INTERNETWORK EXPERIMENT
NOTE 17

INTERNETWORK NOTEBOOK
SECTION 2.2.1.1

Datagrams as a public packet-switched data transmission service

by

Gregor V. Bochmann and Pierre Goyer
Université de Montréal

for the

Department of Communications of Canada

March 1977

This work was carried out on contract (Ref. 02SU.36100.6-0512) for the Department of Communications, Ottawa.

The content reflects the opinions of the authors, and does not necessarily imply a position of the Department of Communications.

Table of Content

Executive Summary

This study on datagrams considered as a possible data transmission service provided by public data networks has been carried out by G.V. Bochmann and P. Goyer of Université de Montréal under contract for the Department of Communications, Ottawa.

The advent of packet-switched data communication networks has brought up the question of what kind of data transmission service should be provided by public packet-switched data networks. There seems to be two basic alternatives: the datagram service and the virtual circuit service. With the CCITT Recommendation X.25, the carriers have recently chosen virtual circuits as the first service to be provided by public packet-switched data networks. Whether the datagram service will be provided in the future is an open question. The purpose of this study is to take up this question, and review and analyse the implications of a datagram service on data communication and processing systems, which implies carriers, computer and terminal manufacturers, and users.

The main objective of the study is to present a comprehensive view of the implications of a datagram service, and the report is divided into the following parts:

- An overview of the history of packet switching and the development of the datagram concept.

- An overview of communication systems architecture identifying the role a datagram service would play within the logical structure of a distributed data processing system.
- A comparison between datagrams and virtual circuits, including such points as field of application, efficiency as a transmission service, interworking of networks, flow control and charging.
- A comparative study of the current proposals for extending X.25 in order to incorporate a datagram facility and/or a "fast select" facility.

The major findings of this state-of-art study can be summarized as follows with supporting details given in the report.

1) As a transmission service, datagrams have a definite advantage over virtual circuits, being simpler to realize, more flexible and within an appropriate systems architecture can be used for obtaining reliable communications.
2) Work and experience is needed to better understand the suitability of virtual circuit and datagram services for different types of applications.
3) For the interworking of several packet-switched networks, datagrams present several advantages over virtual circuits.
4) An integrated, international billing scheme involving carriers and data processing service organizations could be implemented more easily with the virtual circuit service than with datagrams.

5) The problems of flow and congestion control in a network providing a datagram service can be resolved.

6) The different proposals for datagram facility show agreement on the basic functions to be provided and similarity in the protocols.

7) The "fast select" facility is an interesting extension of X.25, but is not an alternative to datagrams.

1. Introduction

This report covers the subject of datagrams in a public data transmission environment. As defined by one of its main proponents [1] "a datagram is a packet of information which is carried to its destination without references to any other packet, nor prior setting of a data path".

Many computer networks are based on a packet-switched datagram transmission facility and in some cases this facility is the basic service offered to the subscriber. However, the data communication carriers have decided, as expressed by th CCITT Recommendation X.25, to provide a packet-switched virtual circuit facility. The virtual circuit service is relatively similar to the service of real circuits, except that charges are calculated by the number of transmitted packets and not by connect time. Because present data processing systems usually use real circuits for their data communication needs, it seems that these systems can be more easily adapted to a virtual circuit data transmission service than to a datagram service. However, the use of a datagram service has several advantages as explained in this report.

Our main objective is to present a comprehensive view of the implications of a datagram service on data communication and processing systems, which implies carriers, computer and terminal manufacturers, and users. The report is divided into the following sections.

Section 2: An overview of the history of packet switching and the development of the datagram concept.

Section 3: An overview on communication systems architecture identifying the role a datagram service would play within the logical structure of a distributed data processing system.

Section 4: A comparison between datagrams and virtual circuits, including such points as field of application, efficiency as a transmission service, interworking of networks, flow control and charging.

Section 5: A comparative study of the current proposals for extending X.25 in order to incorporate a datagram facility and/or a "fast select" facility.

Several public packet-switched data networks are now being built providing only virtual circuits [2]. The discussion is far from closed whether virtual circuits are the ultimate transmission service for computer communications. Apart from considerations of communications systems architecture favouring datagrams, it is generally admitted that datagrams might be more efficient than virtual circuits for certain classes of applications.

The reader will find in Section 7 a list of points that require further study and that were beyond the scope of this report.

2. History of Packet Switching

Since this report deals with datagrams and virtual circuits, it is worth to recall how these two concepts have emerged in the communications community during the early seventies.

The history of packet-switching begins with the advent of "computer networks" and more specifically with the implementation of the ARPA network, started in 1967 and working since 69. In such a network, the host computers as well as terminals exchange "messages" through an underlying communication subnetwork. The subnetwork is based on packet switching, i.e. packets having a fixed maximum length are adaptively routed through a network of nodes and transmission links. Thus for sending a message to a destination host, a host would send the message to its local node which in turn would fragment it into packets and send these packets (possibly through different routes) to the destination node. The latter would reassemble the packets into a message and then transmit it to the destination host. Naturally, there is in addition a host to host protocol to control the transfer of messages.

The later CYCLADES network (1972) is particularly interesting in the fact that its communication subnetwork (CIGALE) accepts only packets and does not guarantee the delivery, nor the order in which the packets would be received. Such a subnetwork is easier to build because the data transmission service provided is less sophisticated. This approach is viable because the computers and terminal controllers connected to the subnetwork implement an end-to-end protocol which builds up a reliable data communication facility on top of the primitive transmission service provided by the subnetwork. This architecture provides the possibility that a host is connected to several different nodes of the subnetwork, which increases the communication reliability. The experience with CYCLADES led to the proposal that public data networks should provide a socalled "datagram" service similar to the service provided by the subnetwork of CYCLADES, and the adoption of this datagram service, together with the service of virtual circuits, as a user facility for public packet-switched data networks [3] .

As new applications grew (time sharing services, remote job entry and file transfer, etc.) more experience was acquired and it became evident that protocols and interfaces were to be studied more carefully in order to eliminate functional redundancy, master flow control and, if possible, define standards before too many incompatible systems would be developed. The telecommunications carriers were particularly active in the CCITT to develop an agreement on a standard network interface providing a reliable communication function based on packet switching.

In Canada the Trans-Canada Telephone System, in the preliminary proposal for Datapac [17] (1974), presented a datagram service and defined a standard network access protocol to interface with such a service. They also suggested a standard end-to-end protocol to be used by the subscriber for taking care of packet sequencing and end-to-end flow control.

In order to provide a more secure and reliable service, including packet sequencing (necessary in many applications) the carrier decided to include some functions of the end-to-end protocol into the service offered by their networks. The discussions in the CCITT resulted in Recommendation X.25 describing a data transmission service of "virtual circuits". It is interesting to note that for many existing or planned networks virtual circuits are built over an internal datagram facility which is not accessible to the subscriber [2] .

During the Special Rapporteurs' Meeting on Packet Mode Switching held by Study Group VII of CCITT in November 1976 the issue of the datagram service was discussed again. No carrier indicated immediate plans for providing a datagram service, but "it was recognized that implementation of the datagram facility in one or more experimental packet-switched networks will provide a valuable contribution to the work of the Study Group VII" [4] .

It is important to note that several carriers have remained sceptical about the future economic viability of packet-switching in general, in view of the new data switching technology of 1980's [2] .

3. Communication Systems Architecture

We give in this section an overview of the principles of communication systems architecture as a background for the discussion of the datagram service. The datagram service, by itself, is not very useful; it becomes an interesting facility in the context of an overall communication systems architecture. In fact, this architecture plays an important part within the normalization process of data communications. Therefore work on this subject has been started in the Data Communications subcommittee (SC6) of ISO/TC 97 [5] and recently, a new subcommittee on Systems Architecture (SC 16) has been formed. The following discussion is partly based on these international developments.

3.1 General overview

The architectural model which is used for describing communication systems is one involving several layers of protocols, where each layer provides a certain set of characteristic functions. Figure 3.1 shows the overall view of two communicating subscriber systems where, based on some kind of end-to-end data transmission facility, three protocol levels are distinguished:

- The transport protocol provides reliable end-to-end data transport, including such functions as recovery from transmission errors and lost packets, packet sequencing, and assembly and disassembly of user messages into packets. (Part of these functions may already be provided by the underlying transmission facility).

- The function oriented protocol provides additional standard functions, for example, depending on the application configuration, access functions for interactive terminals, or file transfer functions, etc.
- The last layer is application specific. It may consist, for example, of the message exchange between an application program and the terminal operator.

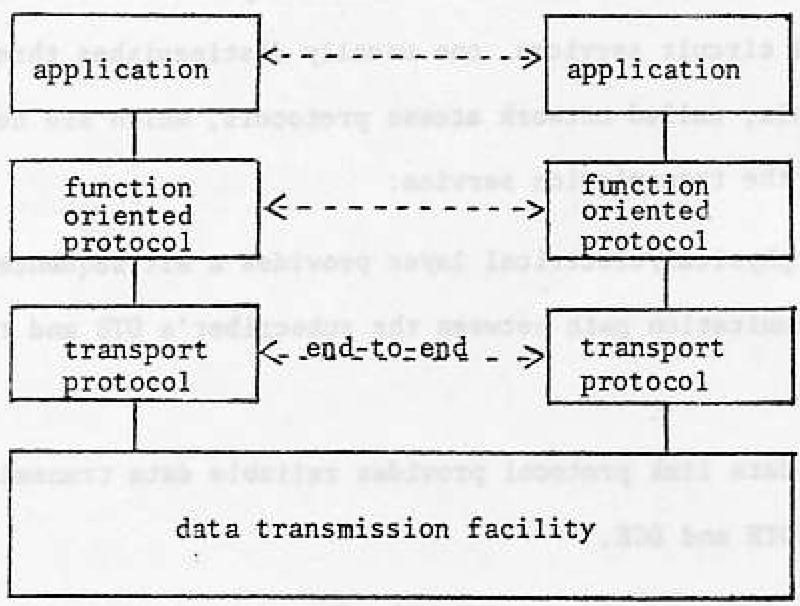

Figure 3.1

Each protocol layer, based on the functions provided by the layer below, realizes some additional communication functions which are finally used by the application.

Depending on the overall system configuration, the data transmission facility can be provided by different kinds of data transmission services; among others, we mention dedicated or switched (real) circuits, a packet-switched datagram service, a packet-switched service of virtual circuits, a satellite based packet broadcast system, etc. Figures 3.3 and 3.2 show the architectural structure of the data transmission facility in the case of (real) circuits, and packet-switched services, respectively. In the case of packet-switched datagram or virtual circuit services, one usually distinguishes three levels of protocols, called network access protocols, which are necessary for providing the transmission service:

- The physical/electrical layer provides a bit sequence oriented communication path between the subscriber's DTE and the network DCE.

- The data link protocol provides reliable data transmission between the DTE and DCE.

- The packet level protocol, via the packet-switched network (and through several networks, in the case of inter-network communications), provides some kind of end-to-end data transmission facility*.

* A problem of Recommendation X.25 is the fact that it specifies mainly conventions for the communication between the DTE and the local DCE. However, the subscriber protocols rely on the end-to-end properties of the transmission service which are not fully specified in the recommendation. The subscriber has to rely on assumptions and assurances of the carriers.

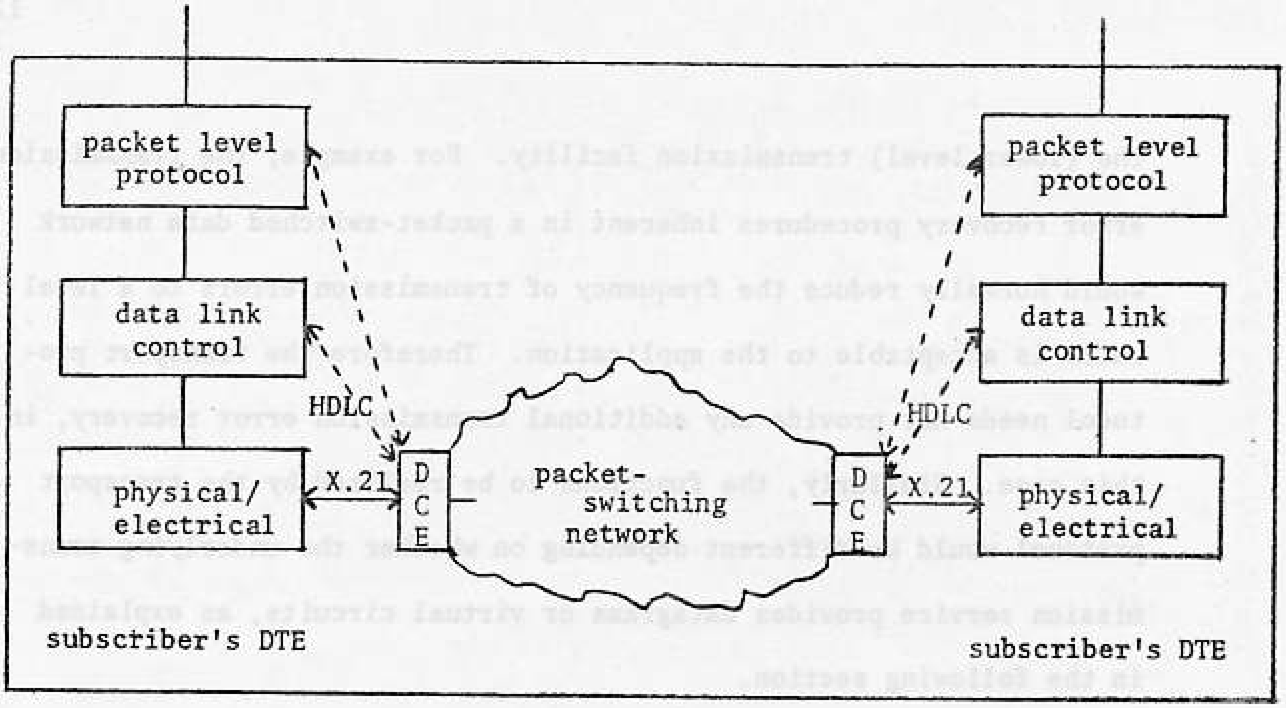

Figure 3.2

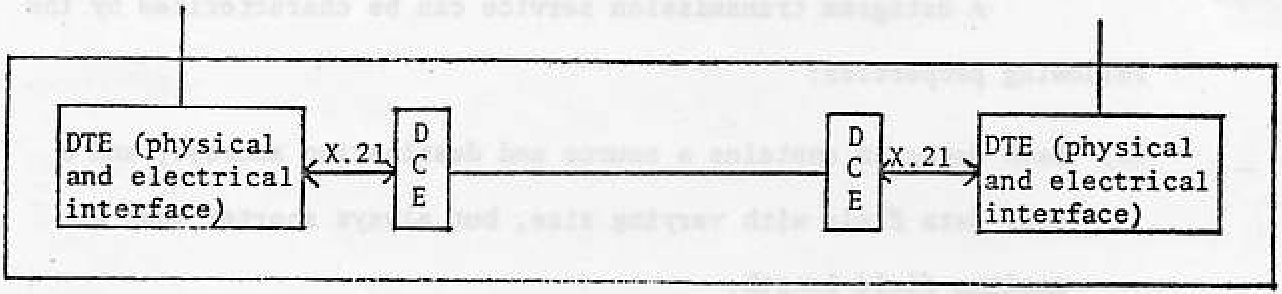

Figure 3.3

The difference between the datagram and virtual circuit services resides in the packet level protocols. They provide different kinds of end-to-end transmission facilities.

It is important to note that the details of the end-to-end transport protocol depend on the facilities and reliability needed by the application and on the communication functions already provided by the (lower level) transmission facility. For example, the transmission error recovery procedures inherent in a packet-switched data network would normally reduce the frequency of transmission errors to a level which is acceptable to the application. Therefore the transport protocol needs not provide any additional transmission error recovery, in this case. Similarly, the functions to be realized by the transport protocol would be different depending on whether the underlying transmission service provides datagrams or virtual circuits, as explained in the following section.

3.2 The datagram transmission service

A datagram transmission service can be characterized by the following properties:

(a) Each datagram contains a source and destination address, and a user data field with varying size, but always shorter than a maximum field length.

(b) The datagrams submitted to the network are transmitted to the destination DTEs independently of one another.

(c) Transmission errors are negligeable, but occasionally, datagrams may be lost by the network.

(d) The sequencing of datagrams sent consecutively by a given source DTE to the same destination DTE is not guaranteed by the network.

For comparing the datagram service with virtual circuits, we give in the following a short characterization of the virtual circuits provided by carriers according to Recommendation X.25. Before exchanging data, a virtual circuit must be established between two DTEs. For this purpose, call establishing (and clearing) packets, together with call information and call progress signals, are exchanged between the DTEs and the respective DCEs. A given DTE may establish several virtual circuits at the same time. Once a virtual circuit is established, data transmission is performed by packets. The network provides transmission error recovery and packet sequencing. Packets are not lost, unless the virtual circuit is "reset", in which case an indefinite number of data packets are lost. At rare occasions, for example in the case of congestion, the network may reset a virtual circuit.

It is clear that the datagram service is less sophisticated than the virtual circuit service. But it is also simpler, and therefore easier to provide and (possibly) less expensive. The different characteristics of the transmission services have important implications on the architectural structure of the overall system. The transport protocol, implemented by the subscriber on top of the network access transmission protocol, depending on the functions provided by the latter, may or may not perform certain communication functions, as outlined in the following.

(1) When the application needs sequential transmission of longer messages, the transport protocol has to perform message assembly and disassembly. The packet sequencing function has to be performed by the transport protocol in the case that the datagram service is used, not in the case of virtual circuits.

(2) For certain applications, such as point of sale transactions, the application only needs interactive transmission of short messages, but message loss and duplication must be avoided. For these applications a simple transport protocol performing recovery from lost packets may be used with the datagram service. A similar protocol is also needed on top of the virtual circuit service, since packets may be lost due to network circuit resets. However, the sequencing function of the transmission service is not needed.

(3) For a number of applications, the flow control mecanism provided by the X.25 virtual circuits (see Section 4.5.3) may not be sufficient. For example X.25 does not provide such functions as acknowledgments for reception, or pacing between the communicating DTEs. Providing these functions at the transport protocol level will normally imply that the basic flow control function provided by X.25 is duplicated at the transport level. In the case of an underlying datagram service, there would be no duplication since the datagram transmission service does not provide an end-to-end flow control mecanism.

(4) Sometimes it is argued that a subscriber does not need any transport protocol on top of the X.25 virtual circuit service for applications such as remote job entry, etc. However, this is only true if the grade of service provided by the data network is sufficient for the application (and we note that the grade of service decreases in the case of interworking of several networks). A critical aspect may be the frequency of data loss due to network generated circuit resets. For applications that are sensitive to data loss (and most applications are) it may be advantageous to use an end-to-end transport protocol (at the level above X.25) that recovers from lost data packets due to resets (by performing data retransmission). We note that in this situation, there are two mechanisms for packet loss recovery, one in the data network and one in the subscriber's equipment.

We can summarize this discussion by noting that the end-to-end transport protocol of the subscriber's equipment has to provide those end-to-end transport functions needed by the application that are not or insufficiently provided by the end-to-end transmission facility of the data network. In the case of a datagram transmission facility, which provides less functions than the virtual circuit facility of X.25, the transport protocol has to perform more functions. On the other hand, the X.25 facility provides some functions that are not always needed (see point 2 above), whereas certain other functions, for many applications, may not be sufficient the way they are implemented in the data network. This leads to duplication of these functions in the transport protocol.

Independently of the above discussion, general reliability principles indicate an advantage of datagrams over virtual circuits for use as a transmission service, especially when several networks are involved. In the case of virtual circuits, certain end-to-end transport functions such as sequencing of data packets, flow, error and loss control are performed in possibly several steps, as shown in figure 3.4 . The functions are performed for each step independently. The more steps are involved, the higher the probability of failure. In the case where datagrams are used, the same functions are performed end to end by the transport protocol in the subscriber's equipment. Only one step, i.e. end-to-end, is involved which implies higher reliability.

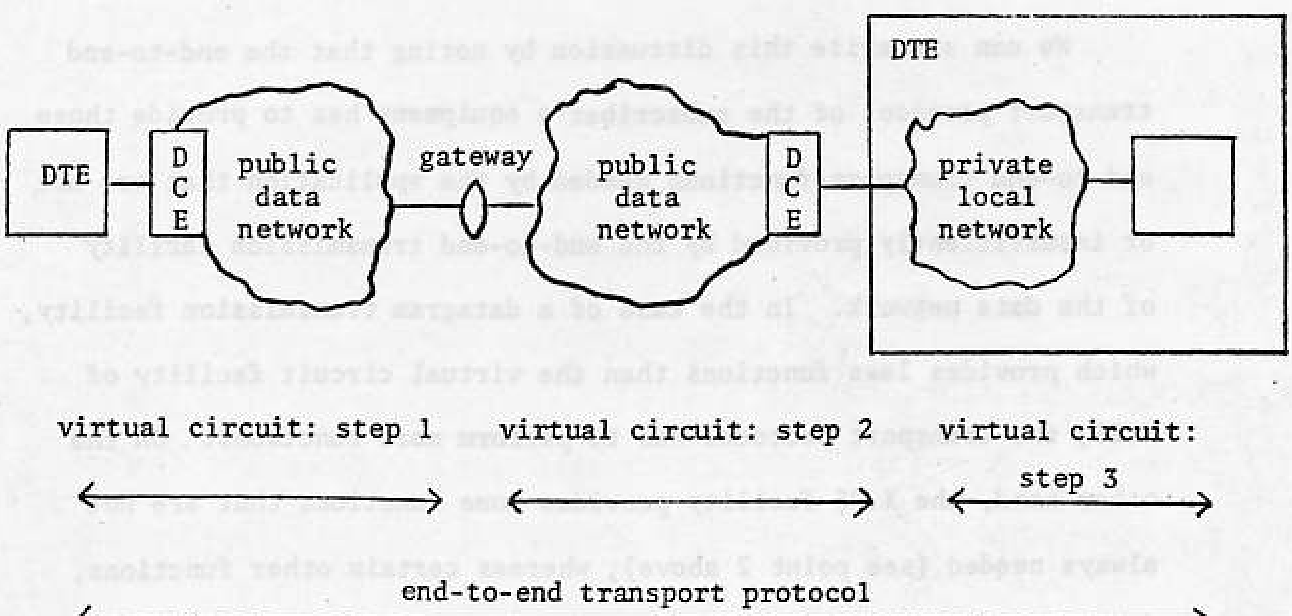

Figure 3.4

3.3 Network interfaces and terminals

Up to now, CCITT has developed two important standard interfaces for public data networks:

- Recommendation X.21 for (real) circuits, and
- Recommendation X.25 for (packet-switched) virtual circuits.

A simple packet-switched interface for single access terminals, called "frame mode interface", is being discussed. The frame mode interface will provide the transmission facility of a single virtual circuit. For providing the datagram transmission service, still another interface must be defined (different propositions for such interfaces are discussed in Section 5).

Most present-day interactive terminals are start-stop terminals. They can be used with a packet-switched data network only through an interface adapter, also called "packet assembler and disassembler" or PAD. The PAD functions can be provided by subscriber's equipment, or a network PAD service.

Intelligent terminals with a synchronous line interface could directly use the packet-switched interfaces X.25, Frame Mode and/or Datagram. Terminals that directly use the datagram interface have to implement, in their micro-computer based communication software, the network access protocol for datagrams together with appropriate end-to-end transport and terminal access protocols to be used with the other communicating DTEs. We note that such interactive terminals have already been built for the computer network CYCLADES [6] . Future widespread use of intelligent terminals will make the use of a datagram service easier.

4. Data Transmission Services: Datagrams and Virtual Circuits

We give in this section a technical comparison of datagrams and virtual circuits considered as data transmission facilities provided by public packet-switched data networks. This comparison does not rely on detailed characteristics of each of these services, but only on the general characteristics outlined in Section 3.2 . Nevertheless we refer sometimes to the CCITT Recommendation X.25 which is the standard for the virtual circuit service of public data networks.

To put this comparison into the right perspective, it is important to note that the objectives of the two services are not the same. The datagram service could be characterized as a simple and inexpensive basic packet transmission service, as expressed in [7] as follows:

"It is considered that the main difference between a datagram service and a virtual call service is that a datagram service is less complex and should only involve a minimum of network provided functions. It is considered that a datagram service should provide an extremely simple transport mechanism upon which customers can implement their own protocols at a higher level to provide any required flow control, acknowledgements, diagnostics, etc."

The objective of the carrier provided virtual circuits seems to be the provision of reliable and sequenced data transmission between two connected DTEs, plus multiplexing of such connections. However, as noted in Section 3.2, it is questionable whether the subscribers can rely on this service without using his own end-to-end transport protocol, as he would do with a datagram service.

4.1 The packet broadcasting data transmission service

The present discussion would not be complete without mentioning the packet broadcasting service which, with the development of the satellite communication technology, will become increasingly important in the near future. The packet broadcasting service would typically be provided by a satellite based communication system, in which a packet sent by a source DTE will be broadcasted to all connected DTEs, but only those DTE to whom the packet is addressed will retain the information, the others will ignore it.

It is not clear which kind of multi-DTE addressing scheme is most appropriate to be used for a broadcasting application, such as teleconferencing, over a packet broadcasting data transmission service. One possibility would be addressing by closed user group. In the case of a single destination DTE it would be sufficient to provide the DTE network address.

It is clear, however, that there is similarity between the datagram service and the packet broadcasting service. The datagram service is actually a special case of packet broadcasting where only one destination DTE is selected for each packet.

4.2 Relation between applications and transmission services

Data processing applications using data communications can be classified according to whether the expected data traffic is characterized by

(a) relatively few and long-lasting end-to-end associations with sequenced data transmission in both directions, or

(b) relatively many and short end-to-end associations, each typically requiring only a single exchange of short messages.

Another aspect is the presence (or absence) of broadcast data traffic from a given source to many destinations. Typical applications for (a) are

- file transfer and remote job entry,
- time-sharing applications, computer assisted instruction, etc.

Typical applications for (b) are

- transaction systems, such as banking systems, point-of-sale systems, reservation systems
- inquiry systems, such as credit checking
- message delivery

A typical application involving broadcast traffic is tele-conferencing, possibly computer assisted.

We consider the following different transmission services

(a) Virtual circuits,

(b) Datagrams,

(c) Packet broadcasting.

The virtual circuits are characterized by the fact that data is delivered in sequence from a source DTE and a destination DTE, except in the case of a circuit reset which could be network generated; and the circuit must be established by a call establishment procedure before it can be used for data transfer. We note that the carriers also provide permanent virtual circuits, which are permanently established. The network can take advantage of its knowledge about the association between the two communicating DTEs, and establish an efficient data transfer path between the two DTEs. (This is done in Tymnet and Transpac [2]). However, it is not clear whether this really represents an advantage, since this approach makes alternate routing (essential for reliability) more difficult. In fact, many public data networks are internally based on a kind of datagram transmission facility, in which no internal paths are established for virtual circuits. (The latter approach is taken in Datapac, EPSS, etc. [2]).

The datagram service is characterized by the absence of network known associations between source and destination DTEs, the absence of packet sequencing and occasional packet loss. It is a basic data transmission service which presents the following advantages:

(1) No call establishment procedure must be executed prior to data transfer. Datagrams are simply sent, one by one, to their appropriate destination DTE.

(2) Multi-homing procedures can be used for DTEs that have multiple connections to the network. Multiple connections of a DTE to a network are used to increase the reliability. In the case of virtual circuits, multiple connections can be established (this possibility is under study by several carriers), but only with the same network node. In the case of datagrams, multiple connections could be established with different nodes of the network. Only in the latter case does the system remain operational during a network node failure.

A superficial comparison of the different types of applications and transmission services, as outlined above, indicates that applications of type (a) may best use transmission service (a), applications of type (b) may best use transmission services (b) and (c), and applications involving broadcasting may best use transmission service (c). However, we believe the situation is not as simple, as shown by the following remarks:

(1) Many applications of type (a) need, in addition to the carrier provided virtual circuits, end-to-end transport protocols which could as well be implemented on the datagram transmission service.

(2) As noted above, many data networks do not take advantage of the virtual circuit structure for optimizing transmission efficiency.

(3) Most present distributed data processing systems are designed around real circuits, available since long time. An adaptation to the datagram service is more difficult for these systems than to the virtual circuits.

(4) The virtual circuit service could possibly be adapted to applications of type (b) (see Section 5.2, discussion of the Fast Select facility).

(5) The advent of cheap intelligent terminals may favor the use of a datagram service.

More work and experience with both, virtual circuits and datagrams, is needed for better understanding the suitability of the different data transmission services for different types of applications. This need is also recognized by the CCITT [4].

4.3 Service charges and billing

4.3.1 Accounting for billing purposes

Communication carriers have the tradition of billing the subscriber providing information about each call established during the billing period. Similar practice can be continued for the virtual circuit data transmission service, keeping a record of all virtual connections established by the subscriber's DTE.

In the case of the datagram service, this practice is not feasible because the network is not aware of virtual end-to-end associations between the DTEs connected to the network. We believe that appropriate accounting schemes can be developed that are reasonably simple and acceptable to the subscriber.

For example, accounting for each individual datagram is probably too cumbersome. However, accounting for all datagrams in each individual distance category and user facility may be feasible. Another possible accounting scheme is by periods for which the DTE is connected to the network. For host computers these periods may be typically one day or longer, whereas for terminals these periods may be shorter.

We note that in the case of virtual circuits, too, the accounting scheme by virtual connections may be found too cumbersome, and schemes like those mentioned for the datagram service may be appropriate.

4.3.2 Reversed charging

The reversed charging facility, foreseen with the virtual circuit service, is difficult to implement with the datagram service because of the following points [7] :

(1) How can the destination DTE refuse charges after it has received the whole packet?

(2) When a datagram is lost (which may happen occasionally) could the network charge the destination DTE?

(3) The destination DTE has no direct control over the number of datagrams that are sent.

4.3.3 The issue of charging for non-delivered packets

For simplifying the accounting mechanism it seems natural to charge a DTE for each datagram sent, irrespective of whether it is actually delivered or lost. We think that this is acceptable to the subscriber as long as the percentage of non-delivered packets is reasonably small. We note that the end-to-end protocol used by the communicating DTEs will usually detect all instances of datagram loss anyway, so that the subscriber can be aware of the momentaneous and average quality of the transmission service.

In the case of datagrams sent with a request for delivery confirmation (a special user facility) charging for non-delivered packets can clearly be avoided very easily.

4.3.4 Towards an integrated system of payments

In this subsection, we consider a system of payment in which the subscriber is billed by a single invoice for all services obtained through a telecommunication connection, not only for national or international telephone or data communication charges, but also for computer service or data base access charges. The different services may be provided by a variety of companies, and it may be the role of the carrier to bill the subscriber and collect the fees. The advantages of such an integrated billing scheme has been eloquently described by I. de Sola Pool [8] with special emphasis on international data communications.

We note that for such an integrated billing scheme a data network accounting scheme by virtual connections (see Section 4.3.1) has the advantage that the additional charges for a computer service or data base access can be simply associated with the virtual connection through the network. In the case when the network only provides a datagram transmission service, different procedures would have to be developed for allowing the computer service bureau or data base administration to communicate to the carrier the service charges for the subscriber.

4.4 Interworking between networks

Different approaches to interworking of data networks have been discussed [9] and many specific proposals have been made [10, 11, 12, 13, 14] of which many have actually been tried out by experiments [11, 13]. All of these proposals assume some kind of gateways which are interposed at the connection points between the networks. The different approaches to interworking can be distinguished by the architectural level at which the interworking between the networks is established. For providing interworking between datagram services of different networks it seems necessary that the interworking is based on the exchange of individual datagrams [9]. Interworking between virtual circuit services of different networks could be realized either by network interworking at the level of datagrams, as above, or at the level of virtual circuits [14], in which case a gateway identifies each virtual circuit that passes through the given gateway. A comprehensive discussion of the difficulties of the latter approach is given by V.G. Cerf [15] and reads as follows:

"Some claims have been made [14] that X.25 can solve the very difficult general network interconnection problem. That is, the problem of communication between two DTEs in different networks. The idea is that the two DTEs could communicate through X.25 interfaces by setting up separate virtual circuits to the network boundaries where a "gateway" would bridge the two circuits and effect the communication. Such a strategy limits the internetworking flexibility since the gateway is a critical part of the circuit. Alternate gateways could not be chosen with the same degree of freedom as, for instance, alternate routes are selected in a packet network. Thus, one of the major advantages of packet switching, dynamic alternate routing, might be lost if X.25 were used to support network interconnection. Even if mechanisms within the network(s) detected a gateway failure and automatically established a new pair of virtual connections between source and destination DTE through a new gateway, it isn't certain how end-to-end robust delivery could be assured by the X.25 protocol, since there would be no way of knowing which packets had or had not been successfully forwarded across the gateway. Consequently, even an X.25 network must rely on higher level DTE-to-DTE protocols for internetwork robustness.

"Moreover, unless non-sequencing networks are explicitly ruled out by the X.25 specification, the existence of an X.25 interface to a network will not guarantee end-to-end sequencing, duplicate detection, and so on. Even if sequencing at the network boundaries were enforced, there remains a potential performance problem, since intermediate sequencing through several networks could easily introduce unnecessary delays.

"X.25 does address the problem of packet size mismatches among packet networks, since the data contents of packets can be fragmented and carried in smaller packets as they cross the gateway. This strategy only works, however, if the networks are able to carry the X.25 "More Data" indicator across the network to the destination DCE. In the X.25 specification, it is not clear whether the receiving DCE should expect to receive such an indication, since its use, as described, is purely to allow for improved data packing in packets issued by the DCE on behalf of those sent to it by the DTE. It would appear that some implicit assumptions concerning the underlying network implementation are necessary before the fragmentation feature of the X.25 protocol can be said to work in the internet environment.

"Alternative strategies for network interconnection have been proposed [10, 11, 12, 13] which depend only on unsequenced datagram services. These strategies allow for substantial variation in network characteristics, provide for packet fragmentation at the network boundaries, concurrent or alternate use of multiple gateways, end-to-end error and flow control, dual homing and loss or duplication of packets in intermediate networks. There is ongoing research to test these ideas through the interconnection of three very different networks [16]. "

Considering in addition to the above discussion the similarity between the datagram and packet broadcast services, we conclude that there seems to be a definite long range advantage of designing the interworking of packet-switched data networks on the architectural level of of datagrams, i.e. on the exchange of individual packets between the networks. However, this approach presupposes that

(a) all interconnecting networks provide the datagram service or, at least, have an internal structure which supports the transmission of individual packets, and

(b) for providing inter-network virtual circuit services, the interworking networks use the same agreed protocol which builds the virtual circuit service on top of the transmission service for individual packets. (We note that such a proposition was contained in the first description of the Datapac access protocol [17]).

4.5 Flow and congestion control

Flow and congestion control are related, since a network needs to limit the incoming data flow in order to avoid network congestion. Therefore these mechanisms have an indirect impact on the quality of service provided by the network.

4.5.1 Flow control mechanisms for datagrams

Depending on where the mechanism applies, we can distinguish the following four mechanisms for controlling, i.e. limiting the flow of datagrams through a network:

(a) At the source DTE:

This is a mechanism by which the network can limit the number of datagrams forwarded by the source DTE to the network. Typically, this mechanism could be provided by the flow control mechanism of the DTE-DCE link access protocol.

(b) In the network:

If a network node (because of congestion or any other exceptional reason) is not able to forward all datagrams in transit, it may drop some. This poses the question whether such a datagram loss should be indicated to the source DTE by a non-delivery indication. This question is related to the option of receiving an acknowledgement for each delivered datagram. Different approaches to these problems have been taken by the datagram proposals, as discussed in Section 5.1. .

(c) At the destination DTE:

This is a mechanism by which the receiving DTE can limit the number of datagrams accepted from the network. Typically again, the flow control mechanism of the DTE-DCE link access protocol is used for providing this mechanism, similarely as for mechanism (a). If the queue of waiting datagrams for a given destination DTE becomes too long, the network is entitled to drop additional datagrams for this destination, possibly returning a non-delivery indication (see mechanism (b)).

(d) Barred access:

This mechanism discriminates, for a given destination DTE, between those source DTEs that are allowed to send datagrams to the destination DTE, and those that are not allowed (i.e. for which the access to the destination DTE is barred). If a datagram arrives at the network node of the destination DTE and its access is barred, the network will drop the datagram, and possibly (see mechanism (b)) return a non-delivery indication. However, it is also possible that the network detects the barred access condition for a given datagram already at the node of the source DTE. This would be a definite advantage in the case when the network is congested due to a large number of datagrams which cannot be delivered.

The closed user group is a particular kind of barred access mechanism. The concept of a closed user group is already been used for virtual circuit networks. Detailed propositions for implementing this facility in X.25 have been made to the CCITT. The idea is that a user group is identified by a user group number and access to a member of such a group is only granted to DTEs belonging to the same group. A given DTE can belong to several user groups. The pertinent user group number has to be indicated for each call establishment in X.25. In the case of datagrams, the number could be indicated with each datagram sent. In the case of the closed user group, the barred access condition can be detected at the source.

A bar access facility has been proposed [18] for providing the possibility for the destination to dynamically bar the access for datagrams from particular source DTEs. This facility could be useful when a particular source DTE floods the destination DTE with a large number of non-relevant datagrams.

We note that, in addition to these flow control mechanisms at the datagram level, the subscriber system will usually contain an end-to-end transport protocol between the communicating subsystems which will control the data flow for each individual communication according to the requirements of the participating subsystems. When these higher-level protocols work correctly datagram loss due to the control mechanism (c) above should be very low.

For providing the options of delivery acknowledgements and non-delivery indications, the datagram protocol has to contain some means for identifying individual datagrams. Generally, a two octet field is foreseen for this purpose. This field can also be useful for the end-to-end transport protocol implemented by the subscriber on top of the datagram transmission service.

4.5.2 Network congestion control for datagrams

For controlling congestion the network can use the four flow control mechanisms described above. A network node experiencing congestion would typically block at the source all datagrams coming from the DTEs directly connected to that node, and would drop, if necessary, datagrams that await delivery to a directly connected destination DTE, and datagrams in transit. Additional strategies for informing the other nodes of the network and appropriate actions for them are much more difficult to elaborate. Although some insight has been gained on overall congestion control [19, 20], there are still many open problems in this area for theoretical and experimental research.

4.5.3 Congestion control: Datagrams versus virtual circuits

Before making a comparison of the flow and congestion control between the datagram and virtual circuit transmission services we explain shortly the flow control mechanism for the virtual circuit service of X.25. For each established virtual circuit, there is a flow control mechanism at the source and at the destination, similar to the mechanisms (a) and (c) above. However, when the destination DTE blocks the incoming packets the network will, after a time depending on its available buffer space, block the source DTE in turn. As far as the number of virtual circuits is concerned, each DTE connected to the network can only support a certain maximum number of circuits at a given time.

The congestion problems for datagrams and virtual circuits (according to X.25) can now be compared as follows.

Datagram service:

Congestion is caused by too many datagrams coming together in a given area of the network. The result is reduced throughput, increased transmission delay and relatively many datagram losses.

Virtual circuit service:

Congestion can be caused by

(a) too many data packets coming together in a given area of the network, and/or

(b) too many call requests coming together to the same node of the network.

In case (a) the result is reduced throughput, increased transmission delay and possibly (hopefully only on rare occasions) loss of data packets indicated by a circuit reset packet. In the case (b) the results are probably delayed call progress signals and call clearing packets indicating the cause of congestion in response to call request packets.

We note that there is a certain similarity of the effect of congestion on the two services: both services have reduced throughput and increased delays, and both may lose data packets. In the case of the virtual circuit service, the throughput can be limited by the network at two levels: (a) limiting the flow individually for each established circuit, and (b) limiting the number of new circuits that are established by delaying the response to call request packets [21] and/or replying to call requests by clear indication packets.

In the case of the datagram service, there is the following additional problem: A malicious DTE may flood a given destination DTE with a large number of datagrams. This may induce a large number of datagram losses, according to mechanism (c) above, also for useful datagrams from other sources, and therefore appreciably reduce the service quality for the given DTE. Possible mechanisms to prevent this from happening are the bar access facility and the closed user group (see mechanism (d) above).

5. Proposals for Datagram and Fast Select facilities

Several different proposals have been recently submitted to the CCITT for providing facilities in packet-switched data networks for supporting applications involving a large number of short end-to-end associations.

Two quite different facilities have been proposed: the Datagram facility [7, 18, 23] and the Fast Select facility [25, 22]. The proposals for a datagram facility are more or less in line with the view that "a datagram service should provide an extremely simple transport mechanism upon which customers can implement their own protocols at a higher level to provide any required flow control, acknowledgements, diagnostics etc." and be "as cheap as possible" [7]. The proposals adhere to the structure and the packet format of X.25 as much as possible. The possibility of interworking between a datagram service and a virtual circuit service has not been considered and remains for further study.

The fast select facility, proposed by Japan, is an extension of X.25 to support more efficiently "simple DTEs which transfer only one inquiry packet and receive one response packet, for example a point of sales terminal" [25]. Because this facility includes the complexity of virtual circuit set up and clearing, it can hardly be considered as an alternative to the datagram facility.

5.1 Proposals for a datagram facility

There seems to be agreement on the general characteristics of a datagram transmission facility as outlined in Section 3.2. There

	Proposal [7] (UKPO)	Proposal [18]		Proposal [23] (ANSI working paper)
		Datagram only	Datagram virtual circuit	
Mandatory function for DATAGRAM INTERFACE				
a) addressing	same as in X.25	same as in X.25	same as in X.25	same as in X.25
b) flow control	insured by Link Access Procedure	insured by Link Access Procedure	insured at level 3	insured at level 3
c) maximum size of a datagram packet	128	128	128	
User facilities				
a) Confirmation of datagram delivery:				
1) acknowledgement not required	default	default	default	default
2) acknowledgement required	opt.	opt.	opt.	opt.
3) discard if not deliverable	always	always	always	
4) return diagnostic if not deliverable	opt.	always	always	
5) deliver to alternate designation if not deliverable				
6) stamp time of delivery				
b) Multiple level of priority				for further study
c) Privacy	closed user group for further study			
d) Bar access	no, but closed user group: yes	yes	yes	yes
e) Echo datagram				
f) Reverse charging	no	yes	yes	

Table 5.1 User facilities to be provided with the datagram service.

is also agreement with the fact that the facility should be provided by a "packet level" protocol at level 3 with an underlying data link and electrical-physical interface (see figure 3.2), as much as possible similar to X.25. Several papers point out the necessity of having a separate "channel" between the DTE and DCE for transmitting datagrams in the case that the DTE also uses the virtual circuit service of the data network. Different schemes for obtaining such a separate "channel" have been proposed [7, 18].

The main functions and facilities* to be provided by the proposed datagram service are summarized for the different proposals in Table 1. Different proposed datagram packet formats are shown in figure 5.1. These differences do not seem very significant. We note, however, that although many of the proposals suggest a maximum user data field length of 128 octets, a recent US document [24] suggests a maximum length of 1024 octets. An universal maximum length must be adopted if datagrams are to be used for internetwork data transmission.

Two proposals give almost complete specifications [7, 18].

The UKPO proposal

We consider this proposal [7] the best because

(a) the protocol structure seems to be as simple as possible,

(b) it provides the basic functions of a datagram facility,

(c) the paper is well written, giving good justifications and explanations.

* The list of functions and facilities is taken from [24].

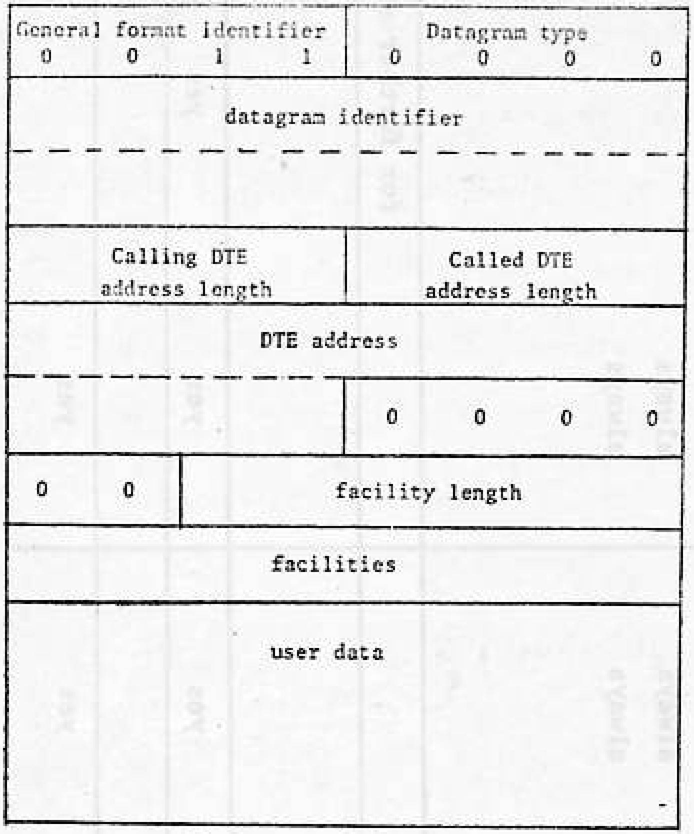

a) Proposal [7] (UKPO)
Data carrying datagram packet

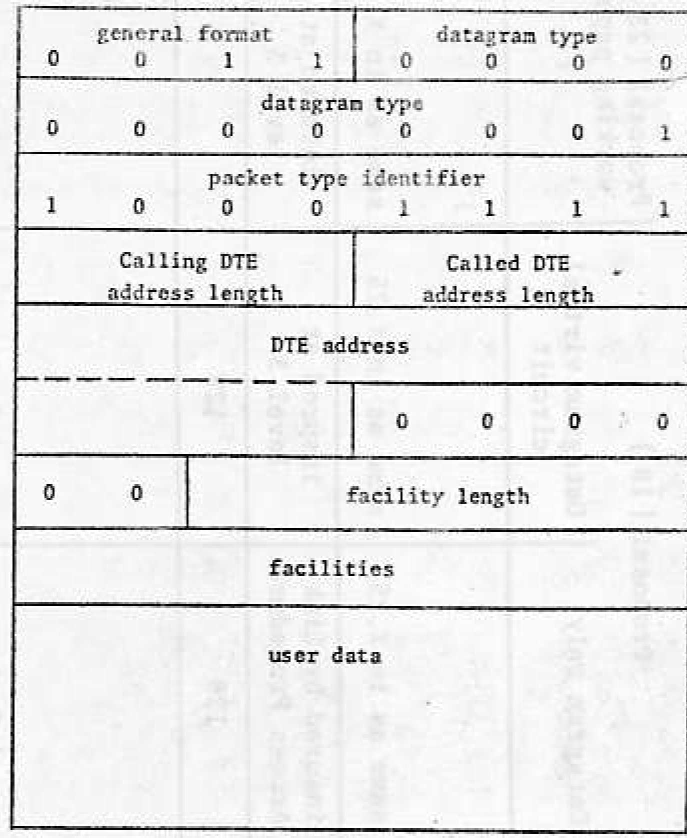

b) Proposal [18] (Solomonides)
"datagram only" datagram packet

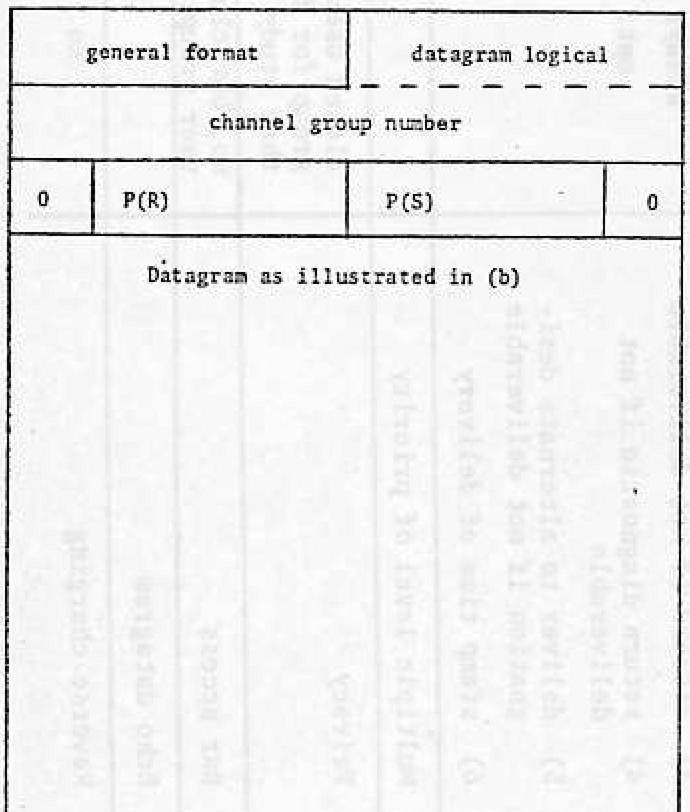

c) Proposal [18] (Solomonides)
Virtual circuit datagram packet

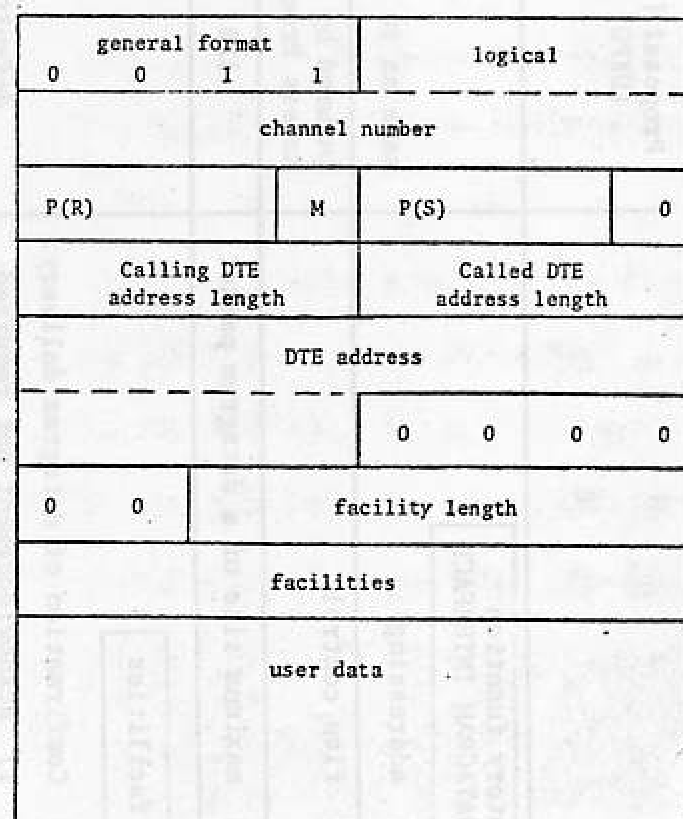

d) Proposal [23] (ANSI working paper)

Figure 5.1
Formats of data carrying datagram packets

As far as the separate "channel" for datagram packets is concerned, the paper proposes to use two identical, but separate link access procedures, one for the virtual call packets and one for the datagram packets. The two identical links would be distinguished by using different values as station addresses in the HDLC header of the frames (see also figure 5.2).

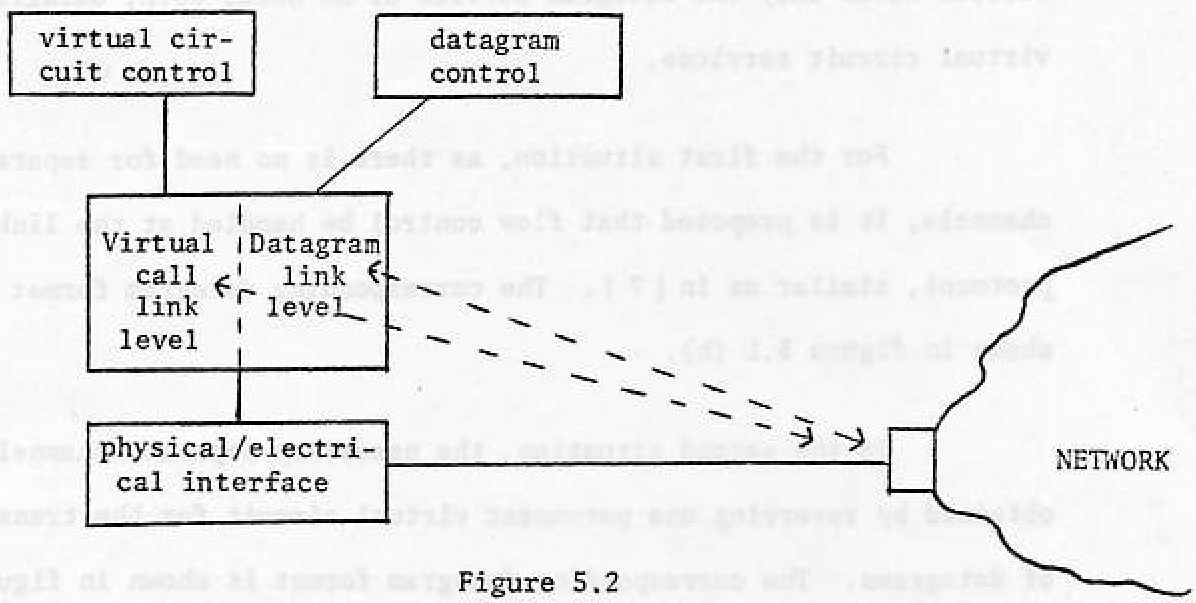

Figure 5.2
Several identical link level procedures for different packet level services.

The datagram flow control at the source and destination DTE (see mechanisms (a) and (c) in Section 4.5.1) are performed by the link access procedure of the datagram facility, thus eliminating duplication of the flow control function in the link and packet levels.

We also note that the datagram header provides space for a two octet "datagram identifier" (figure 5.1). This identifier can be used by the subscriber's implemented end-to-end protocol, and is returned by the network in acknowledge or non-delivery diagnostic service signals.

The proposal by C.M. Solomonides

This proposal [18] distinguishes two situations: The subscriber needs only the datagram service or he needs both, datagram and virtual circuit services.

For the first situation, as there is no need for separate channels, it is proposed that flow control be handled at the link level protocol, similar as in [7]. The corresponding datagram format is shown in figure 5.1 (b).

In the second situation, the necessary separate channel is obtained by reserving one permanent virtual circuit for the transmission of datagrams. The corresponding datagram format is shown in figure 5.2 (c). Datagrams are transmitted over the permanent virtual circuit between the DTE and DCE like ordinary data packets, except that the flow control procedures is slightly changed: "A datagram is always promptly acknowledged by the local DCE to allow uninterrupted flow of datagrams into the network". In this proposal, facilities such as "delivery confirmation requested", "no-delivery confirmation required" and signals such as "inaccessible destination", "bar access request" etc. are coded as different datagrams types (see figure 5.2. (b)) and not as user facilities as proposed in [7].

The user facility field is used to select reverse charging. This facility is made possible by the provision of an authorization code obtained at subscription time which, when confronted with the combination of calling and called DTE addresses, becomes the key to the reverse charging facility.

5.2 The fast select facility

Two proposals [25, 22], not very different from one another, are made to support what the Nippon Telegraph and Telephone (NTT) public corporation has called a "fast select facility".

The overhead of call set up and call clearing necessitated by X.25 becomes very important for applications requiring a simple data transfer such as inquiry/response (figure 5.3) or message/acknowledgement.

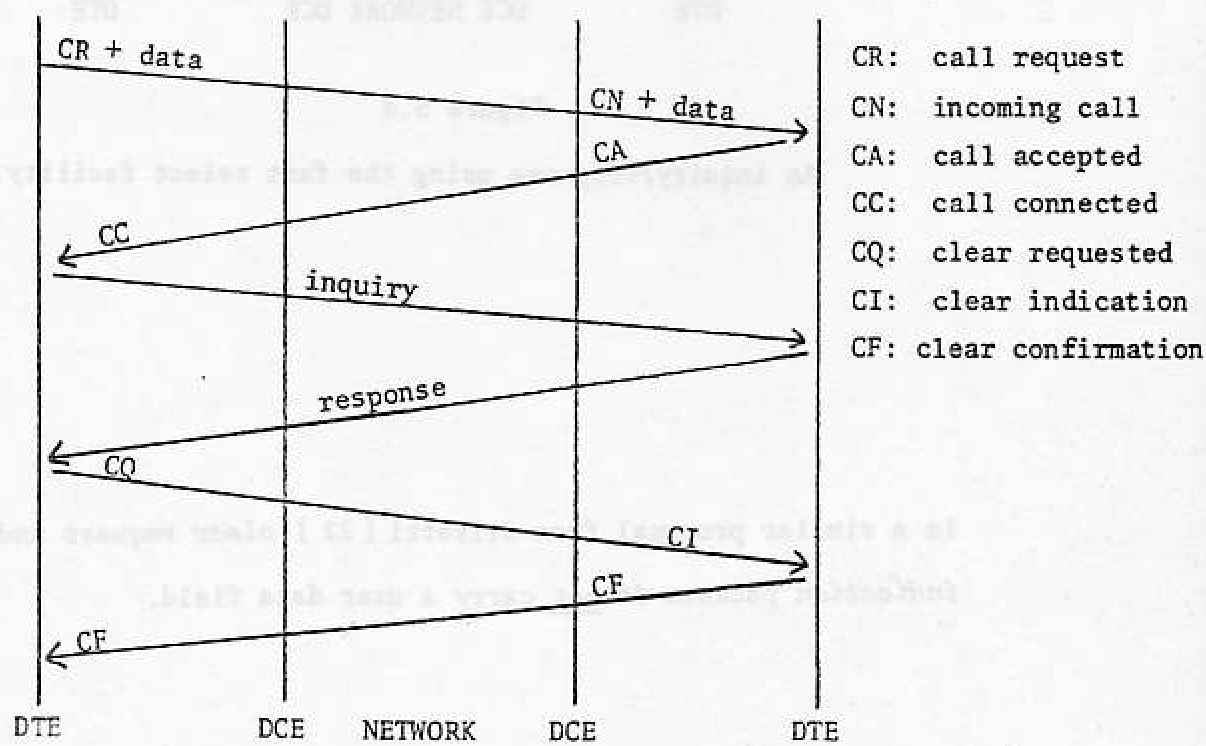

Figure 5.3
The mechanism of an inquiry/response with X.25 virtual circuits.

It was thus proposed by NTT [25] to provide user data fields with the call control packets: *call request*, *incoming call*, *call accepted*, *call connected*, *clear request* and *clear indication*. The user data field should have a maximum length of 128 octets. A simple transaction, with this proposition, is now reduced to a few exchanges as illustrated in figure 5.4 .

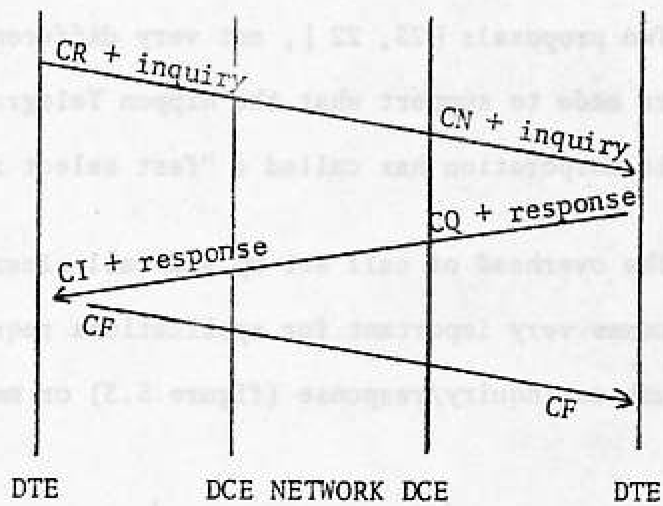

Figure 5.4

An inquiry/response using the fast select facility.

In a similar proposal from Olivetti [22] *clear request* and *clear indication* packets do <u>not</u> carry a user data field.

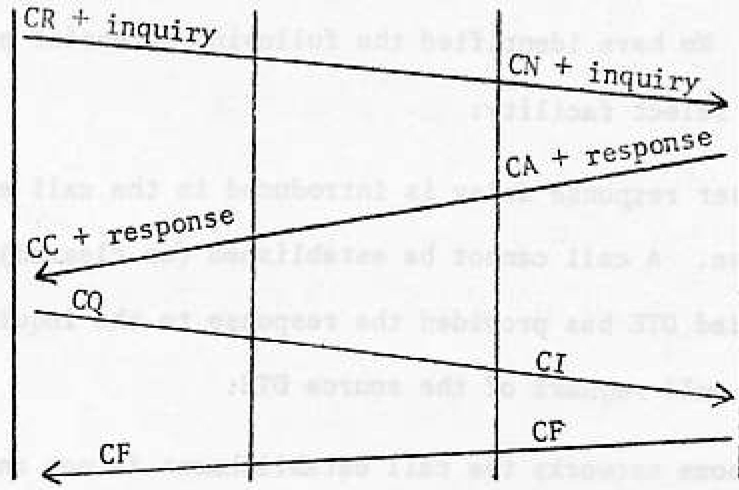

Figure 5.5

An inquiry/response using Olivetti's proposal [22] .

The proposal does not suggest a maximum length for the user data field. Rather it is stated that "the maximum user data field length is agreed for each DTE/DCE interface for a given period of time".

Some minor modifications to allow for end-to-end user controlled handshaking and user specified clearing/resetting causes were also included in the proposal.

5.3 Implications of the Fast Select facility

The implications of the fast select facility on network design and other user facilities such as reverse charging need further study.

We have identified the following potential problems for the fast select facility:

(1) A user response delay is introduced in the call establishment phase. A call cannot be established (or cleared) until the called DTE has provided the response to the inquiry included in the *call request* of the source DTE.

(2) In some networks the call establishment is not an end-to-end procedure. For instance in DATAPAC, the destination DCE sends a *call connected* packet before receiving a *call accepted* packet from the called DTE (for implementing faster call set-up). Such a procedure is incompatible with the fast select facility.

(3) It has been claimed that the fast select facility provides a solution to flow control problems encountered with datagrams. However, there can be congestion due to a large number of call request packets. NTT [21] has proposed two methods to solve this problem:

- To allow the network to delay the transmission of *call connected* or *clear indication* packets just not to exceed a predetermined DTE time-out. (This might be difficult to realize: there is not yet an agreement on DTE time-outs).

- To allow a network to decrease temporarily the number of logical channels available to packet mode DTEs when the network is congested. (This seems to contradict Recommendation X.25).

(4) Reverse charging is difficult. The problem (1) of Section 4.3.2 applies to the fast select facility.

(5) The presence of the user data field in call control packets increments the overhead associated with call set-up, and call clearing. Undoubtedly this will influence the network design.

(6) Is an end-to-end transport protocol necessary on top of the fast select facility? Can packets be lost? For example, what happens if, while an acknowledgment packet is on its way, a restart is initiated by the network?

We conclude that the fast select facility is an interesting extension of X.25 which may be useful for applications such as point of sale systems, etc. However, due to its complexity it can not generally be considered as an alternative to datagrams.

As far as point of sale and similar applications are concerned, the choice between a fast select or datagram transmission facility will depend largely on questions such as:

- does the subscriber want to leave the responsability for the detection of packet loss or duplication to the carrier, or should these problems be handled in the subscriber's equipment?
- Are the terminals used in the application intelligent enough to handle such problems?

The grade of service and tariffs of the transmission facility may also influence the choice.

6. Conclusions

The present study is mainly based on experience gained with different experimental packet-switched data/computer networks. The main conclusion is that a datagram transmission service is an interesting alternative to the virtual circuit service provided according to X.25 by many public packet-switched data networks. It is particularly interesting (a) for applications involving short exchanges of messages, such as point-of-sale systems, and (b) in the cases where the communicating subscriber DTEs are sophisticated enough to implement an end-to-end data transport protocol providing the necessary reliability, such as in the case of intelligent terminals and/or distributed systems of (mini-) computers. For internetworking, the datagram service has important advantages over the service of virtual circuits.

It is important to gain more experience in the use of packet-switched data transmission facilities for every day data processing applications, as well with virtual circuit as with datagram facilities. With the advent of Datapac and Infoswitch, many Canadian users may gain experience with a packet-switched virtual circuit service. We hope that in the near future, in addition, some data communications users will identify the advantages of the datagram service for some particular applications, and a carrier will agree to provide such a service for these applications, which should not be difficult in the case of a data network that internally uses a kind of datagram facility.

We give in the following a list of more detailed conclusions:

- Used with an appropriate end-to-end protocol implemented in the subscriber's equipment, the datagram transmission service can provide very reliable data communication, even in the case of intermittent network malfunction. With multi-homing arrangements, communication can be maintained even in the case of local network node failure, which is not possible with the virtual call facility.

- There is no doubt that a datagram service can be more cost effective than virtual circuits for applications requiring a simple and fast transport service for a small number of short messages. However more work and experience with both virtual circuits and datagrams is needed for better understanding the impact on the terminal equipment, and the suitability of the different data transmission services for different types of applications.

- Interworking of packet-switched data networks can be realized at the architectural level of datagrams or virtual circuits. Interworking on the basis of datagrams seems to be simpler to realize, specially in the case of multi-gateway configurations where high communication reliability is obtained without any complications.

- Sufficient methods for handling flow and congestion control within a datagram providing network seem to be known. In this respect, the subscriber's end-to-end protocol and the closed user group and bar access facilities may play an important role.

- We believe that it is possible to develop accounting schemes for datagrams that are reasonably simple, and acceptable to the subscriber.

- Virtual circuits may present a certain advantage over datagrams for an international integrated billing scheme based on calls.

- The proposals for a datagram facility, which were available for the study, show agreement on the basic characteristics of a datagram service, and on the fact that such a service should be provided by a protocol similar and complementary to X.25. Similar packet formats and the same link and physical level procedures should be used. The remaining differences seem to be of secondary importance. Among the proposals made, the one from United Kingdom Post Office [7] seems to be the most appropriate.

- The fast select facility is a proposition for more efficiently supporting applications such as point of sales systems, with only a minimum modification to the X.25 specifications. However, because of its inherent complexity, it can not be considered as a general alternative to the datagram facility.

7. Points for further study

As mentioned in the conclusions, more experience is needed with the use of packet-switched transmission services for data processing applications. For evaluating the relative merits of virtual circuits and datagrams as a transmission facility within the communication system architecture of data processing systems, we propose the following lines of action:

(1) Determine the data transport requirements (including grade of service) for certain important classes of applications, in particular those classes for which the datagram facility seems especially interesting, i.e. inquiry and transaction systems, distributed processing and broadcast applications.

(2) Partly in relation with point (1), elaborate standards for end-to-end transport protocols, terminal access and other higher level protocols.

(3) Gain practical experience with the use of a datagram transmission facility.

The area of internetworking needs further study. Considering the advantages of using the datagram facility as the basis for the interconnection of data networks, it would be interesting to investigate the possibilities of building datagram gateways for the interworking of public data networks. The following points need particular attention:

(4) Detailed evaluation of different proposed schemes for the interworking of data networks, and their application to existing and planned networks.

(5) Interworking between datagram and virtual circuit services [26].

(6) Establishment of international accounting and billing schemes in the context of a datagram transmission service.

In addition we note the following points for which further studies would be useful:

(7) Intelligent terminals directly interfacing with a datagram transmission service.

(8) Congestion control of in data networks. (There are many problems for theoretical and experimental research.)

(9) The implications of the "fast select" extension to X.25 on the network architecture and its resource management.

(10) Optimal maximum datagram size for universal utilization.

8. Glossary

The following glossary (extracted from [15]) may be useful for clarifying certain notions which occur frequently in the present report.

Communication Protocols

Sets of conventions (formats, control procedures) which facilitate all levels of data communication. Includes electrical interface conventions, line control procedures, digital communication network interfaces, inter-process communication conventions, and application level (e.g., file transfer, database retrieval) standards.

Computer Network

A collection of autonomous but co-operating computers sharing common transmission and switching media. This definition could include computers sharing memory through a common data buss, computers sharing a common terminal concentration network, and computers using a common carrier network to communicate among themselves, for example.

Data Network or Data Communication Network

A network providing data transmission facilities.

Congestion

A situation in a packet network in which all or nearly all buffer capacity is used up, leading to excessive retransmission and potential lock-up.

Datagram

A packet of data to which sufficient addressing and control information is affixed to allow the packet to be routed and processed independently of all others in a store-and-forward network. Delivery of a datagram does not depend upon its arrival at the destination in any particular order, for example.

DCE (Data Communication Equipment)

Usually taken as the extreme boundary of a data communication network.

DTE (Data Terminal Equipment)

The device which connects to the DCE and acts as a source or sink of data crossing the boundary. A DTE can be a computer, a telecommunication device, etc. Often the distinction is made on the basis of ownership whether a device is a DCE or DTE.

Multi-Homing

A method of networking in which a DTE is connected simultaneously to more than one DCE. A routing method used to accommodate this sort of multiple connection.

End-to-End

Usually taken to mean from DTE to DTE in a network environment. In the case of multiple networks, it means spanning any intermediate networks as well.

Flow Control

A technique used by a receiver to limit the amount (or rate) of data flow transmitted by a sender.

Gateway

The logical or physical interface between computer networks. The mechanism by which data passes from one network to another.

Internetworking

The process of interconnecting two or more distinct computer communication networks. (See Gateway.)

Multiplexing

Time or space-sharing of a frequency band among several transmitters.

Packet

A short (128-2000 bits) block of data, prefixed with addressing and other control information, which is used to carry information through a packet switching network.

Packet Switching

A form of message switching which facilitates asynchronous time-division sharing and switching of transmission resources. Particularly suited to bursty communication requirements; permits the sharing of common switched transmission resources among both high bandwidth and low delay applications.

Switching

Provision of point to point transmission between dynamically changing sources and sinks (i.e., different sources and different sinks are connected and disconnected as required).

Virtual Circuit

A synthetic equivalent of a real circuit (point to point) derived from a store-and-forward packet network. Data packets are kept in sequence at the delivery point (though not necessarily in transit).

List of references

1. L. Pouzin, "Virtual circuits versus datagrams, technical and political issues", AFIPS, National Comp. Conf., New York (Jan. 76), 32 p.

2. P.J. Kirstein, "Planned new public data networks", Computer Networks, Vol. 1, No. 2, Sept. 76.

3. CCITT, Recommendation X.2 .

4. CCITT, COM VII 54, Part IV, Annex 2, "Possible procedures of datagram facility" (Nov. 1976).

5. ISO, TC 97, SC 6, Rapporteur's report on Communication Systems Architecture and End-to-End Protocols, No. 1417, 1977.

6. N. Naffah, "Presentation du système Tipac", Doc. Cyclades - Ter 524 (Nov. 1975).

7. UKPO, "Proposals for a datagram service", CCITT, COM VII 24 (Oct. 1976).

8. I. de Sola Pool, "International implications of computer communications", Telecommunications Policy, Dec. 1976, p. 33.

9. D. Lloyd, P.T. Kirstein, "Alternate approaches to the interconnection of computer networks", Proceedings European Computing Conference on Communication Networks, England, Online, 1975.

10. V. Cerf, R. Kahn, "A protocol for packet network intercommunication", IEEE Trans. on Communications, Vol. COM-24, No. 5, May 1974.

11. V. Cerf, A. McKenzie, R. Scantleburg, H. Zimmermann, "A proposal for inter-network end-to-end protocol", IFIP WG 6.1, General Note 96, Sept. 1975.

12. L. Pouzin, "A proposal for interconnecting packet switching networks", IFIP WG 6.1, General Note 60, March 1974.

13. D. Lloyd, M. Gallard, P. Kirstein, "Aim and objectives of internetwork experiments", IFIP WG 6.1, Experimental Note 3.

14. L. Roberts, "International interconnection of public packet networks", Proc. ICCC, August 1976, p. 239.

15. V.G. Cerf, "Research topics in computer communications", in Computers and Communications, Proc. (US) Federal Communications Commission Planning Conference, Nov. 1976, p. 56.

16. This research is supported by the Defense Advanced Research Projects Agency and uses its ground packet radio network, packet satellite network, and the ARPANET as a testbed.

17. TCTS, "Standard network access protocol for Datapac", Nov. 1974.

18. "A proposal for the addition of the datagram facility in the CCITT draft recommendation X.25", C.M. Solomonides, IFIP TC 6.1, Protocol Note 47 (1976).

19. W.L. Price, "Simulation studies of an isarithmically controlled store and forward data communication network", Proc. IFIP Congress 1974, pp. 151-154.

20. L. Pouzin, "Flow control in data networks - methods and tools" Proceedings ICCC (Aug. 76), pp. 467-474.

21. Japan, "Study report on fast select facility", CCITT, COM VII 64 (Jan. 1977).

22. Olivetti, "Comments on Recommendation X.25, inclusion of the datagram user facility", CCITT, COM VII ..., (Nov. 1976).

23. USA, "Working paper on datagram procedures", CCITT, COM VII 25, (Nov. 1976).

24. USA, "Generic requirements for datagram service", ISO, TC 97, SC 6, N 1403 (Feb. 1977).

25. Japan, "A proposal for addition of data transfer function to call set up and clearing phase (fast selecting facility)", CCITT, COM VII 29 (Nov. 1976).

26. J.H. Sexton, "End-to-end protocols, virtual calls and the IIASA Network", IFIP, TC 6.1, Protocol Note 46 (June 1976).