks and Euronet

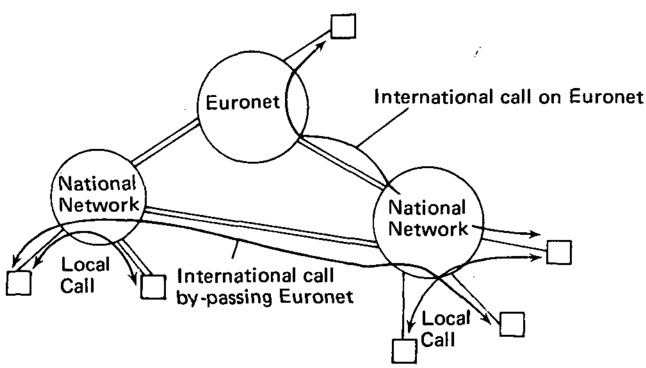

FIG. 4. Euronet by-passed

2.2.1. Total calls

About 1000 calls are made per day (autumn 1981). Obviously usage is higher during the week than at the weekend. During weekdays some 1200 to 1300 calls are made, and at weekends between 100 and 200. There are peak periods in the late morning and early afternoon.

2.2.2. Calls per NUI

Approximately 15 calls per month are made by each NUI.

2.2.3. Effective calls

Between 75 per cent and 80 per cent of calls are 'effective'. An effective call is one which results in traffic, i.e., is 'potentially billable'. Ineffective calls are those which fail to get established or used, for one reason or another. Of course, even an effective call may prove useless to the user, because it gets cleared before completion due to an error, or because he can find nothing of interest during his search on the host.

2.2.4. Domestic calls

Some 20 per cent to 25 per cent of all calls are domestic. That is, they are between DTEs in the same country. Such calls will cease to be carried by Euronet when a domestic network is available in that country.

2.2.5. Calls which pass through only one node

A further 5 per cent to 10 per cent of all calls pass through only one node. For example, an Irish user accessing a UK or a Dutch host.

2.2.6. Number of nodes per call

The average call passes through 2 nodes. The average call, which does not pass through only one node, passes through 2½ nodes. Some 40 per cent of all calls pass through 2 nodes; some 30 per cent through 3 (or conceivably more) nodes.

2.2.7. Call duration

The mean duration of an effective call is about 10 minutes, when the call is made from an asynchronous terminal. Calls made from lowspeed (300 bps) terminals are, on average, slightly longer than calls from higher speed terminals (600/1200 bps).

There is a large proportion of total calls which are very short: under 1 minute—30 per cent, under 2 minutes—40 per cent. These represent testing, demonstrations, network monitoring (see below), etc. From the point of view of the network they are valid calls. From the point of view of the user, it may be doubted if these are really information searches. If we neglect calls lasting less than one minute, the mean call duration is higher than 10 minutes; some 15 per cent to 20 per cent last longer than 20 minutes; and some 10 per cent longer than 30 minutes.

2.2.8. Loading of nodes by calls

If a 9-hour day is assumed for Euronet usage, and we assume 800 effective calls averaging 10 minutes in this period, then there are some 15 calls active at any moment. Spread over the 5 nodes, assuming a call on average uses 2 nodes, a given node handles some 6 simultaneous calls. Euronet is obviously lightly loaded, even if we allow for two or three times this figure in busy periods. The provision of, on average, 60 asynchronous ports per node appears to be excessive.

The most heavily loaded nodes are London (even discounting network monitoring traffic) and Paris, followed by Rome, Frankfurt and Zürich. (It must be remembered that at the time to which these figures apply, domestic French traffic was on Transpac, not Euronet.) London is the node which is the source and destination of most calls. Paris handles most through calls.

2.2.9. Data flows

Some 160000 to 180000 data segments per day are carried by Euronet, so per effective call there are perhaps 200 to 220 segments. The problem of relating segments to octets or packets has been mentioned. If we assume a segment is in practice about 40 octets (it could vary from 1 to 64), a call averages about 8000 to 9000 octets. The ratio of ingoing (to the host) to outgoing (from the host) octets is unknown, but a figure between 1:15 and 1:20 may be hazarded; remembering that the average call lasts 600 seconds, few users can type faster than 10 characters per second, and pauses are in the nature of on-line searching.

In terms of packets, we assume a segment is 2/3 of a packet; since a packet could contain 1 or 2 segments. Thus Euronet handles perhaps 100000 to 120000 packets per day.

2.2.10. Loading of nodes by data

Packet-switching nodes are conventionally dimensioned in terms of the number of packets they can handle per second. Figures from 20 (for small systems) to several 100 packets per second (for large systems) are normal. The designed capacity of Euronet's nodes is not known, but they are large and powerful computer systems.

Assuming a 9-hour day, Euronet is handling 3 to 4 packets per second. If, on average, each packet traverses 2 nodes, and there are 5 nodes, each node handles 1 to 2 packets per second. Allowing for busy periods, bursty traffic and uneven geographical loading it would still appear that Euronet is very lightly loaded—perhaps at most 10 packets per second per node at peak loading.

2.2.11. Loading of trunk lines by data

Trunk lines are full duplex with a capacity of 48K bps. Each call, on average, uses one line (2 nodes), and there are 6 trunks. Thus, on average, any one direction on one trunk carries 1/12 of the traffic. Total traffic (see above) is estimated at 5 to 6 segments per second, which is equivalent to 200 to 240 octets per second. To this we add 30 octets per data packet to allow for the Data Packet Header, an RR-packet, the X25 Level 2 framing, and the Level RR frames. (This is generous, and should cover other packet types as well, e.g., Call Request.) There are 3 to 4 data packets per second, so, very approximately, the traffic on Euronet averages 300 to 500 octets per second. Divided by 12 (see above) we arrive at a load figure of 25 to 42 octets per second per trunk line, or 200 to 350 bits per second on 48000 bps lines! Even allowing for busy periods, bursty traffic and geographical anomalies, it would appear that the trunk lines are grossly over-dimensioned; 9600 bps circuits would be adequate, although these would affect response times marginally.

The above analysis, although approximate, suggests that the Euronet network is very lightly used. It could probably handle 10 times the existing traffic without problems.

3. THE PERFORMANCE OF EURONET

The technical performance of Euronet has been the subject of a separate paper (Alton *et al.*, 1981). Here we only summarize the results of that study, under various headings. The study was made using a computer at the National Board for Science and Technology (NBST), Dublin, with an X25 interface to the network. Traffic from and to this computer is of course in packet form, and can be likened to traffic out of a PAD into the network and vice versa. The study does not cover the behaviour of asynchronous lines to the PAD. These lines can consist of two sections:

1. Access to the PAD or to a remote multiplexor via an external network (e.g., the public switched telephone network, PSTN).
2. Access to the PAD from the remote multiplexor (if applicable).

In Dublin, all access to Euronet is via the Irish remote multiplexor, to which the NBST is connected via a 2400 bps line leased from the PSTN. Thus all component circuits of Euronet are covered by the study, plus access via the Irish PSTN—but not on a dial-up line (Figure 5).

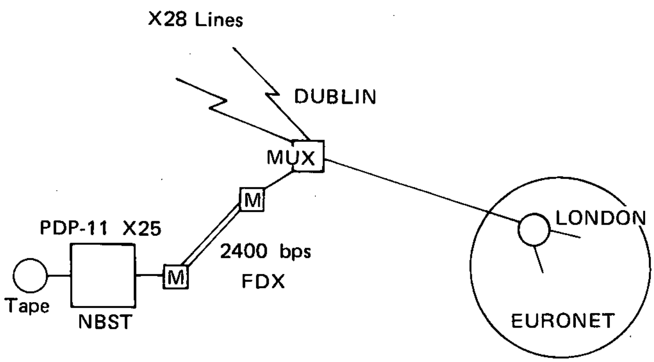

FIG. 5. Euronet monitoring—Dublin

Performance, in particular response times, is clearly a function of distance. The study measured performance by making calls automatically and randomly to the echo facilities in the Frankfurt, London, Paris and Rome nodes; and also to the echo facility at the Ispra JRC (Italy) system. The effect of topology/geography on perceived performance is mentioned where appropriate.

3.1 Call set-up delays

The mean call set-up delay was found to be 0.58 seconds. Some 8 per cent of successful call set-ups suffered long delays (up to 10 seconds) due to transmission errors, and the resultant recovering retransmissions. If these are excluded, mean call set-up delay is 0.35 seconds. If the transmission delays associated with the 2400 bps line to London, and NBST internal processing delays are excluded, this figure drops to 0.24 seconds.

More detailed analysis showed that there is a 50 msec delay per node involved, with an extra 50 msecs at the *first* node (in the NBST case, London), presumably for address validation, etc.

3.2 Trans-network delays

The NBST system actually measures round-trip delays, e.g., Dublin to Paris and back. This is equivalent to a trans-network trip from Dublin to a similar computer attached to the Rome node by a 2400 bps line. The average delay is found to be about 1.0 second. As for call set-up, a percentage (10 per cent) of these delays is long, due to transmission errors. If these are excluded, the average round-trip delay is 0.5 to 0.6 seconds.

A more detailed analysis shows that delays in the nodes themselves are very short, less than 10 msecs, provided that there is no queueing. Queueing is predominantly at the entry-node to the network, where delays can occur relatively easily. As mentioned previously, Euronet does not recognize the receiver's window size, and if the sender uses window size of 3, for example, to talk to a receiver with window size 2, he will start to experience X25 Level 3 Resets and hence further delays. Furthermore, Euronet's use of X25 is inefficient. The 'piggybacking' facility for Level 3 Acknowledgements (RRs) is not used. Queueing, errors and other inefficiencies all result in the round-trip delays experienced being well in excess of the theoretical values of 0.3 to 0.4 seconds.

Trans-network delays are fairly insensitive to the time of day, indicating that network loading is always light (see above).

3.3 Maximum throughput

If an attempt is made to send packets as fast as possible, it has been possible to attain a throughput rate of 94 per cent of the theoretical maximum, taking into account link capacities. On average, however, 75 per cent is attained, due to errors necessitating retransmissions.

An important point to make is that the performance of Euronet (and indeed any X25 network) in the presence of errors is very sensitive to the values used for the various timeouts governing retransmissions. This appears to be a forgotten area. In the CCITT Specifications most timeout values are 'for further study'. When will these studies take place?

3.4 Call failures

Calls once established are occasionally prematurely cleared by the network (LIB NC). The probability of this happening is, not surprisingly, related to the length of the calls. Using a not very large sample, it appears that perhaps 1 per cent of 1–2 minute calls are cleared in this way, while 2 per cent of 15 minute calls experience clearing by the network.

3.5 Reliability of data transfer

Excluding transmission errors which are recovered from automatically, and which only affect response, the user is interested in whether or not the data he receives are correct. Have packets been lost, duplicated or their sequence changed? Packets are lost when Resets occur, and in the NBST's experience Resets are nearly always due to the window problems mentioned above. When the sender's and receiver's window sizes match, no Resets occur. Out of some million packets sent, only one packet was duplicated. None were received out of sequence or otherwise corrupted.

3.6 Tariffs

Euronet's tariff structure was determined by the PTTs with the input from the Commission. The detailed tariffs are different in each country. The structure may be summarized as follows (figures are European Accounting Units, EAUs, translated from Irish charges and are indicative only) (Irish Government, September 18th 1978).

3.6.1. Fixed tariffs

There is a fixed charge for 'connection' to Euronet, plus a quarterly charge. Users of asynchronous terminals receive in return an NUI (Network User Identification), with which they must identify themselves to the network when they make a call. (35 EAU plus 7 EAU per quarter.) Users connected via an X25 interface, e.g., hosts, pay according to the capacity of their connection. For example, a fixed charge plus 5 EAU per channel supported per quarter, plus 200 EAU per month for a 2400 bps X25 port.

3.6.2. National access tariffs

To access Euronet most users employ the PSTN. Typically they pay normal national charges for their leased line or their dial-up call, plus rentals for modems. In some countries fixed distance-independent tariffs apply for national access. In others, national access is via a non-PSTN connection, e.g., a national data network.

3.6.3. Use of Euronet

This is the important tariff. There are two types of charges. Duration charges are typically small (e.g., 0.03 EAU/minute) and are merely a deterrent to holding calls open for ever. The volume charge is the critical one. A typical figure is 1.5 EAU per 10 segments. There are various comments to make.

1. The segment is an artificial unit. A careless user could consistently send small packets (for example if he uses a short packet-forwarding delay as defined by his PAD parameters) and be charged for 64 octets when he sends 1 or 2. A strong case could be made for reducing the size of a segment to, say, 16 or 8 octets, and its charge accordingly. This would protect users of interactive terminals from being robbed.
2. Certain forms of usage are automatically discouraged, for example, remote echo from a host, with no echo from the PAD. To avoid paying for the echoed traffic from the host a user will use echo at the PAD, and as a consequence have his host password (if applicable) displayed or printed at the terminal.
3. Euronet does not support reverse charging. Users must identify themselves with their NUI and pay their own network bill. Other packet-switched networks support *only* reverse charging for users of synchronous terminals, the user having no NUI valid for the network; he must identify himself to the host. This presents obvious problems when networks are interconnected.
4. Unlike many other packet-switched networks, Euronet has no cheap rates. Cheap rates could apply both to time of day and to volume transmitted. For example, the German Datex–P charges (FTZ–Darmstadt, March 1979) from 0.33 pfennigs/segment for small volumes during the day, down to a marginal 0.03 pfennigs/segment for traffic in excess of 400000 segments per month late at night. This possibility of cheap rates is crucial if one considers delivery of substantial volumes of information, e.g., the full text of a document, as critical to the utility of an information retrieval network. One Megaoctet currently costs 20 to 25 EAU on Euronet, whereas on Datex–P it could cost 2 to 3 EAU at the cheapest rate. The cheap rate is even more valuable when (as in Datex–P) it applies to *all* traffic for a given host for the month, in the sense that all traffic is added together, thereby producing higher volumes and thus qualifying for the volume discounts.

4. INTERFACING TO THE NETWORK

Terminals and host DTEs interface to Euronet. Their interfacing problems are discussed separately.

4.1 The user terminal

The basic 'Euronet terminal' is an asynchronous, ISO–7, device obeying the X28 procedures. It is a 'scroll-mode' device with output to the screen or printer on succeeding lines. Essentially it operates in two modes: Data or Control. In Data mode, what is typed at the keyboard is built into data packets by the PAD (in accordance with the parameter values) and forwarded into the network. Output to the terminal from the PAD is a stream of characters for printing, the PAD having the ability to insert Line Feed, Carriage Return if required in the output stream, to observe the width of the device. Control mode is entered by typing Control P. In this mode, the user can clear the call (LIB), etc., as well as manipulate the PAD parameters. Control output from the PAD simply consists of printed or displayed messages. Exit from Control mode is made automatically when the command sequence is ended.

When a user connects to Euronet initially the terminal is in Control mode. He must then enter a much-criticized, and recently revised string of 20 to 30 characters which include:

1. The destination address.
2. His own NUI.
3. Any Call User Data required by the remote host.

The Call User Data field contains characters included in the Call Request packet, used by some hosts to identify the service required.

If a user wishes to change hosts, he must clear the call (LIB) and make a new one.

For frequent users of Euronet, these procedures are cumbersome and various proposals for improving them have been put forward. A study (NETSERV2) (Commission of the EEC, 1980a) has been undertaken for the Commission to that end. At the network level there are already interfacing devices available, based on a microprocessor and memory, to allow a user to store addresses, passwords (and even search strategies) which can be transmitted directly as a result of the user depressing one or two keys. Such devices greatly simplify the user's procedures for connection, particularly if they also include a facility for automatically dialling a Euronet port via the PSTN.

Another direction taken, to increase the facilities available to the user, is to allow for a wider range of terminal types to be supported. A commonly available type of terminal that can be used with Euronet is the Videotex terminal. This is a page-mode device, but a Euronet host which is adapted to insert the equivalent of a 'Home' character after every *N* lines (*N* being the height of the screen in lines) will be able to drive it. Clearly the host should know what sort of terminal it is handling; this information can be supplied on Call set-up, e.g., in the Call User Data field. Further minor modifications to host software would enable it to make use of the bottom lines of the screen, where system messages are conventionally displayed. A further example of this sort of approach is the interfacing of the French PTT's electronic directory terminal to Euronet.

A more ambitious approach was taken by the Commission when the design and implementation of the so-called Data Entry Virtual Terminal (DEVT) was made. The DEVT is a conceptual ('virtual') device which operates in page mode (Zimmerman and Naffah, 1978). A page on the screen is composed of fields which have attributes, e.g., 'protected' (the user cannot type into those areas), 'blinking', 'suppressed' (no echo for passwords), etc. The host can control the cursor on the screen, and can define and locate the fields. The user essentially fills in the fields allocated to input, and then presses a transmit key to send his data to the host. The concept is familiar in conventional data entry (hence the name DEVT), and indeed the DEVT was conceived initially as a means of solving the problems of certain hosts proposing to connect to Euronet which were already supporting page-mode terminals.

To support the DEVT a Virtual Terminal Protocol (VTP) was designed which would enable all hosts to drive all such terminals using *standard* commands. A particular terminal manufacturer could map these commands, by firmware in the terminal, into the actual commands his device required. Similarly, a particular telecommunications access method (terminal handling software) in a host would map its output commands into VTP commands for transmission across the network.

The DEVT requires an X25 interface with as many channels as it has screens-plus-keyboards. Additionally the DEVT supports ancillary devices, e.g., printer, with appropriate flow control facilities, etc.

In practice a DEVT was actually built (real not virtual!) and the VTP implemented on a host. However little interest has been shown by the users in this product, largely because hosts have not the software to support it properly. It requires not only the VTP software but, more importantly, the actual application (database retrieval) software to make *use* of the fields and other facilities. The prototype DEVT is also expensive and is not attractive unless it can be used in a wide range of applications, or hosts.

Another, similar, development is the Remote Printing Protocol (RPP) (Commission of the EEC, 1980b) and its associated printing station. The RPP, like the VTP, is a protocol enabling a standard conceptual device to be driven by any host which has the appropriate software. Such software has been developed for IBM, Siemens, CII-HB, and Univac computers. The printing station itself has also been implemented. The objective is to permit high-speed printout, at or near the user's site, over the network. Accordingly an X25 interface is needed for the printing station. Unfortunately, as discussed above, Euronet tariffs do not encourage the use of high-volume outputs—and all the RPP headers and control strings are, of course, chargeable to traffic!

Concluding the discussion of the terminal interface: despite interesting developments, the old asynchronous teletype-compatible terminal holds its own. Fancier devices are more expensive; most hosts in practice do not support them; and the tariffs discourage the use of the high volumes, which are inevitably associated with more sophisticated facilities and higher data rates.

4.2 The host interface

The designers of X25 were not computer specialists. They gave little thought to how X25 might fit into existing hardware, and more importantly software, architectures. The Virtual Call technique, which supplies a Virtual Circuit (VC), has been seen by some mainframe manufacturers as the supplier of a *real* circuit. On this circuit they send low level line procedure frames, end-to-end, containing all their higher level protocols! The overheads involved in including, for example, SDLC within X25 data packets are horrendous. Other approaches have been to use X25 Level 3 as a kind of Transport Layer in the well known ISO OSI model, ignoring end-to-end flow control problems and using *ad hoc* message fragmentation techniques.

The problems of implementing X25 on a computer are two-fold:

1. How is the protocol itself implemented?
2. How does it interface to higher level functions whether applications, or protocols for end-to-end communications?

'Doing it yourself' to solve the first problem could involve one to three man-years of skilled programming, depending on the higher level systems with which it is supposed to interface. Buying an existing product to resolve the first problem confronts one with a second problem: is the product compatible with existing higher level functions?

X25 has been implemented on front-end processors (FEPs) for various machines, either by the manufacturer or research groups, but the majority of Euronet hosts have shirked the problem and selected the famous 'black box'. The basic idea of the black box (Commission of the EEC, 1979) is that it interfaces to the X25 network, demultiplexes the X25 channels, disassembles the packets, and presents each channel's traffic on a separate asynchronous line to the host (and vice versa). The host thinks it has a point-to-point asynchronous line to the remote terminal. The black box effectively obeys Call Request, Call Clear and similar procedures and only transfers data packets to/from the host. It is a kind of PAD.

In more detail, incoming calls to the black box are automatically accepted provided an asynchronous line to the host can be found, and the host has circuit 108 (Data Terminal Ready) asserted on that V24 interface. Provision is made for the black box to establish parameters of the remote user's PAD using X29.

Outgoing calls are not supported.

In the data transfer state flow control is required. Flow control on incoming data is done using X25 according to the buffers available in the black box. The black box does not recognize flow control on the host interface; it simply pumps characters into it. For outgoing data, the black box controls flow to it from the host using circuit 106 (Clear to Send) according to the buffers available in it for queueing data into the network. X25 provides flow control at the network interface. Call clearing from the network is confirmed by the black box, which pulls down Circuit 109 (Carrier Detect) on the host V24 interface. Call clearing by the host is done by lowering Circuit 108, and then the black box generates a Call Clear for the network. The black box handles Resets essentially without informing the host.

Various other facilities and options are provided, for example, for handling Interrupt Packets, for handling Restarts, for 'partial clearing' of VCs, for automatic speed recognition on the host interface.

The black box is cheap and convenient. However it is a confession of defeat. It needs as many asynchronous ports on the host as there are simultaneous X25 channels to be supported. It avoids all the problems of incorporating X25 into the host by providing a protocol convertor; asynchronous line to X25.

As has been discussed, various designs and projects for putting higher level end-to-end protocols on top of X25 in the host have been put forward and indeed implemented. As is well known, there is intensive work under way on standardizing transport layer, file transfer and other end-to-end procedures in accordance with ISO proposals. However it is difficult not to be sceptical about their implementation in specific systems, which are firmly rooted in manufacturers' individual architectures and which have to maintain a continuous commercial service. When the problems of implementing an *interface* protocol on large commercial systems are considered, it is hard to believe in the immediate likelihood of making end-to-end ones compatible with, for example, IBM's Systems Network Architecture.

5. HIGHER LEVEL FACILITIES FOR THE USER

The higher level facilities provided by Euronet for the user include (where implemented) such items as interfacing boxes for easier access, and the VTP and RPP protocols, already discussed. At the Application Level there is the important Common Command Language (CCL) (Negus, 1979) available, with various local differences, on several hosts. The CCL is outside the scope of this article, since it is essentially a Diane feature, rather than a network one.

On the network itself there are various other facilities available to the user. These include:

1. Testing facilities. These are really for the benefit of implementors of protocols, etc., rather than for the end user. However mention should be made of the Echo, Traffic Sink, and Traffic Generator facilities at nodes, with which the implementor of, for example, an X25 interface can converse when testing. Facilities are also available at JRC Ispra for testing, in particular, for testing VTP implementations.

   A newly available facility is the Trace facility provided via the NBST computer to authorized users. A call to this facility allows the user to ask for 'Trace' to call any specified host, and relay all packets received back to the original caller. The information relayed is supplemented with a hexadecimal dump of all data packets. This enables a user to investigate, if he suspects a host is sending him strange data such as inappropriate X29 commands to his PAD.
2. Help facilities. The Network Information Centre is a free service currently available at a London address, and shortly to be provided at Luxembourg. The user can call this service via the network and be given information about hosts available, about PTT contacts in case of problems, the latest Euronet Diane news, etc. The information provided is in all Community languages (except Greek).
3. Facilities not related to Information Retrieval. Two such facilities are available in Dublin via the NBST. One is a general Mailbox facility using the standard MS package on DEC-20 computers. The other is a Teleconferencing facility, allowing many users to converse 'simultaneously'. Essentially one user at a time has the 'turn' to type messages which are forwarded to all other logged-in users. There are algorithms for allocating the 'turn', a chairman with privileged rights (e.g., the right to interrupt) is supported, etc. Mailbox facilities are also available on other networks (Sweden) accessible via Euronet.
4. Interconnection to other networks. Besides the 'official' gateways to national networks already discussed, several research networks may be accessed via Euronet. An example is Centernet in Denmark.

These higher-level facilities are, of course, secondary to Euronet Diane's main objective: information retrieval. The user is primarily interested in accessing databases on normal Euronet hosts. It is appropriate here perhaps to put the availability of 'extras' in perspective, by quoting the results of a 'Blitz Study of Host Performance' made in February 1981 (Euronet Launch Team, 1981b). Thirty hosts were called at random in some 900 calls. Some 20 per cent to 30 per cent of all calls to hosts were unsuccessful, depending on the time of day, for one reason or another. On average, 22.4 per cent were unsuccessful. Of the failures, 1.3 per cent were due to network problems; 19 per cent were due to host problems. The balance of the failures (2.1 per cent) is due to no reaction whatever being given to calls.

The distribution of problems was:

1. Three hosts had 100 per cent availability.
2. Nine hosts had 93 per cent to 97 per cent availability.

3. Eight hosts had 81 per cent to 90 per cent availability.
4. Five had 78 per cent to 80 per cent availability.
5. The remaining hosts were available for less than 60 per cent of the time.

It is clear that the availability of *basic* service on the network is more critical than that of higher level facilities.

6. THE FUTURE OF EURONET DIANE

The future of Euronet Diane has been discussed at some length by the Commission, CIDST, etc. It is important to separate the network itself from the Diane concept.

6.1 The network (Euronet)

As has been discussed Euronet itself seems destined to see its role decreasing as national networks take its traffic. It will be used to link national networks, and then, when direct links between national networks are provided, it will merely serve to provide access from/to DTEs in countries which have not got their own national networks.

It is ironic that a network which is clearly over-dimensioned for the traffic it is carrying should see its traffic *removed* to national networks. This waste of an important capital resource is not unconnected with the fact that the Commission is committed to financing Euronet's operational deficit. Euronet's 4 million segments carried per month plus standing and other charges, probably produce less than 20000 EAU per month. This does not cover the operating costs (which include 48K bps line rentals), let alone amortize the some 7 million EAU of invested capital. Apart from some minor technical problems discussed under Network Performance (above) the network performs well, and traffic should be diverted onto it, rather than off it.

6.2 The Diane concept

The Diane concept of on-line access to information retrieval services is not, of course, dependent on the Euronet network. Any technically acceptable network will do. The future of Diane as a whole is outside the scope of this article, but certain network-related problems deserve careful attention whatever the future of the underlying network(s) may be:

1. Tariffs. The concept of distance-independent tariffs was basic to Euronet. Information should be available at the same price throughout the Community. Additionally, the need for cheap rates for bulk traffic (at night? volume discounts?) has been emphasized. Perhaps this can be achieved more easily on national networks subject to commercial pressures, than on Euronet itself.
2. The provision of extra network facilities specifically for Diane users. These include all the facilities already available at the NIC, plus possible new features such as call transfers between hosts, referral facilities in which users are advised which databases or hosts best meet their needs, automated access services to

assist in connecting to hosts, etc.

3. Better accounting information for the user. Some hosts can inform a user on-line what his search has cost him on the host. However the PTTs do not inform him on-line of his network costs. Even in the absence of reverse-charging, hosts could calculate a user's network costs and advise him on-line, which would be of great use to small users who are alarmed at the prospect of future invoices of unknown size. Many people have also advocated a central 'subscription' and billing system in which a user subscribed to, and would be billed by Diane, the individual hosts apportioning costs, etc., between them according to usage.
4. A higher degree of standardization by hosts. All Diane hosts should be capable of supporting the same range of user terminals, and using the same protocols, for example, Videotex terminals, the RPP, as well as X28 terminals. They should also all offer similar degrees of availability. At a higher level, a genuine *common* Command Language is clearly desirable.

In the absence of these features not only will the Euronet network cease to exist but so will the Diane concept. Hosts will become a collection of independent, unrelated services and inevitably the Community-wide ideal will suffer.

6.3 Document delivery

A final topic is document delivery. The Commission has undertaken a series of studies on the delivery of documents in electronic form (Little, 1980). These cover the production of electronic documents, their archiving, retrieval, transmission and reception. The aim is to provide Diane users with actual source documents on-line, or at least very quickly, rather than only with bibliographic references.

From the 'network' point of view there is an obvious need for protocols for handling this bulk traffic, and for presenting it. The concept of 'browsing', in which a user pages back and forth through a document, on-line, has also been proposed.

Probably packet-switched networks, even with cheap tariffs, are not really suitable for this bulk traffic which will include facsimile with several 100000 bits per page even with compression. Satellite transmission has been considered for the forward direction, archive-to-user. Alternatively all-digital networks could be used.

The integration of document delivery into the Diane concept is the most important step forward to be proposed.

7. CONCLUSIONS

As a network Euronet is technically successful. It performs well, nearly all transmission problems in practice being attributable to the local access (PSTN) part of the connection. The fact that adequate end-to-end protocols are not generally available is regrettable. However this only reflects the confusion reigning at that level in the whole networking field, not particularly Euronet. The Commission has made a series of positive, if not very successful initiatives, in this area. The tariff structure and the rates are acceptable, except for the absence of any form of cheap rate for bulk traffic.

Unfortunately Euronet is sadly under-utilized, perhaps as low as 10 per cent of capacity. There is little point in apportioning blame for this, but a lesson should be learned: forecasting studies should be treated with considerable caution. Few would claim that the hosts as a whole provide a consistent and coherent high-quality service. Few would deny the effort put into the project by the Commission, not least, in the last few years, by the Launch Team.

The attitude of the PTTs to Euronet has always been somewhat equivocal. On the one hand they have provided a good network. On the other, they have never actively 'sold' it to the public, and are now effectively running it down. Their financial relations with the Commission are bizarre.

To the smaller countries of the Community the Euronet project was and is of considerable importance, serving to reduce both their technical and informational isolation. It is to be hoped that the problems Euronet has encountered are not allowed to outweigh its achievements, and that the Community-wide cooperation it has engendered continues to grow.

REFERENCES

Alton, B., Patel, A., Purser, M. and Sheehan, S. (1981) The performance of a packet switched network—a study of Euronet. *Performance of Data Communication Systems and Their Applications.* (G. Pujolle, ed.) Amsterdam: North Holland.

CCITT Seventh Plenary Assembly (1980) *The Yellow Book. Volume VIII, 2.* Geneva: ITU.

Commission of the EEC (1979) *Euronet Connection Black Box.* Luxembourg: DG XIII, EEC.

Commission of the EEC (1980a) *Feasibility of a Network Facility to Simplify Information Retrieval.* Luxembourg: DG XIII, EEC.

Commission of the EEC (1980b) *The Remote Protocol.* Luxembourg: DG XIII, EEC.

Euronet Launch Team (1981a) *Euronet Statistics.* Luxembourg: DG XIII, EEC.

Euronet Launch Team (1981b) *Blitz Study on Host Performance.* Luxembourg: DG XIII, EEC.

FTZ-Darmstadt (March 1979) *Introduction of the Public Packet-Switched Datex Service.*

Irish Government (September 18th 1978) *Official Journal.*

Little, A. (1980) Electronic Document Delivery, *The Artemis Concept.* Oxford: Learned Information.

Negus, A. (1979) *Euronet Guideline: Standard Commands for Retrieval Systems.* London: INSPEC.

Roberts, L. G. (1978) The evolution of packet switching. *IEEE Proceedings.*

Sanders, R. W. (1978) Comparing network technologies. *Datamation.*

Zimmerman, N. and Naffah, J. (1978) *The Data Entry Virtual Terminal Protocol.* Paris: INRIA.