

INET INTERCONNECTION TO STELLA

A.Endrizzi
J.Locquet
L.Lenzini
M.Mannocci

Rapporto interno C82-7

Istituto CNUCE
PISA
Gennaio 1982

STELLA/CNUCE/82/02
march 1982

INET Interconnection to STELLA

A. Endrizzi, J. Locquet
CCR EURATOM
Ispra

L. Lenzini, M. Manocci
CNR - Istituto CNUCE

1. Historical context and purposes

- a) The teleinformatics laboratory of JRC at Ispra owns a PDP 11/44. This computer is mainly aimed at being a link driving computer for a satellite earth station (Stella project).

However, at present, an antenna is not available at JRC: the price of its installation is judged prohibitive. For this reason the PDP 11/44 cannot be used as a link driving computer.

The working satellite (OTS) will be replaced by a new one. JRC is waiting for it. It is hoped that an antenna will be then available at a lower price and a true satellite link will be implemented at Ispra.

- b) CNUCE at Pisa is involved in the STELLA project and owns a PDP 11/70 which presently works as a link driving computer.
- c) Digital Equipment, the manufacturer of PDPs, has developed a X.25 interface. Software and hardware for this interface is available on both PDPs. Jrc (at present) and CNUCE (in the near future) have a connection to EURONET, the European X.25 network.
- d) The dialogue between JRC and CNUCE is therefore technically feasible across the X.25 link.

This X.25 link can be used to provide JRC with an access to the satellite network. This access cannot be thought of as an efficient link for file transfer: the reason is the differences of speeds between the satellite link (1 Mbauds) and the terrestrial link (EURONET lines leased by JRC and CNUCE have a speed of 9600 bauds). The terrestrial link constitutes a bottleneck.

However the X.25 link gives the possibility to develop and implement an internetworking experiment. Both ends of the link are presently connected (and will be more and more in the future) to existing local networks which provide a lot of different resources.

CNUCE at Pisa is a node of the STELLA network and of a CERNET network. A Cambridge ring will be installed in the near future. The IBM computers of CNUCE are also part of the Italian CNR network.

JRC at Ispra has its local X.25 network (called INET). This has a working connection to EURONET and several entries in the Amdahl mainframe computer. Several PDPs are spread over the Centre and some of them are interconnected. A X.25 microprocessor network has been brought and is to be installed. An optic fibre ring is under development.

The aim of the connection is double:

- provide the subscribers of both ends with an easy and quick access to resources of the other end,
- develop, implement and test an internetworking protocol according to the O.S.I. model.

2. General overview of INET: the Ispra network

a) Hardware configuration (see Figure 1)

Six Solar 16/65 minicomputers (16-bit processors from SEMS, France) form a mesh network. They are internally connected by HDLC line drivers (speed up to 48000 bauds).

The Solars are at the same time network processors (packet switching nodes) and host computers offering services to users.

Asynchronous terminals are connected to the various computers and are managed according to the ESP20 protocol which is software implemented in each Solar.

The network is internally organized as a X.25 like network but provides the users (local tasks on Solars, or tasks on communicating external computers) with a true X.25 interface. Any additional computer can be connected to the network provided that it offers a X.25 interface (i.e. the PDP 11/44).

b) Connections to the environment

INET also provides the users with

- TSO entries in the Amdahl mainframe at JRC;
- RJE entries in the Amdahl mainframe;
- a gateway to EURONET across a dedicated line at 9600 bauds and 16 SVC;
- a connection to the PDP 11/44. This is viewed by INET as a subscriber and is also connected to the PDP 11/70 of the teledetection laboratory.

c) Services offered by INET

An interactive user (local or remote) may

- gain access to a particular Solar specifying its physical address and then ask for a service provided by this Solar;
- ask for one of the services offered by the various

servers of INET without specifying a physical address. INET searches for the available computers which are able to give the requested service and chooses the best suited one.

The presently offered services are:

- "TSP": time sharing system of Solar; available languages are Assembler, PL 16, Basic, Fortran and APL.
- "TSO": time sharing system of the Amdahl mainframe.
- a remote user can test its own implementation of the NIFTP protocol (Network Independent File Transfer Protocol).
- "ECDIN" and "EURPI" are two accesses to data bases of the Amdahl computer. These data bases are parts of DIANE (Direct Information Access Network for Europe).
- A file transfer between INET nodes or any other connected subscriber using the same FTP.
- A remote job entry in the Amdahl mainframe.
- Synchronous terminals (IBM 3270) could be managed by the DEVT protocol and used to communicate with IMS of the Amdahl computer.

3. Internetworking

a) Theoretical approach

The purpose of internetworking is to allow communications between subscribers from networks of different architectures.

To do this, it is required to:

- build gateways between the networks (a network of networks is created);
- implement in each network an intermediate layer of internetworking (layer 3B situated between layer 3A (network) and layer 4 (transport)).

An internetwork packet is composed of

- the data field;
- the internet header;
- the local header.

In a particular local network, the internet header is

considered as data. The local header is peculiar to each crossed network.

The internet header contains the information which is useful for the gateway especially for routing.

The addressing scheme is local net address, local address in the destination network.

In any case, gateways have to be implemented. However this approach has a drawback: the introduction of a new layer implies the rewriting of existing software in each local network.

b) CERNET approach

Two CERNETS and the STELLA satellite network have been interconnected. CERNET is a datagram network and has an efficient transport manager.

The CERNET designers did not permit changes to their existing software.

Instead of introducing an intermediate layer inside the network software, it was decided that applications (running above the transport manager) should use a free byte of the CERNET packet. This byte is used for addressing.

This solution provides a small addressing space but no changes are made to existing software.

c) INET approach

NOTE:

In the following lines EURONET is not considered as a local network but is only a means to establish a physical line between Pisa and Ispra. The PDP 11/44 at JRC is considered as a normal subscriber of INET.

End of note.

c.1 First gateway between EURONET and INET can be used to manage the internet flow. It must recognize the incoming (to INET) internet packets and route them to the destination subscriber in INET. The outgoing (from INET) internet packets must be recognized; all of them must be routed to Pisa through a single X.25 virtual call.

An internet layer (3B) must be implemented in each subscriber of INET (Solars and PDPs).

c.2 Second alternative (see Figure 3)

- 1- A "permanent" X.25 virtual circuit is established between the PDP 11/44 of INET and the PDP 11/70 of Pisa.

NOTE :

In reality there is a virtual circuit between the PDP 11/44 and the EURONET gateway across INET and another one between the EURONET gateway and the PDP 11/70 across EURONET).

End of note.

- 2- When an INET subscriber wants to communicate with an external correspondent (reachable from Pisa), he must address the local subscriber PDP 11/44. There will be a local convention inside INET for addressing an external subscriber: the address is composed of the INET address of the PDP 11/44 followed by a number identifying the remote subscriber.
- 3- The bulk of internetworking is under the responsibility of the PDP 11/44. This manages the various virtual calls which are addressed to itself, recognizes the destination subscriber (performs an address translation), builds the internet header and sends the packets to the PDP 11/70 through the X.25 channel.

When the PDP 11/44 receives internetworking packets coming from the PDP 11/70 at Pisa, it analyses the internet header, retrieves the corresponding virtual circuit in INET, strips the internet header and sends the packet to the INET subscriber.

- 4- At the other end of the X.25 channel, the PDP 11/70 plays the same role. Incoming packets are routed to the destination subscriber of STELLA according to the contents of the internet header. Outgoing packets which are destined to an INET subscriber are sent to the PDP 11/44 at Ispra through the X.25 channel.

The PDP 11/44 at Ispra is the first gateway half of the interconnection between INET and STELLA.

The PDP 11/70 at Pisa is the second gateway half of this interconnection. The X.25 channel can be seen as a wire linking these two gateway halves.

NOTE :

If we consider that several networks are reachable at Pisa, the situation becomes a bit more

complicated: (STELLA but also CERNET and soon a Cambridge Ring).

One gateway half must be implemented for each network and these gateway halves must be interfaced to the X.25 channel by a special task: its function is to route incoming packets from Ispra to the appropriate gateway half according to the information contained in the internet header.

End of note.

c.3 Comparison of the two alternatives

This second alternative is preferred because it presents a great advantage: no changes must be made to the existing software of INET. (Same constrain as for CERNET). The new software which is required for internetworking will be developed on both PDPs.

However there is a loss of efficiency: the internet traffic in INET is doubled. An INET subscriber must first send his packets to the PDP 11/44 and this resends the packets which cross INET again.

Logically the link between the two gateway halves (PDP 11/44 at Ispra and PDP 11/70 at Pisa) can be seen as external to INET (see Figure 4). This vision changes nothing to the logical design of the interconnection.

If the internet traffic is low, the loss of efficiency has little importance. If it becomes high, then EURONET leased lines will also be saturated and an increase of its capacity will be required.

If necessary a new line connected to EURONET should be connected directly to the PDP 11/44. This will alleviate the internal traffic of INET and the physical connections will be the same as the logical connections.

4. Additional problems

a) Datagrams and virtual circuits

The gateway half at JRC will have to manage the interconnection between virtual circuits and datagrams and it must be transparent to both senders and receivers.

INET has no datagram possibilities: so a virtual circuit must be opened. The other networks which are reachable at Pisa may be datagram networks. Between the two gateway halves no virtual circuits are really used. All the

communications between the two gateway halves are multiplexed in the same X.25 channel.

What should be the policy of the PDP 11/44 for a communication between an INET subscriber and a STELLA subscriber for example?

If the gateway half receives a request to open a virtual circuit from the INET subscriber, it must accept the opening of this circuit immediately and no message will be sent to the correspondent. When the first data packet is sent, the gateway half sends it to the receiver as a datagram.

In the other direction, if the gateway half at Ispra receives a datagram destined to an INET subscriber it must first verify if a virtual circuit (between the two correspondents) is already opened. If yes it sends the packet in this virtual circuit. If not it must first open the virtual circuit and when this one is properly opened, send the datagram packet through the created virtual circuit.

If the opening of the virtual circuit is not accepted by the INET subscriber, the datagram packet is dropped.

When should the closing of the virtual circuit between an INET subscriber and the gateway half occur?

If both ends of the communication are allowed to close the virtual circuit, it results that special datagram packets should be sent and recognized by the gateway half in order to close the circuit.

This special mechanism should not be implemented because it requires a dialogue between a subscriber and the gateway half which must be transparent. The users may ignore the presence and implementation details of the gateway.

The result is that the closing of the virtual circuit should be made only by the INET subscriber.

b) Flow control

When the gateway receives a packet coming from INET, it may:

- hold the packet for retransmission to the other side and, when this packet is properly sent, send an acknowledge to INET.

This solution is dangerous because (according to the traffic) the packet may stay in the gateway for a long time and may reach the destination when the receiver is no more waiting for it.

- send immediately an acknowledge to INET.

The packet will be retransmitted if possible and dropped if not. This second solution is preferred because it also avoids the saturation of the gateway.

It results that the flow control mechanism must be implemented end to end by transport managers. Their task is to retransmit packets if a time out has expired (that means that the packet has been dropped somewhere) and to send an acknowledge to the conversing transport manager when a packet is properly received.

c) Packet fragmentation

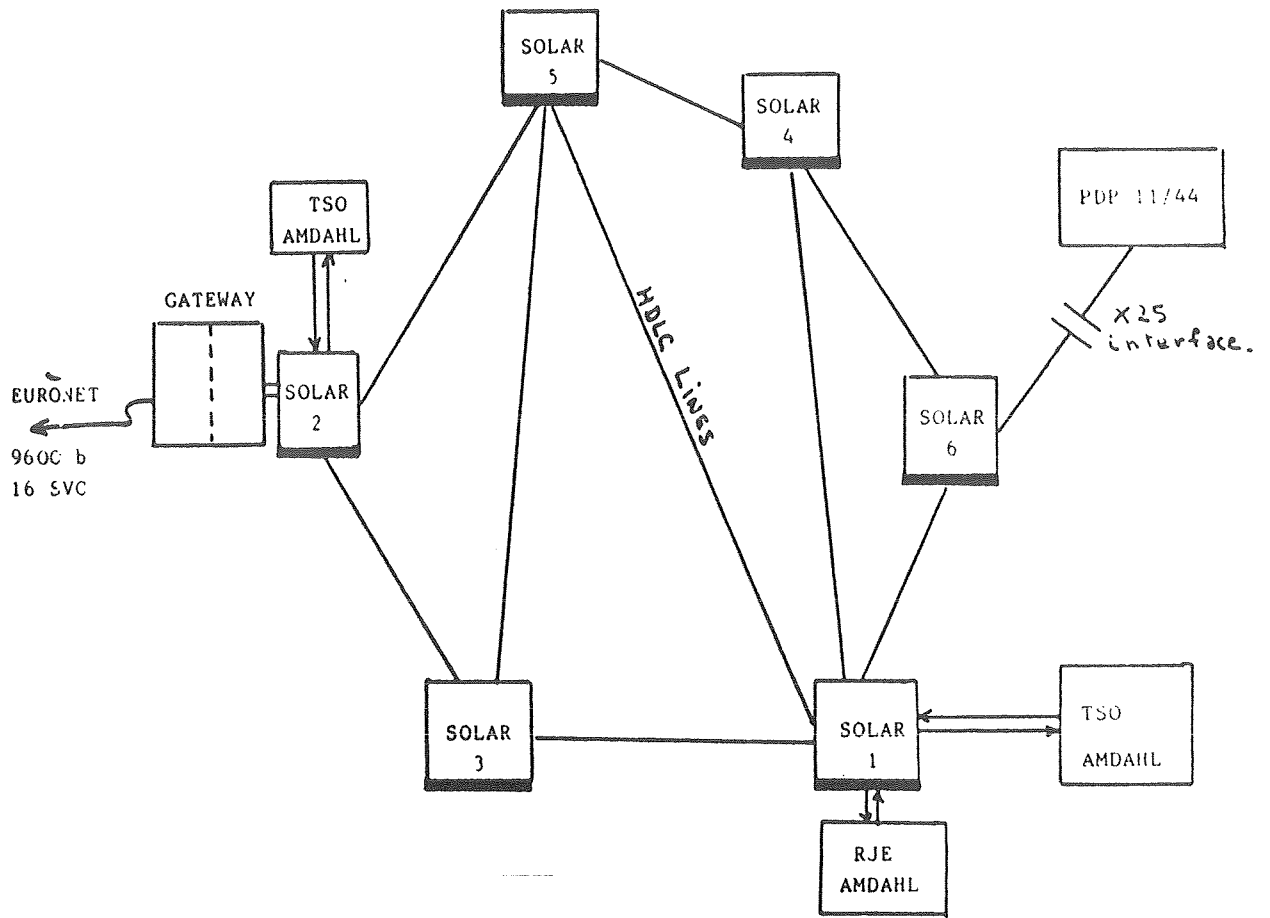
STELLA packets are much bigger than X.25 packets: this implies that a packet fragmentation must take place at the Pisa gateway half for packets coming from a satellite subscriber.

c.1 A single packet from STELLA will create a lot of packets that must cross the X.25 channel between the two gateway halves. There is the danger of hogging this X.25 channel for a long time: the flow control of this single virtual circuit is global and does not distinguish between the various communications.

It is therefore necessary to implement a round Robin mechanism between the various communications which must share the single X.25 channel and to implement a selective flow control mechanism between both ends of this X.25 channel.

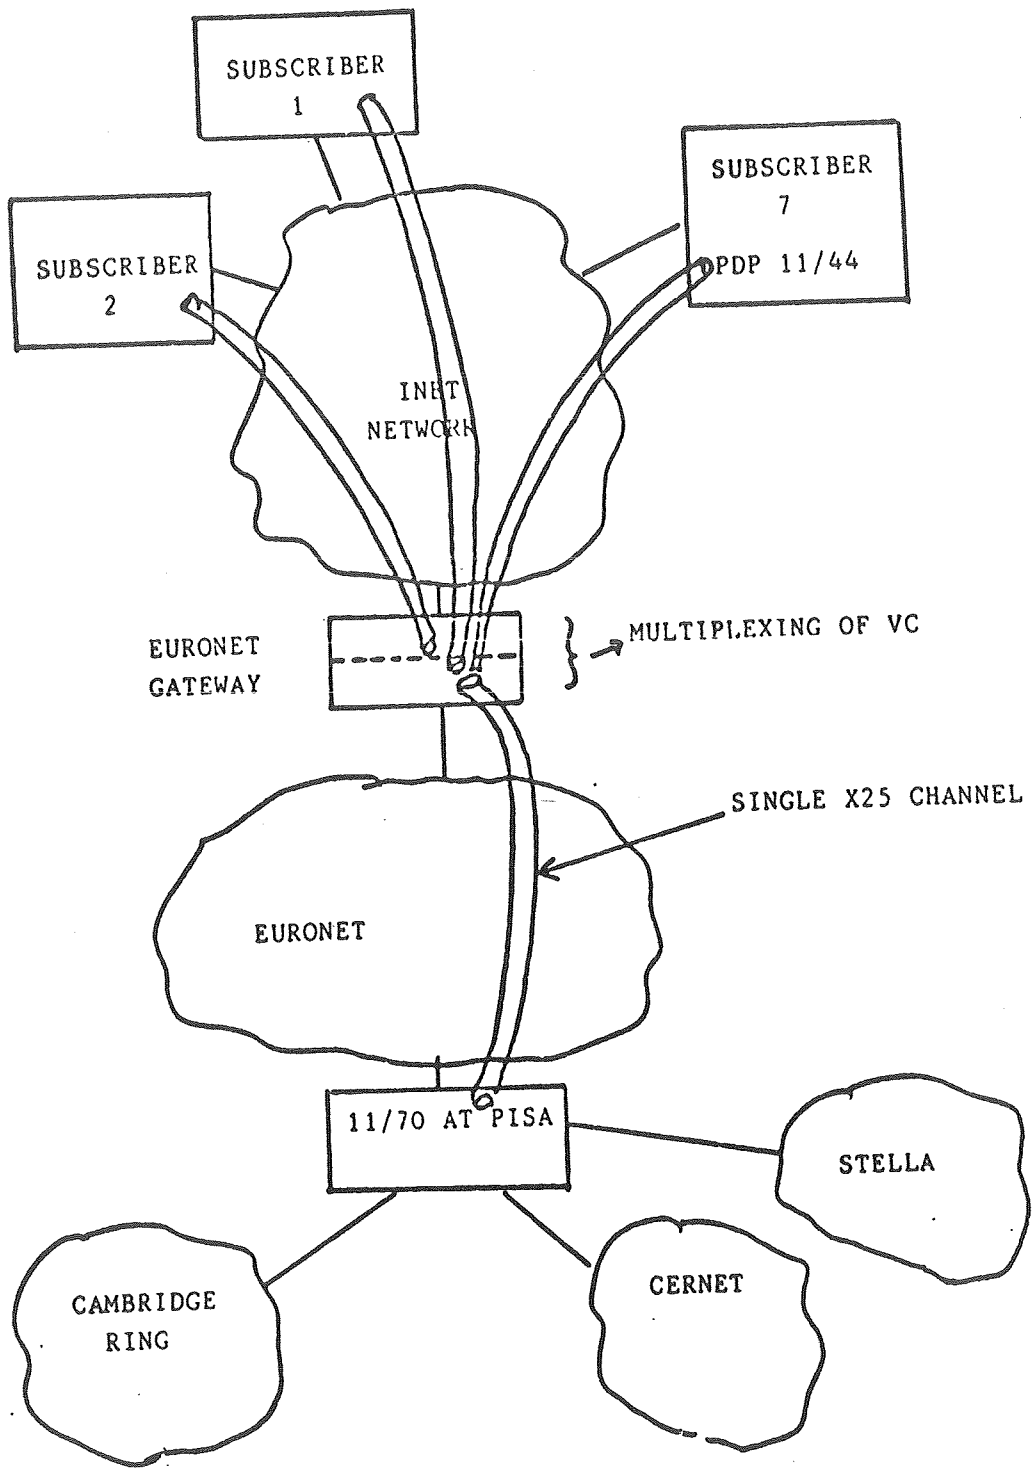
c.2 In the other direction, small packets will be sent from INET to a STELLA subscriber. Three policies are possible:

- When the Stella gateway half receives a packet coming from Inet it sends it immediately in a big STELLA packet. This solution wastes the possibilities of the satellite link.
- The STELLA gateway half buffers the packets and waits for a particular packet which indicates the end of the X.25 message. This implies that the INET task must be written according to the STELLA peculiarities.
- The best solution is to introduce a time out mechanism. When the STELLA gateway half receives packets from INET it sets a time out and buffers the packets. A STELLA packet will be sent if the cumulative size of the buffered packets has reached the size of the STELLA packet or if the time out has expired between two successive packets. The expiration of the time out means generally that the sender has no more packets to send.



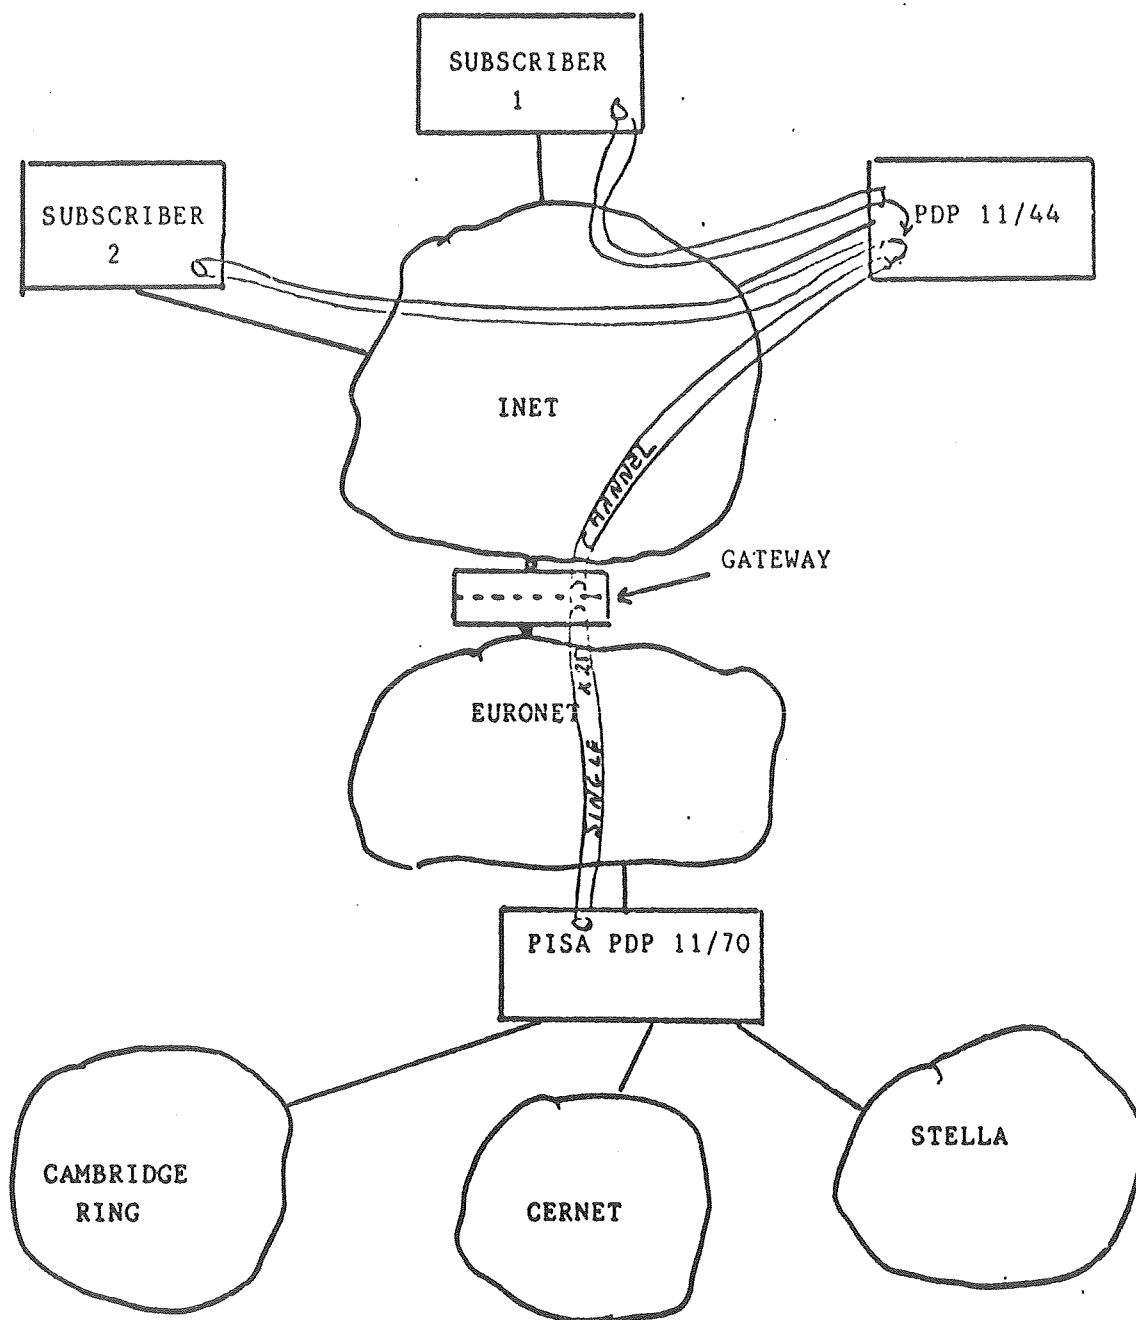
(SOLARS = NODES + HOSTS)
 INET TOPOLOGY (AT ISPRA)
 =====

Figure 1




FIRST ALTERNATIVE

Figure 2



SECOND ALTERNATIVE

Figure 3



LOGICAL REPRESENTATION OF THE SECOND ALTERNATIVE

Figure 4