TECHNICAL MEMORANDUM

67

Local area networks; Why What How

W. Jensen

Operations Department

October 1982

This paper has not been published and should be regarded as an Internal Report from ECMWF.
Permission to quote from it should be obtained from the ECMWF.

European Centre for Medium-Range Weather Forecasts
Europäisches Zentrum für mittelfristige Wettervorhersage
Centre européen pour les prévisions météorologiques à moyen

CONTENTS

Abstract

The paper gives a brief review of some of the background for the rapid development of local area networks and distributed computing.

A discussion on the basic characteristics of local networks concludes with a proposal for an overall taxonomy. The low level data transfer aspects define only a tiny fraction of all problems associated with distributed computing. Finally, a presentation of some of the issues presently discussed at ECMWF is given.

Talk given at
2nd meeting of Computing Representatives - 15 October 1981
ECMWF, Reading

Remarks

Bob Metcalf, the constructor of Ethernet, has written "anybody can build a local network - and that's what everybody does". This reflects the current situation which is one of confusion. A lot of so called local networks are running in laboratories, but very few are commercially available. Discussion focuses on bit rates of 5, 10, 20, 50 M bps, but less on what to do with the "flying" bits, i.e. how to achieve meaningful application-to-application communication. Moreover, the fact that effective date exchange rate is 1-10% of the maximum cable speed is seldom mentioned.

I will try to review some of the background for local computer communications. I will say a few words about the technical state of the art for local networking and I will pay attention to the very complex problems involved in obtaining multivendor compatibility. Some parts of my talk will have direct relevance to the present situation at ECMWF where planning has just started.

1. Interconnected Computer Systems: WHY?
2. Local Area Networks: WHAT?
   - basic characteristics
   - taxonomy
   - examples
   - commerical availability and ECMWF relevance
3. Stepwise construction of powerful computer services: HOW?
   - constructive long term planning is necessary
   - standardisation is important
   - technological development is fast but industrial development is slower

1. Interconnected Computer Systems: WHY?

We will consider computer services as tools for general improvement of quality and efficiency in information work. This includes:

- processing services: creation, modification, calculation
- archiving services: storing, structuring, retrieving
- presentation services: displaying, reproducing, printing
- communication services: message exchange between humans, file exchange between computers, etc.

By and large, these classes of services are equally important. It is unlikely that one single computer could be the "best" provider for services which are very different in nature (cf . number crunchers, interactive service providers, word processors, etc.). As technology develops and experience is gained, the principal user requirements to computer services become:

availability - reliability - adaptability (and economy).

The monolithic (one single mainframe) and the anarchic (each user group its own mini) approaches have to be rejected for rather obvious technical and economical reasons. We can now watch a controlled fragmentation process based on interworking computer resources (see Figure 1). The various resources and user terminals are interconnected - physically and logically - according to sets of communication protocols. ISO's reference model for Open Systems Interconnection defines the architectural framework for specification of such standard protocols. Computer systems designed according to this functional structure with a common set of protocols implemented are able to cooperate through meaningful information exchange. Such systems are called Open. Some of the characteristics of Interconnected Open Systems are (see Figure 2):

Technical

- quality of service could be improved through use of dedicated, specialised resources.
- duplication of resources and parallel servicing could increase reliability.
- decomposing a very complex total system could be of benefit for designers and users as well as for maintenance.

User availability

- the human interface could be individualised by use of "intelligent" terminals;
- some of the most often used services - like program development, document preparation, mail - could be located to particular resources (e.g. the terminal itself) thus improving availability and approaching personal computing.

Economy

- shared use of expensive resources (storage, printers, number crunchers, etc.)
- the extremely high degree of modularity invites stepwise upgrading and expansion by adding new resources according to needs and economy.
- the multivendor compatibility requirement implies that benefits from reduced costs through competitive procurement could be achieved.
- the distributed nature and the emphasis on communication should promote collaboration between user groups locally and remote, thereby reducing overlapping efforts and making exchange of data, programs and results easier.

Organisational flexibility

- the architectural aspects of interconnected open systems could be adapted to any organisational structure (centralised, hierarchical, decentralised, etc.);
- distribution of resources and local "intelligence" introduce elements of local control and influence in the users' environment, thus increasing their self-confidence and, consequently, efficiency.

- use of computer-assisted communication services (electronic mail, message systems, teleconferencing, etc.) as supplements to ordinary telephone, telex, post and face-to-face communication are considered to be of great value in most cases. Experience indicates that these very fast and highly available computerised communication services develop a new "style" of working.

Obstacles

- computer interconnection requires comprehensive agreement on various protocols. This is standardisation which normally is a very slow process. Work in ISO, CCITT, ECMA, etc. on protocol standardisation is very promising (e.g. see Figure 3). However, many are tempted to (or have to) invest in non-standard intermediate solutions which of course lead to subsequent incompatibility.
- the complexity (and therefore cost) of interconnecting computer systems is significant. The physical interconnection (i.e. the cable, wires, modems, etc.) is a very minor part of the total application and system compatibility problem including distributed management and resource control. The cost will, however, clearly decrease as the components and protocols become standardised and thus widely spread.

2. Local Area Networks - WHAT?

Traditionally, so called local networking emerged from aims slightly different from those of remote networking 10-15 years ago. From rather straightforward terminal (RJE) access to distant mainframes, networking now involves highly dynamic inter-computer communication. Discussions on computer communication very often focuses on characteristics like:

separation	bit rate	
> 10 km	< 0.1 Mbps	remote networks
10 - 0.1 km	0.1 - 20 Mbps	local networks
< 0.1 km	> 20 Mbps	multiprocessors

Personally, I think there are characteristics which could be associated to local networks that are equally important:

- private: local networks are not subject to PTT constraints and regulations (legal, technical standards, economical);
- new technology: the private and local nature encourages early use of new technology in experimental set ups.
- simple: the local character and experimental nature of most of these networks makes extremely simple interfaces and protocols feasible.

- low cost: one of the basic requirements for local networks is low cost of the transmission medium itself as well as for the various interfaces to the medium. While most of the cables and wires are cheap, the actual interfacing of resources, terminals and gateways typically leads to significant expenditures. This is reflected in the current commercial situation when very few "local networks" are available.

- new services: in parallel to the development of local networks - and partly caused by it - we are watching the emergence of "new"services: archive services based on dedicated, advanced file servers, graphical presentation services, exchange of programs, e.g. by core swapping, human oriented communication services, etc.

3. Taxonomy and examples - HOW?

I think that the establishment of a taxonomy could be of help in the present somewhat confusing situation. Most of the existing attempts at classification of local area networks are based on very low-level characteristics. This is indeed no surprise; it is at the lowest signal propagation level and in the medium access methods that the existing local networks differ. One can develop different taxonomies depending on a chosen criteria and the actual context. I have chosen one which should be related to ISO's OSA, and I hope it can contribute to conceptual clarification.

Family 1: General store and forward networks (Figure 4)

- point-to-point based general topology
- usually V.24 or X.21 physical interface
- general error-correcting data link control (X25, HDLC).
- packet switching at Network Layer to compensate for partial physical connectivity.
- datagram or virtual circuit based network services (ARPA-protocols, X25).
- end-to-end protocols at Transport Level introducing additional error control, address and flow control capabilities.

These networks are similar/identical to the long-haul classical packet switched network, like ARPA, X25 public networks, a.s.o.

They have the same functional structure and protocols as their long-haul counterparts but the physical parameters like distance, delay, number of stations a.s.o. are scaled down.

Examples of these networks are local ARPAnet (incestuous traffic) DECnet, local UNINETT (Norwegian), NPLnet (British), HMInet (German).

Their major advantage is that (more or less) standardised protocols are used. These can be "bought off the shelf" (e.g. X25) and internetworking is particularly easy. They are very appropriate for interactivetraffic but, though the speed is often 9.6-64k bps locally, these networks are not suitable for many of the new applications that require very fast transfer of (large) files.

Family 2: Circuit switched and strict star networks (Figure 5)

- point-to-point based star shaped physical topology.
- X21, X22 or V24, RS 232 interfaces for medium access and FDM/TDM techniques for trunk circuit sharing.
- general error-correcting and flow controlling protocols across the data link between end stations (X75, HDLC).
- in the circuit switched case, the Network Layer is responsible for the call set up operation (X21 call procedure), but the circuit concatenation in the switch is invisible during transfer phases.
- data links could be multiplexed, e.g. between time-sharing resources and terminal concentrators.

This is a centralised structure where the circuit switch or the centre of the star acts as master controlling the slave stations (terminals) through some polling technique.

The homogeneous computer manufacturers' terminal networks are examples of Family 2 elements. The ECMWF Gandalf system is another example. Of particular interest for the coming years is the use of PABXs for local voice and data traffic.

These networks utilise normally well known technology; they depend, however, upon some central control mechanism and the data transfer speed is, so far, too slow (up to 64k bps) for advanced computer-to-computer traffic.

Family 3: Circular Networks (Figure 6A)

- point-to-point based circular topology.
- various techniques for medium access are used:

  loop - central station masters access through polling

  ring - (almost) distributed access like empty slots, token passing, register insertion.
- different link control protocols from full SDLC with error and flow control to the simple mini-packet approach (Cambridge ring - see Figure 6B) with rudimentary flowcontrol.
- the address recognition, i.e. the "routing" - is normally performed at Data Link Layer. Virtual circuits, including end-to-end error control could be provided.

The most well-known example of a circular local network is the Cambridge ring which is one of the few local networks commercially available. The Cambridge ring operates at speeds around 10M bps, and it seems to satisfy many requirements to local networks. However, the ring depends on active elements in the medium itself (repeaters) and this could lead to lower reliability for the total system of interconnected stations.

Because the ring is uni-directional, fibre optics may be used as medium (replacing twisted wire or coax cable). There is an experimental ring at Toshiba running at 100M bps.

Family 4: Broadcast packet switched (contention) networks (Figure 7A)

- multiple-endpoint (broadcast) physical connections using a variety of media (coax cable, CATV, satellites, radio, ...).
- various medium access techniques based on contention strategies to achieve sensible sharing of the broadcast channel:
  - pure ALOHA strategy: collisions are assumed when acknowledgement is not received.
  - slotted ALOHA strategy: sending is only permitted at the beginning of certain time intervals, slots. This improvement doubles the channel utilisation.

. CSMA: stations are forced to listen to the channel before transmission. Collisions may now occur if two stations begin to send at nearly the same instant. Resolution must be taken care of by higher level protocols.

. CSMA/CD: stations listen before transmission and they listen while transmitting. This strategy makes the channel available for retransmission sooner.

- collision detection and retransmission are controlled from the Data Link Layer. Address recognition, i.e. establishment and maintenance of system-to-system connections is also the responsibility of the Data Link Layer. The high capacity and low error-rates of most of the media have lead to use of very simple link protocols (e.g. compared to HDLC).
- most "Ethernets" are based on datagram services with end-to-end control introduced at higher levels. Because of the limited range of most of these bus networks, internetting is a most important issue. The PUP Architecture for interconnected Ethernets (mainly through Arpenet) is well known (and, in my opinion, not particularly well structured).

The broadcast baseband bus networks are conceptually very simple and lots of them are running in various laboratories around the world. It has, however, been difficult to construct low-price controllers for higher speed than about 5M bps. The DIX Ethernet running at 10M bps will become available during the next 18 months. This version conforms with the forthcoming IEEE 802 standard. The medium itself is passive which makes it tolerant to station failures.

The bus network seems to be appropriate in most office environments. However, as the number of stations increases, collisions will be more frequent. It is important to stress that the effective data exchange rate for two computer processes communicating through a contention bus network will be much below the maximum burst speed. This is due to collision handling, protocol overhead, computer capacity, etc.

User demand for a local network that interconnects processors from several vendors has increased significantly in recent years. One of the few commercial offerings that supports this is Hyperchannel, developed by Network Systems Corporation in 1976 (see Figure 7B). Like Ethernet, it is a serial-bus architecture, but with different design goals. It was designed to handle the data-movement needs of large-scale mainframes and large minicomputers and to support communications at high computer-data-channel rates.

The Hyperchannel bus is a single strand of coaxial cable, up to a mile long, that can be multidropped at up to 64 different locations. Each computer site has a "network adapter" that interfaces the cable and functions as a high-speed controller on the computer's data channel. An adapter can broadcast on the cable in "frames" that vary in size from 12 to about 4.1k bytes, and the burst speed could be about 50M bps.

Broadband media are attractive because users can multiplex digital, voice, and video signals on a single communications carrier. On a broadband cable used for computer communications, a frequency range of about 150 MHz (or 150 Mbit/s, if all traffic is data) can be allocated.

Broadband networks are potentially a fine solution to users who need to carry a mixture of low- and medium-speed links (100kbit/s or less) between their computers and terminals.

Broadband employs CATV, not digital, technology, and this means that any interface to be broadband cable requires an RF (radio frequency)modem to perform the digital-to-RF conversion. While these modems vary significantly in price, depending on efficiency, the sheer number required in local network with many nodes might prohibit the implementation of this medium for some time to come. It might be argued, however, that the RF transmission used for broadband could be the most viable technology for future-office-type networks - where not only digital data but analog voice and video will need to be carried as well.

Mitre Corporation has developed several CATV based bus networks.

4. The ECMWF situation

The ECMWF computer services are already, to some extent, based on interconnected and interworking Open Systems. The CRAY-1 is a dedicated resource and, though the two CYBERs are general purpose computers, they will have partly different responsibilities (front-end for batch processing and interactive). The plans for purchasing a common mass storage resource and also additional processing power and linking all these resources by a local network are just in line with current thinking on computer service architectures.

These initial phases are by far the most important and difficult. It is not only a problem of deciding which cable and interface unit is most appropriate, it is more important to select an architecture, protocols and software products that facilitate future extensions and replacements. The necessary analysis must include considerations on:

- ECMWF computer services development, e.g. to what extent will the various office services be computerised (text, archives, mail, etc.), what kind of interconnection strategy should be applied to the meteorological workstations, what development could be foreseen for the remote network, its protocols and its relationship to the local network?

- Standardisation efforts within ISO, CCITT, NBS, IEEE, ECMA and other bodies will have significant impact on future computer systems architecture and software. Thw work within ISO and ECMA on so-called higher level protocols is aimed at providing multivendor compatibility and the services necessary for distributed application processes to interwork through fast, reliable and flexible exchange of information. The problems associated with physical interfaces and logical access methods to a local communication medium are only a very tiny fraction (say 1%) of the total multivendor compatibility and distributed management problems. So, it is very likely that "we will be climbing the ISO-ladder during the eighties" (Robert M. Metcalfe).

- Technological developments take place on several levels. We do know something about what will be generally available in computing and communication the next five years because this technology already exists in laboratories. It is more difficult to predict exactly when and at what price. We know less about the likely development in the laboratories.

- However, fibre optics, infrared radiation, CATV and digital PABX systems will strongly affect local networking.
- Public networks, satellites and the emerging standards will be key issues for all remote exchange of data.
- VLSI technology is certainly becoming important for efficient and cheap implementation of complicated functions.
- New comfortable and powerful devices for presentation and registration of data will have significant impact on man-machine interaction (rastergraphics, laser printers, voice synthesizers, etc.).
- Methods for software engineering will still be improved. The immense value of portable software, e.g. when implementing protocols, is generally accepted. I am, however, afraid that programming still develops much slower than the other fields.

These coming tools will most likely increase the quality of computer services significantly. For computing centres, including ECMWF, it is most important that these components can be easily included as soon as economy permits.

5. A Possible ECMWF Strategy (Figures 8A, 8B)

Depending on the available economical resources, including technical staff, there are different ways of working to start the construction of an open and interconnected computer systems architecture. The interconnecting mechanisms (the local "networks") are very basic parts of this structure. So are the mass storage resources, too.

Some of the local network products on the market are merely tables, taps and network access units with very low level protocols. This is, for example, the approach taken by NSC in their HYPER-channel. One then has to specify and implement the higher layer protocols, e.g. advanced error and flow control, dialogue synchronization, format and presentation conversions, file transfer and remote job entry control. Other products are more system-oriented and include implementation of higher level services and protocols. A few of these products apply to multi-vendor environments.

In the present situation, the following points indicate a proposal for a possible ECMWF strategy:

1. Select system oriented "local network" products, e.g. the Massnet of MASSTOR Corporation which also can provide mass storage resources (shared virtual storage system) and communication software for many hosts (resource systems). Massnet uses HYPERchannel hardware.

2. Study the potential products carefully with respect to future replacement/addition of stations, in particular, special purpose mini, workstations and gateway stations to other local subnetworks.

3. Study the potential products and their vendors carefully with respect to the standardisation issues, in particular their relationship and attitudes to the work in CCITT/ISO (for interconnection with long haul networks) and ISO/IEEE 802 for interconnection with other local networks.

4. The adaptation of the presently running computer services to a distributed architecture is of course of utmost importance. This process will probably involve development of distributed management services, including common user interfaces allowing the user (program) to be unaware of the actual physical fragmentation of the resources.

This very rough proposal is meant to represent a pragmatic compromise between the ideal top-down design approach, and the most used ad hoc bottom-up approach. An overall conceptual architecture serves as the framework for the actual development taking place among available products and preliminary standards with limited economical and human resources.

Figure 1: OPEN SYSTEMS INTERCONNECTION, NO STANDARD

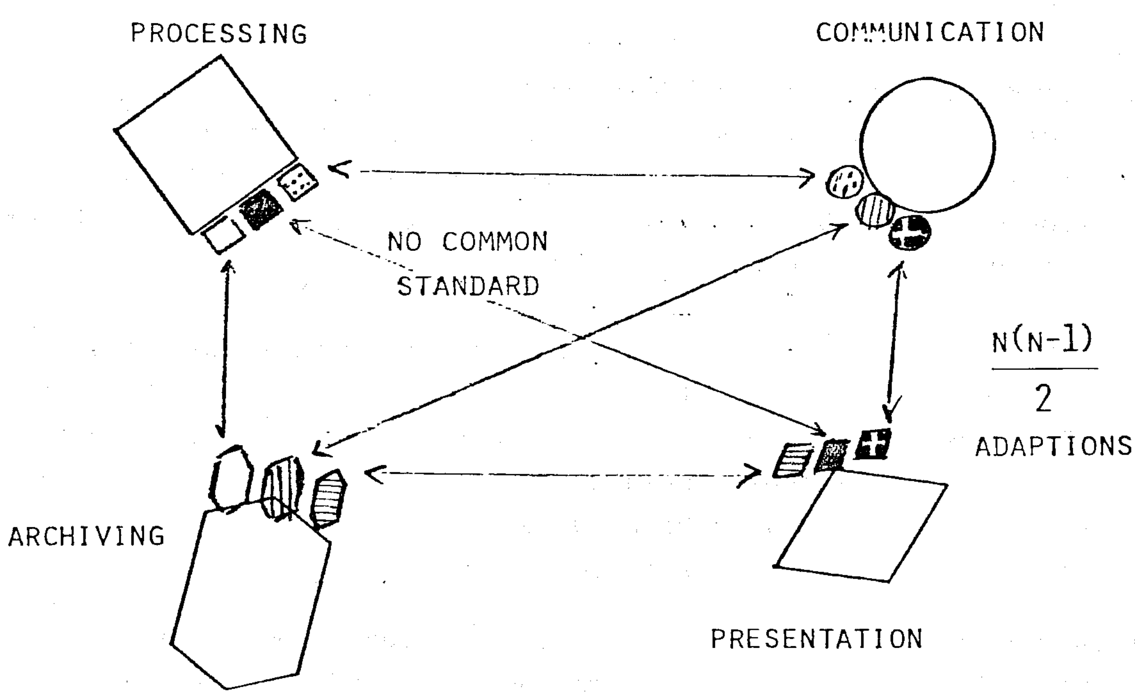

- DEDICATED SPECIALIZED RESOURCES
- SHARING OF RESOURCES
- COMMUNICATION SERVICES

Figure 2 : OPEN SYSTEMS INTERCONNECTION, WITH STANDARD

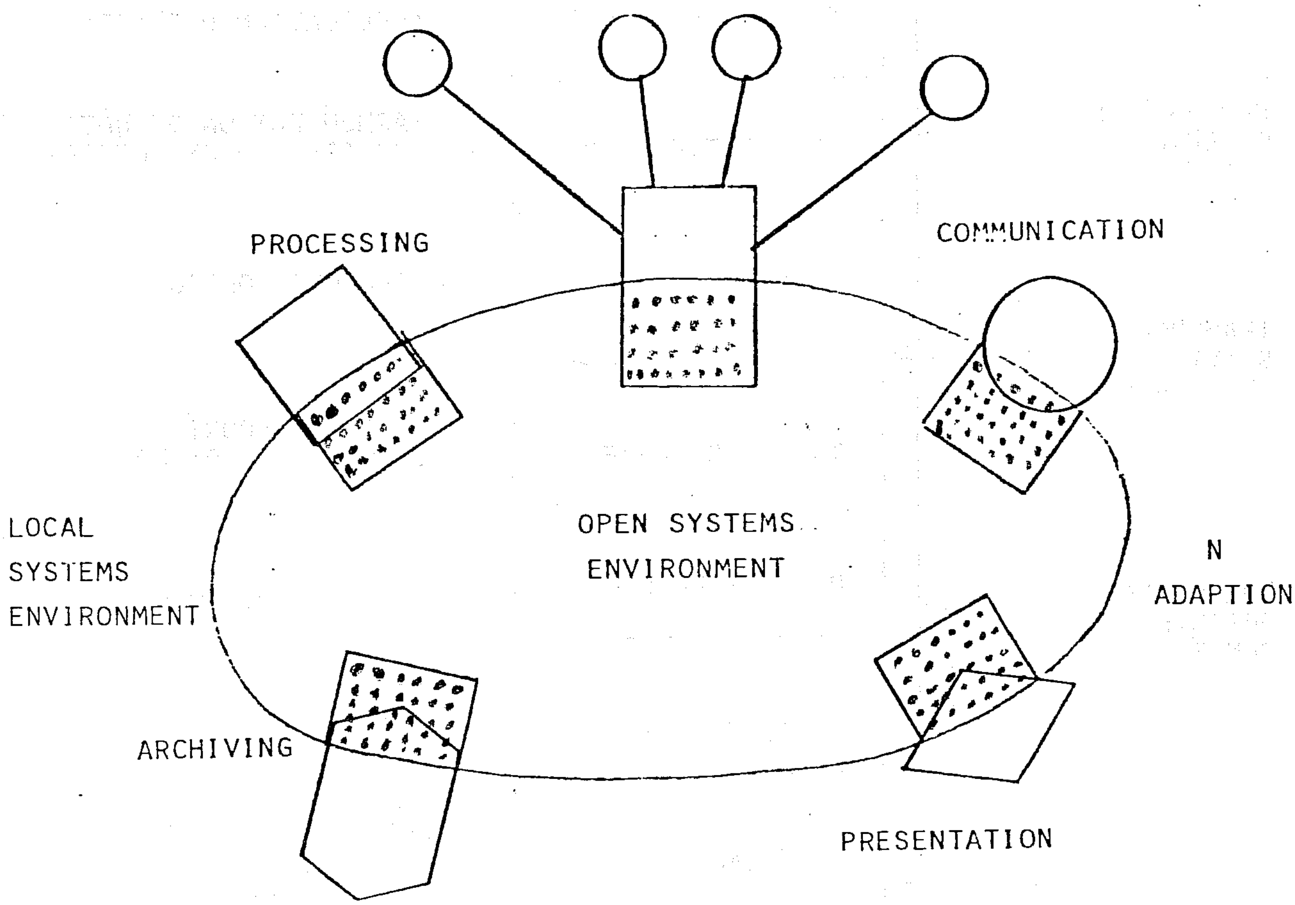

- DEDICATED SPECIALIZED RESOURCES

- SHARING OF RESOURCES
- COMMUNICATION SERVICES
- MODULARITY
- USER AVAILABILITY
- ORGANIZATIONAL FLEXIBILITY

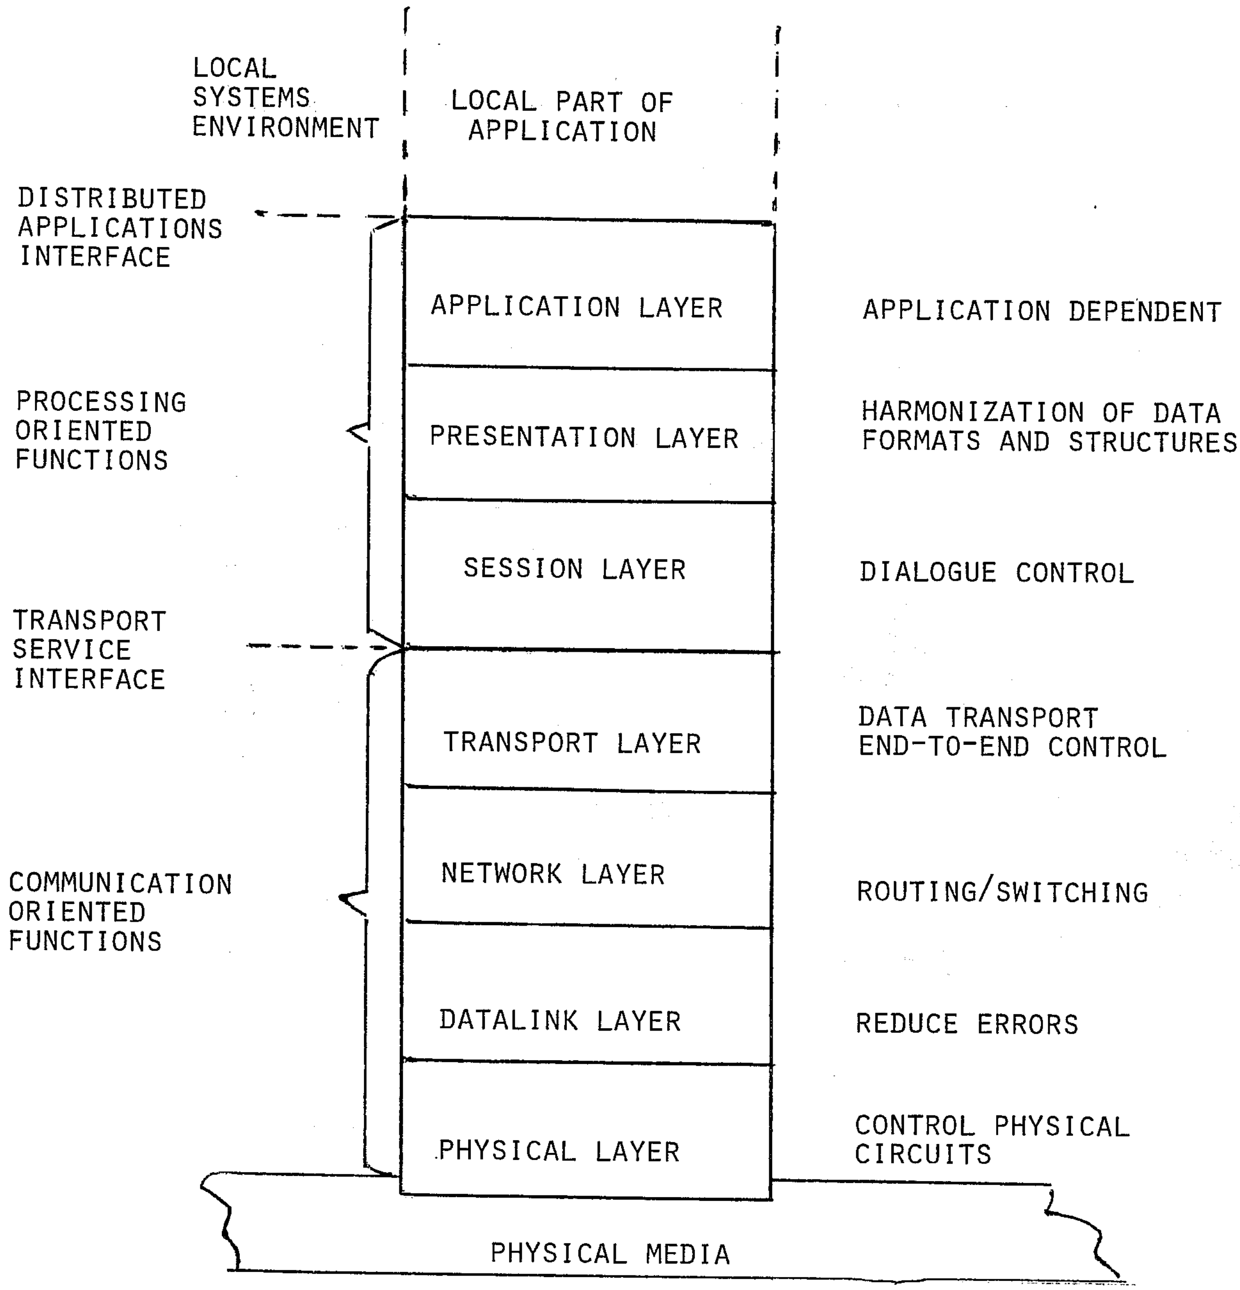

Figure 3:FUNCTION MODEL FOR OPEN SYSTEMS INTERCONNECTION

Figure 4: GENERAL STORE AND FORWARD NETWORKS

- POINT-TO-POINT BASED MESH TOPOLOGY
- V24, X21
- ERROR REDUCING LINK CONTROL
- DATAGRAM OR VIRTUAL CIRCUIT
- PACKET SWITCHING
- END-TO-END TRANSPORT PROTOCOLS

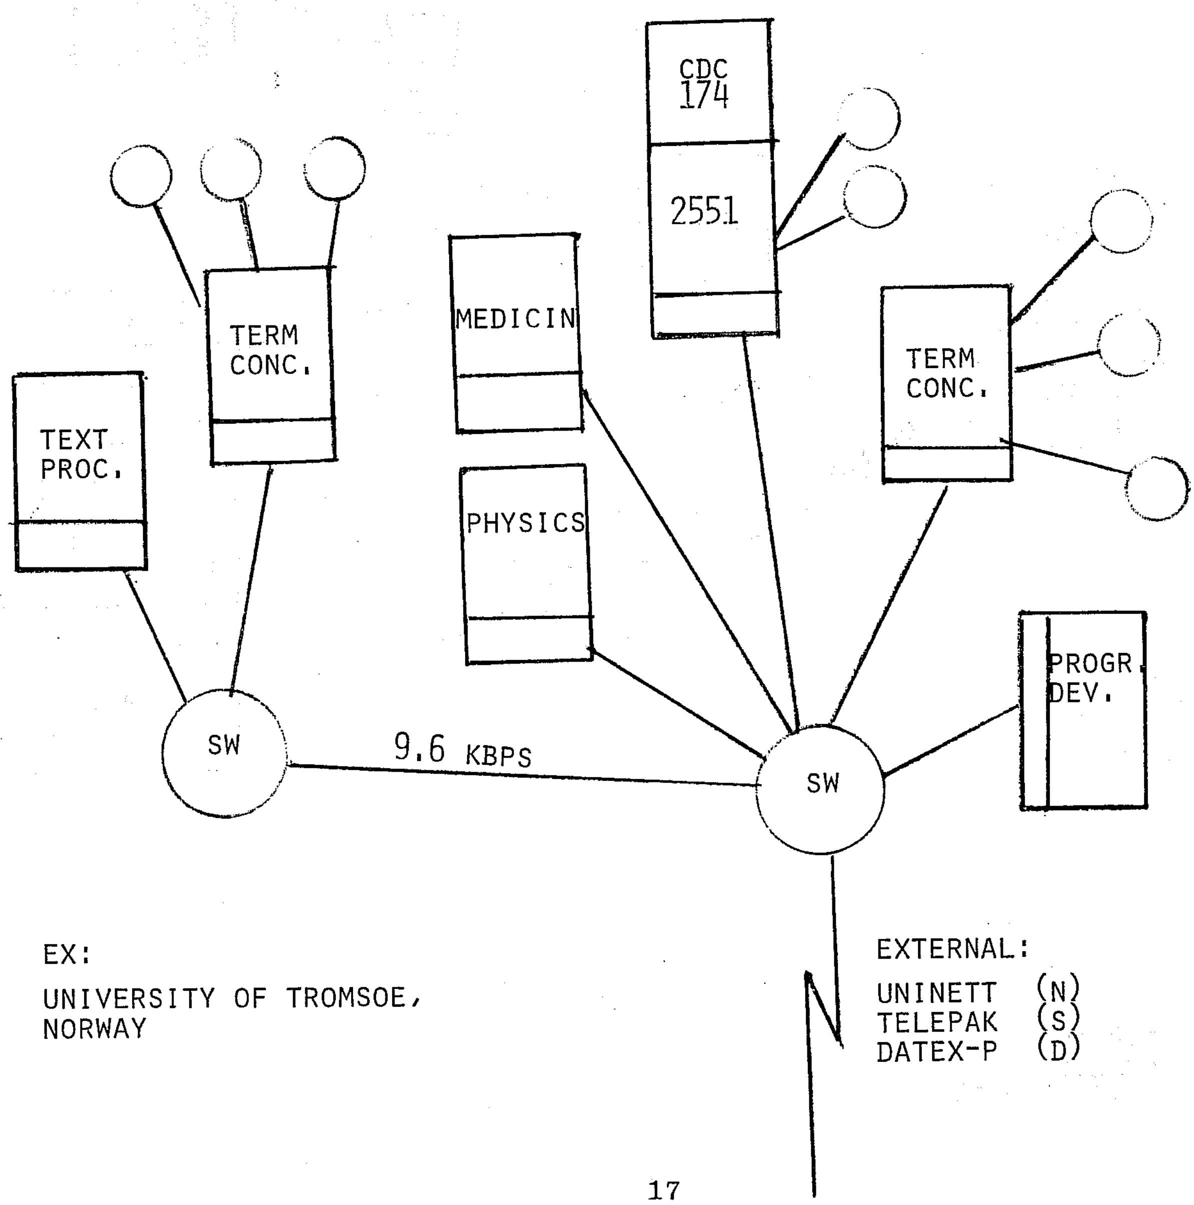

Figure 5 : CIRCUIT SWITCHED AND STAR NETWORKS

- X21, X22, V24 PHYS. INTF
- FDM, TDM FOR TRUNK SHARING
- ERROR REDUCING LINK CONTROL BETWEEN END SYSTEMS
- CENTRALISED NETWORK CONTROL STRUCTURE

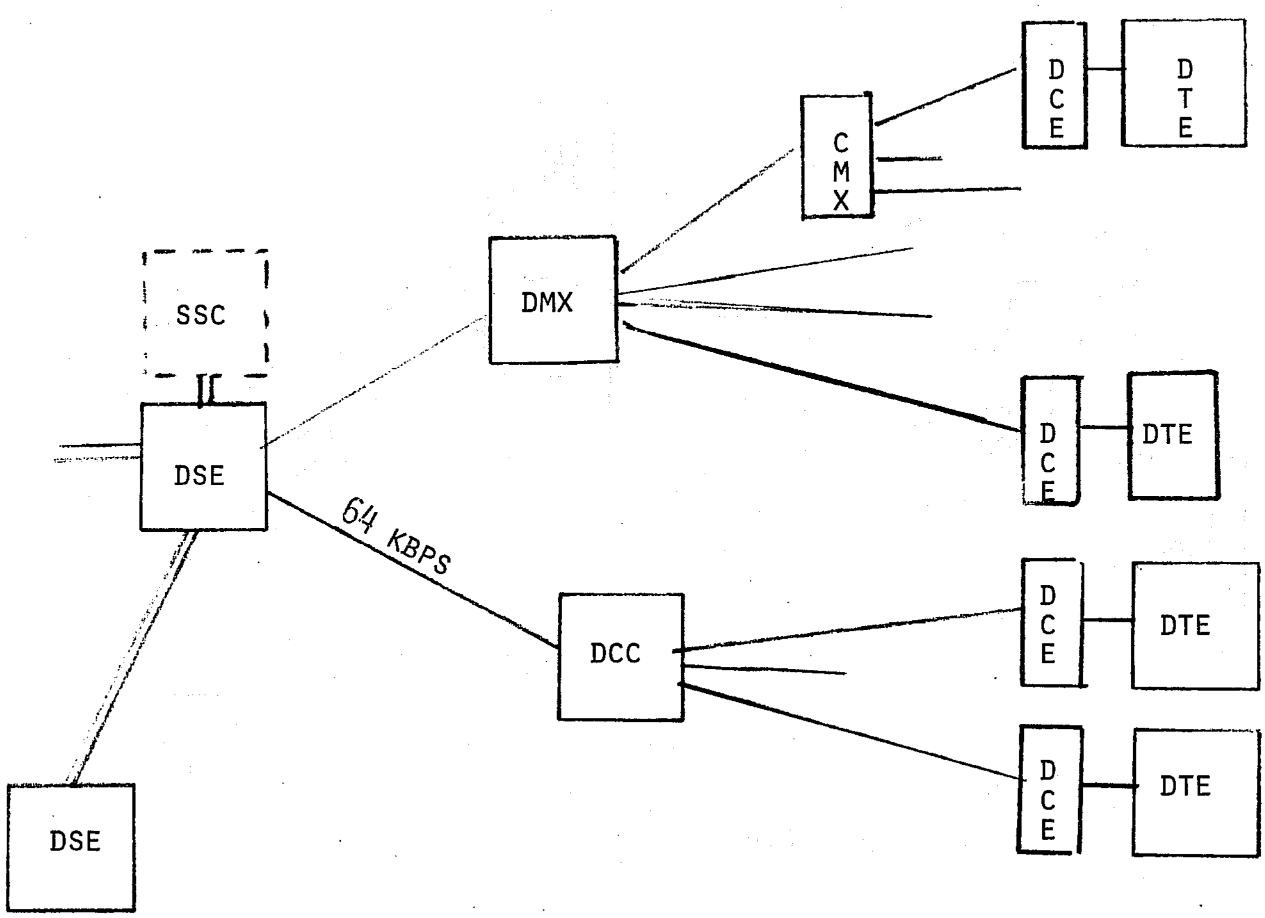

EX: NORDIC
PUBLIC
DATA NETWORK

Figure 6A : CIRCULAR NETWORKS

- POINT-TO-POINT BASED CIRCULAR TOPOLOGY
- VARIOUS MEDIA AND ACCESS TECHNIQUES
  - LOOP: CENTRAL STATION POLLS
  - RING: EMPTY SLOTS, TOKEN PASSING,...
- VARIOUS LINK CONTROL APPROACHES:

  SDLC, MINI-PACKET,...

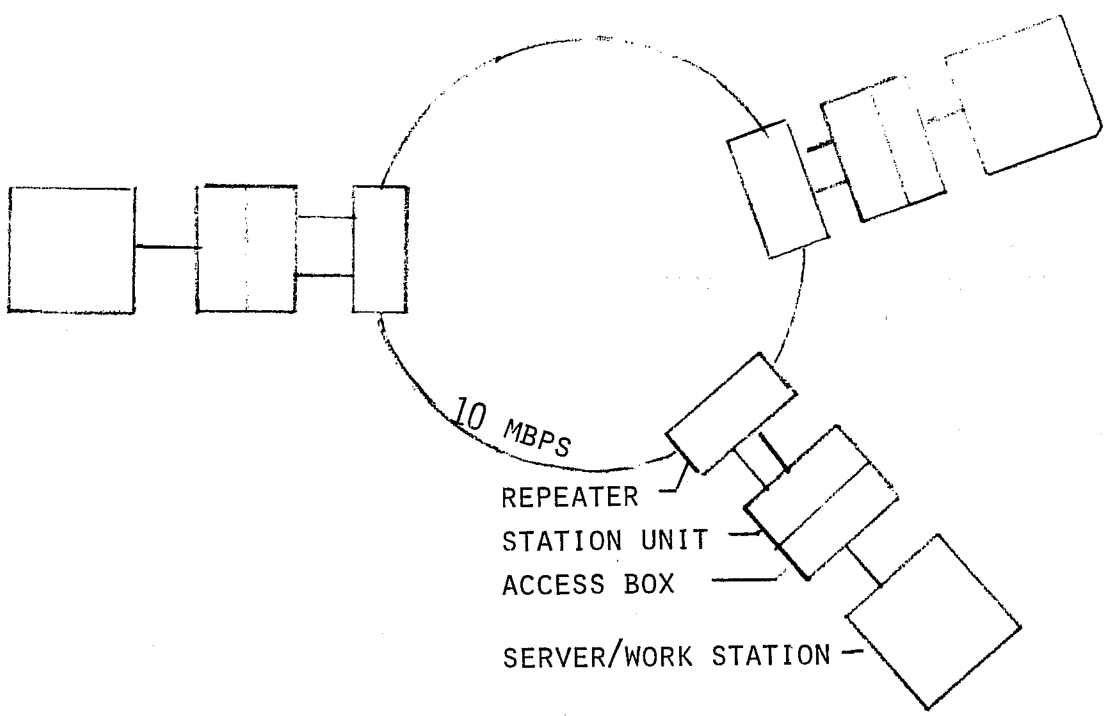

EX: CAMBRIDGE RING

Figure 6B : DATA LINK CONTROL

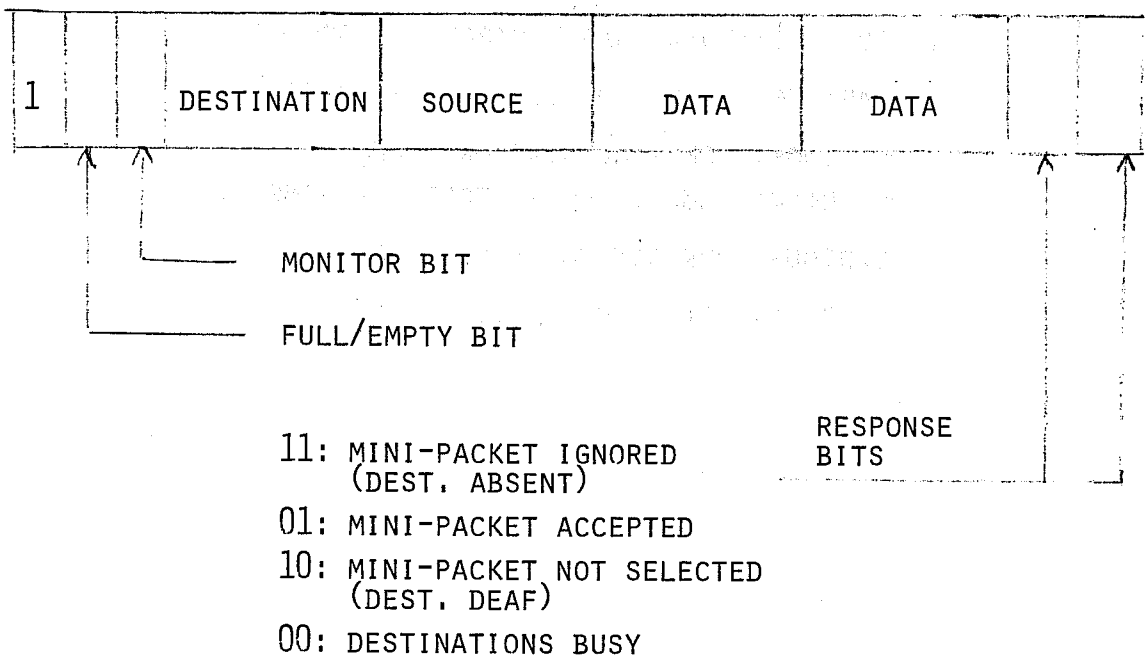

EX: CAMBRIDGE RING MINI-PACKET

Figure 7A : BROADCAST PACKET SWITCHED CONTENTION NETS

- MULTIPLE ENDPOINT PHYS. CONNECTIONS (BROADCASTING)
- VARIOUS MEDIA (WIRES, COAX, CATV, RADIO, SATELLITE...)
- CONTENTION STRATEGIES TO SHARE MEDIUM:
  - ALOHA, SLOTTED ALOHA
  - CSMA
  - CSMA/CD AND DERIVATIVES
- VERY SIMPLE LOW LEVEL PROTOCOLS (DATAGRAM)

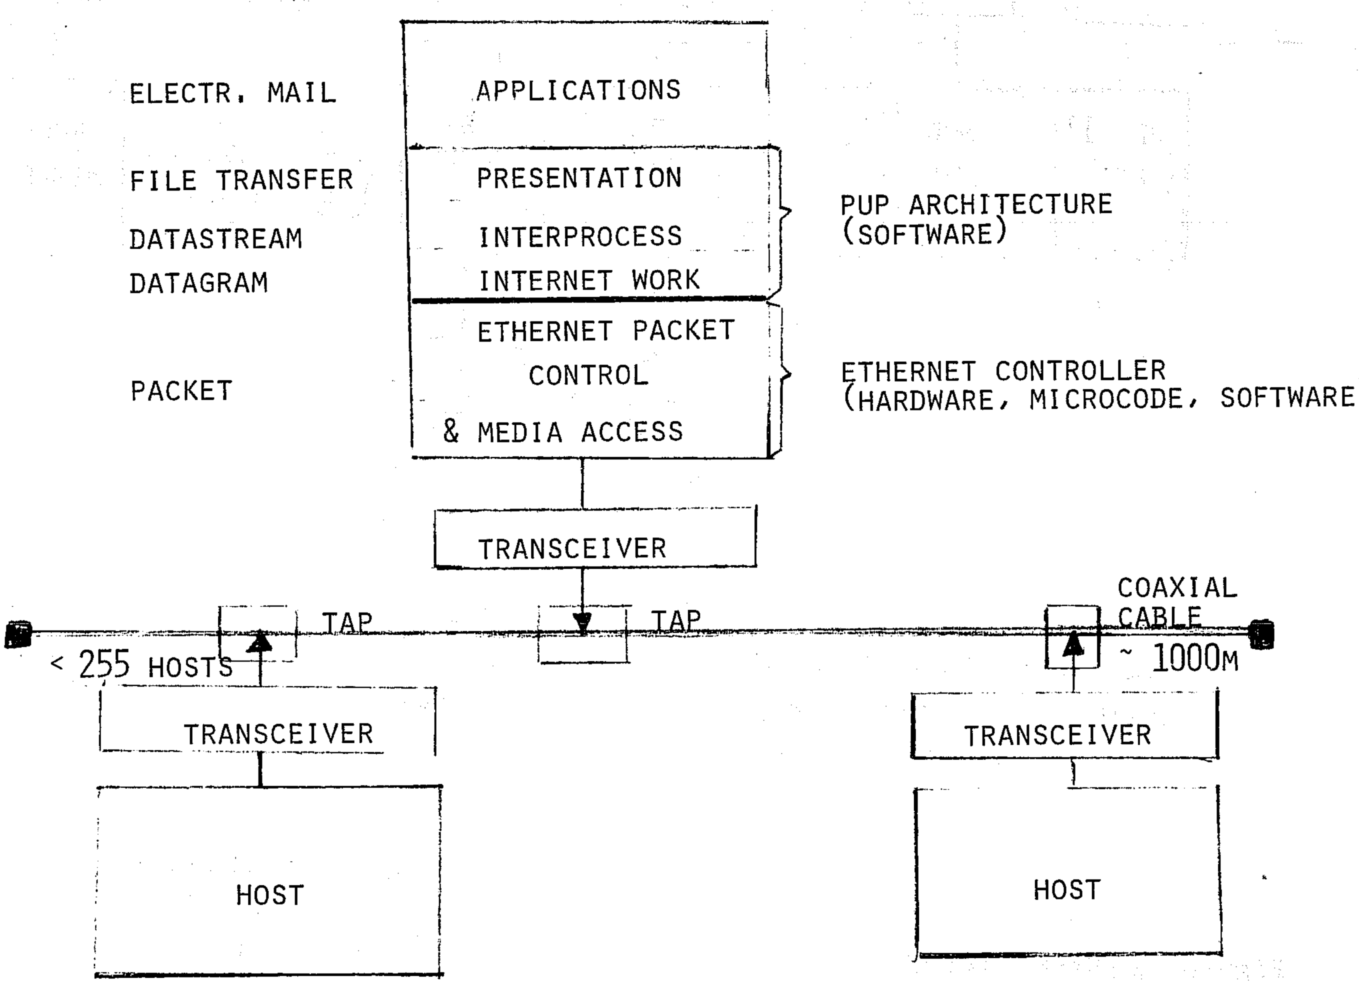

EX: ETHERNET AND PUP ARCHITECTURE

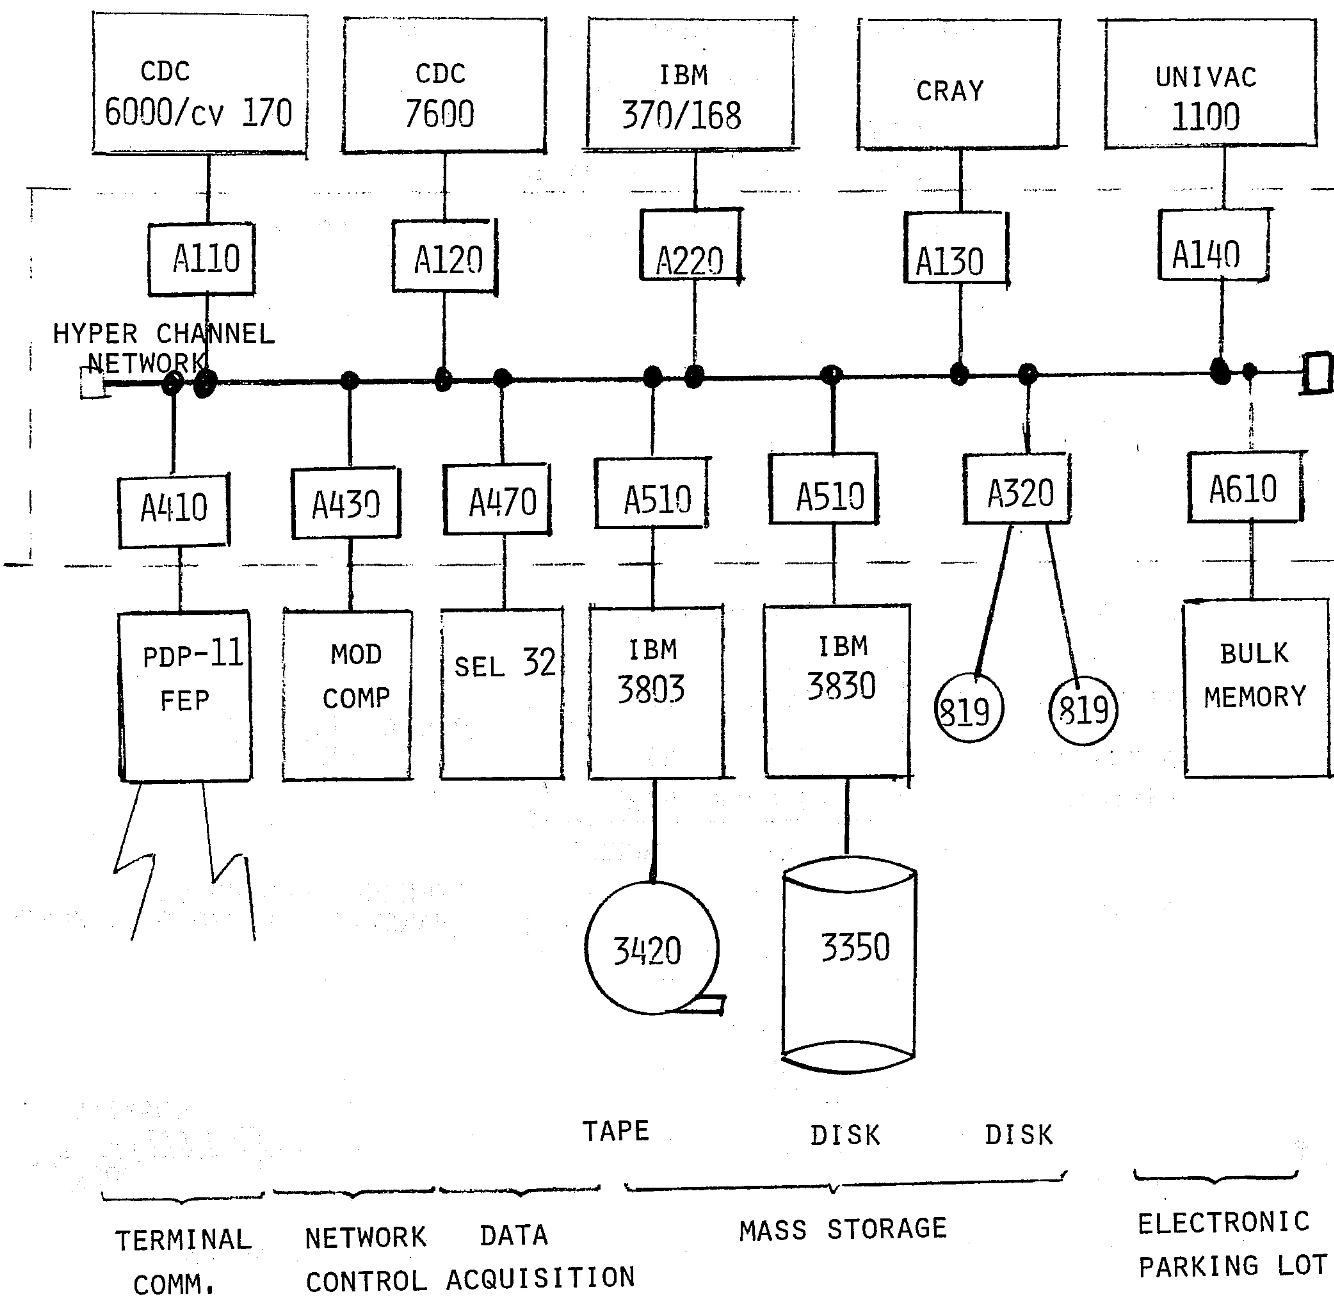

Figure 7B : HYPERCHANNEL APPLIED IN A LARGE COMPUTER SERVICES ENVIRONMENT

<u>NOTE</u> THAT HYPERCHANNEL IS ONLY THE MEDIUM AND PHYSICAL INTF. INCLUDING CONTENTION ACCESS TECHNIQUE.

Figure 8A:
ECMWF - SITUATION

EXAMPLE 1

ECMWF POSSIBLE DEVELOPMENT

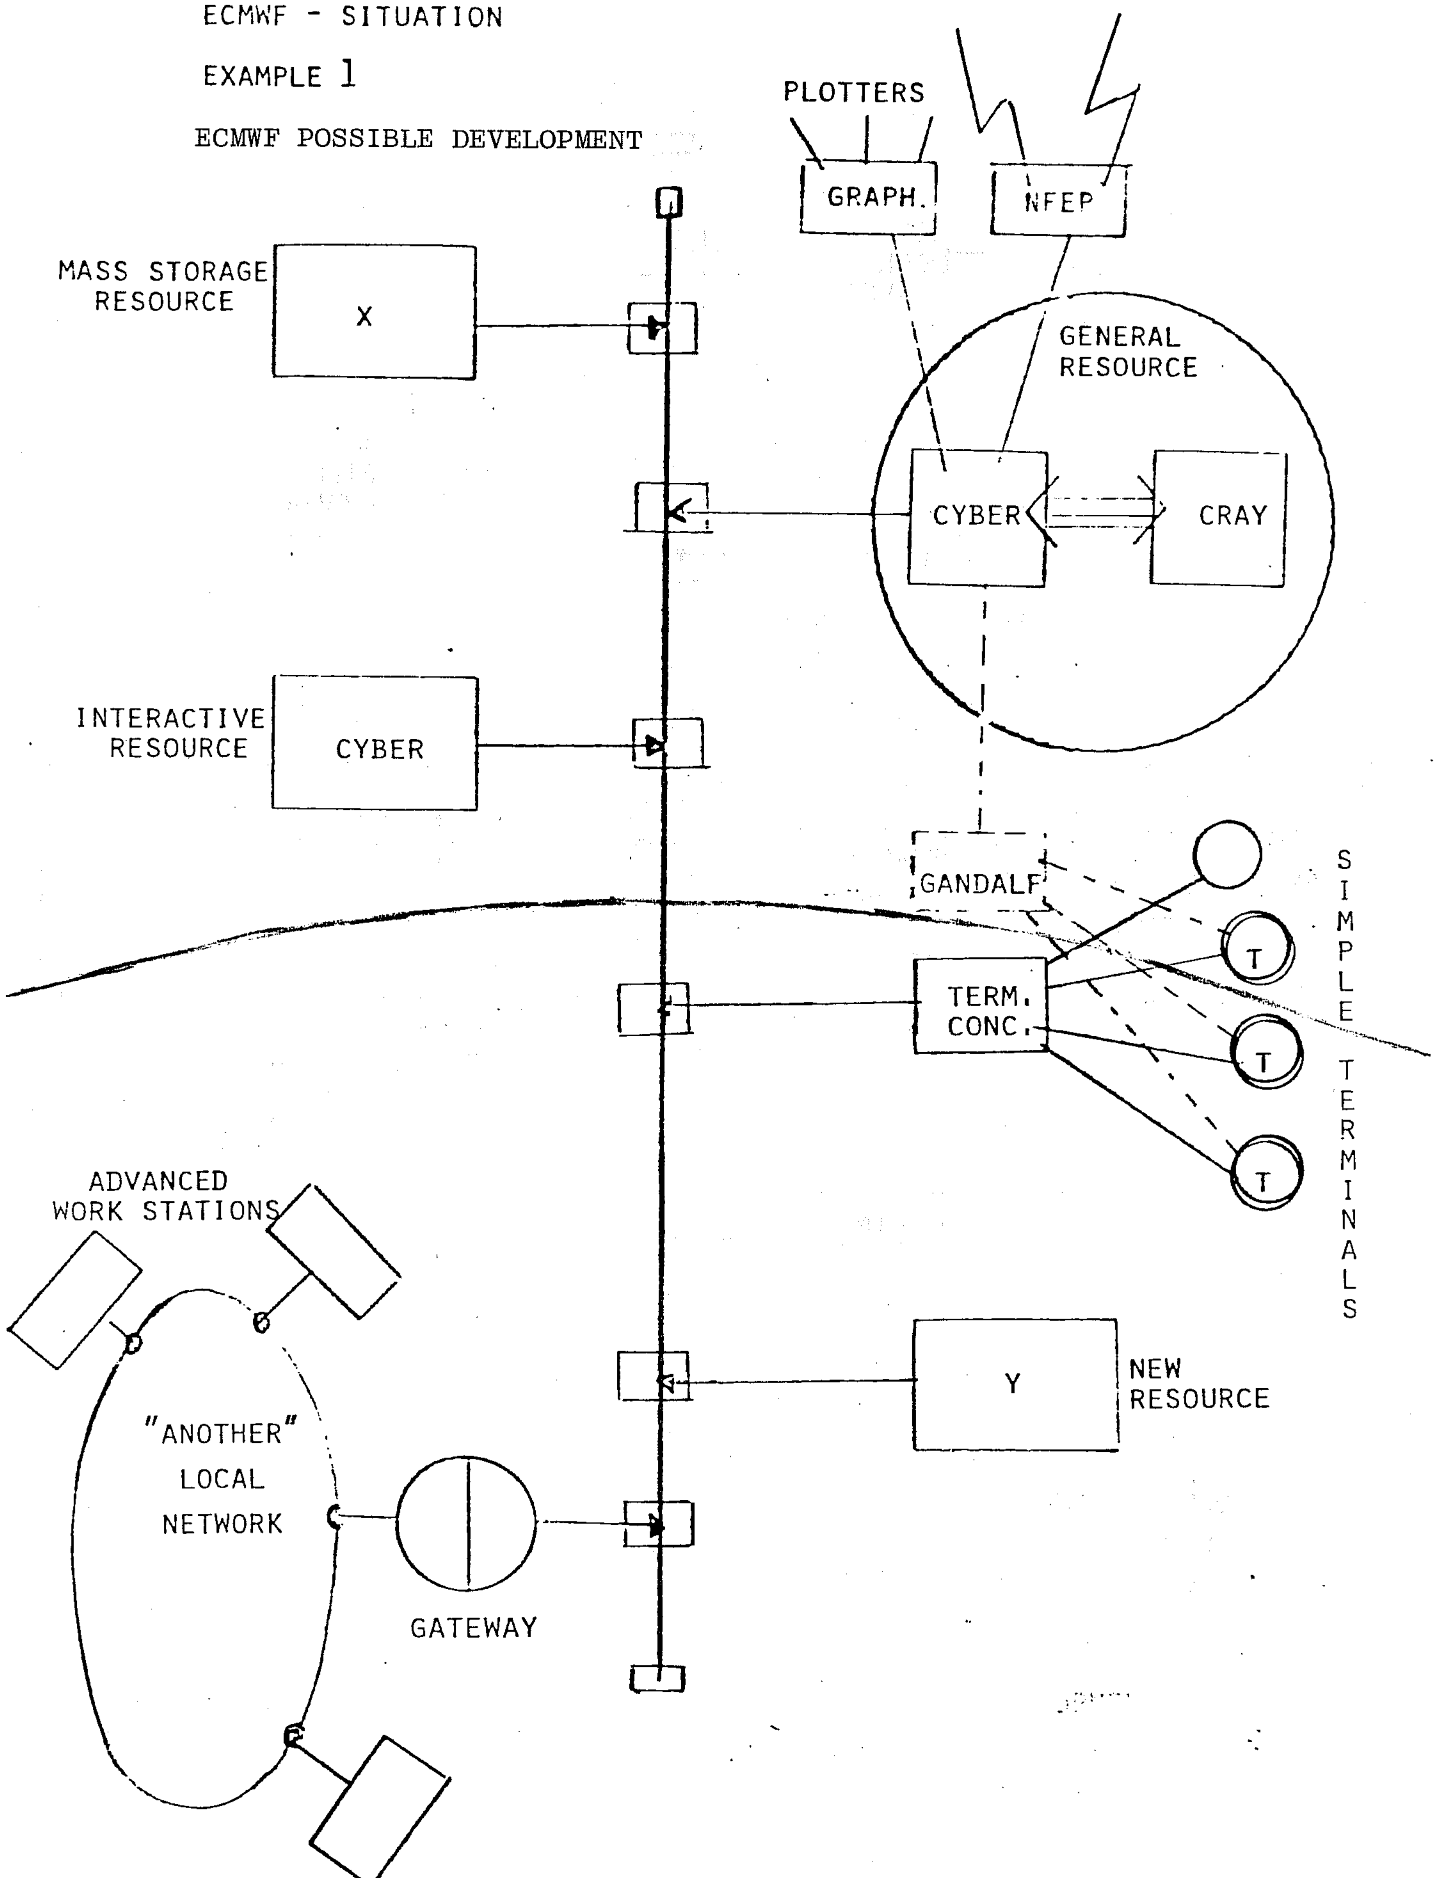

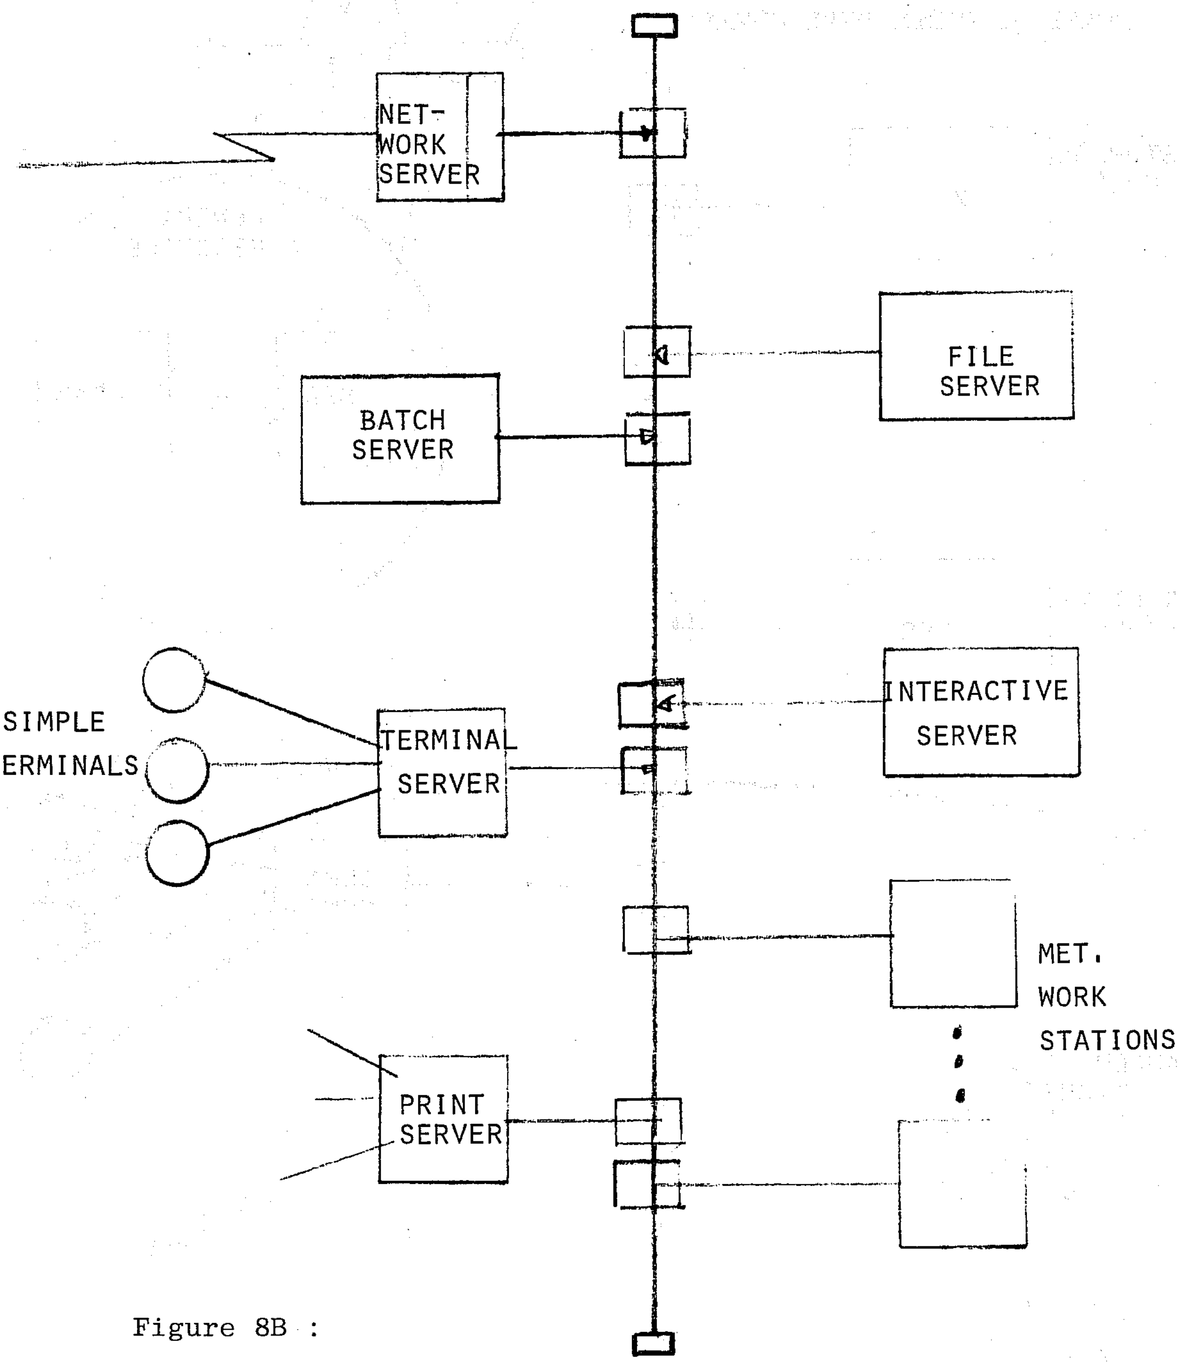

Figure 8B :

ECMWF - SITUATION
EXAMPLE 2
A WELL COMPOSED FUTURE SOLUTION

Acronyms

ISO	International Standards Organisation
NBS	National Bureau of Standards (USA)
IEEE	Institute of Electrical and Electronics Engineers
CCITT	Comite Consultatif International Telegraphie et Telephonie
ECMA	European Computer Manufacturers Association
DIX	DEC - Intel - Xerox (USA)
DEC	Digital Equipment Corporation (USA)
NSC	Network Systems Corporations (USA)
ARPA	Advanced Research Project Agency (USA)
PUP	Parc Universal Packet (USA)
UNINETT	University Network (Norway)
HMI	Hahn-Meitner Institute (Germany)
NPL	National Physical Laboratory (UK)
PABX	Private Area Branch Exchanges
CSMA/CD	Carrier Sense Multi Access/Collision Detection
HDLC	High Level Data Link Control (ISO-Standard)
SDLC	Synchronous Data Link Control (IBM)
V24) X.21) X.25) X.75)	CCITT standard interfaces and protocols
M bps	Mega bits per second
K bps	Kilo bits per second
VLSI	Very Large Scale Integration.

REFERENCES

1.	Metcalfe, Shoch	Local Computer Networks McGraw-Hill Conference 1980.
2.	Metcalfe	A strategic overview of Local Computer Networks. On line Conference on Local Networks and Office Automation, London, May 1981.
3.	Jensen	Local Area Networks. Taxonomy and architectures related to ISO OS University of Tromsø ISSN 0373-4854, May 1981
4.		Data Processing - Open Systems Interconnection Basic Reference Model ISO DP 7498
5.	Clancy, jnr.	A Status Report on the IEEE Project 802 Local Network Standard On line Conference on Local Networks and Office Automation, London, May 1981