ISTITUTO NAZIONALE DI FISICA NUCLEARE - ISTITUTO NAZIONALE DI FISICA NUCLEARE - ISTITUTO NAZIONALE DI FISICA NUCLEARE - ISTITUTO NAZIONALE DI FISICA NUCLEARE - ISTITUTO NAZIONALE DI

Laboratori Nazionali di Frascati

LNF-88/31(R)
3 Giugno 1988

G. Di Pirro, L. Trasatti:

ETHERFAR: A LONG DISTANCE FIBER OPTIC ETHERNET CONNECTION

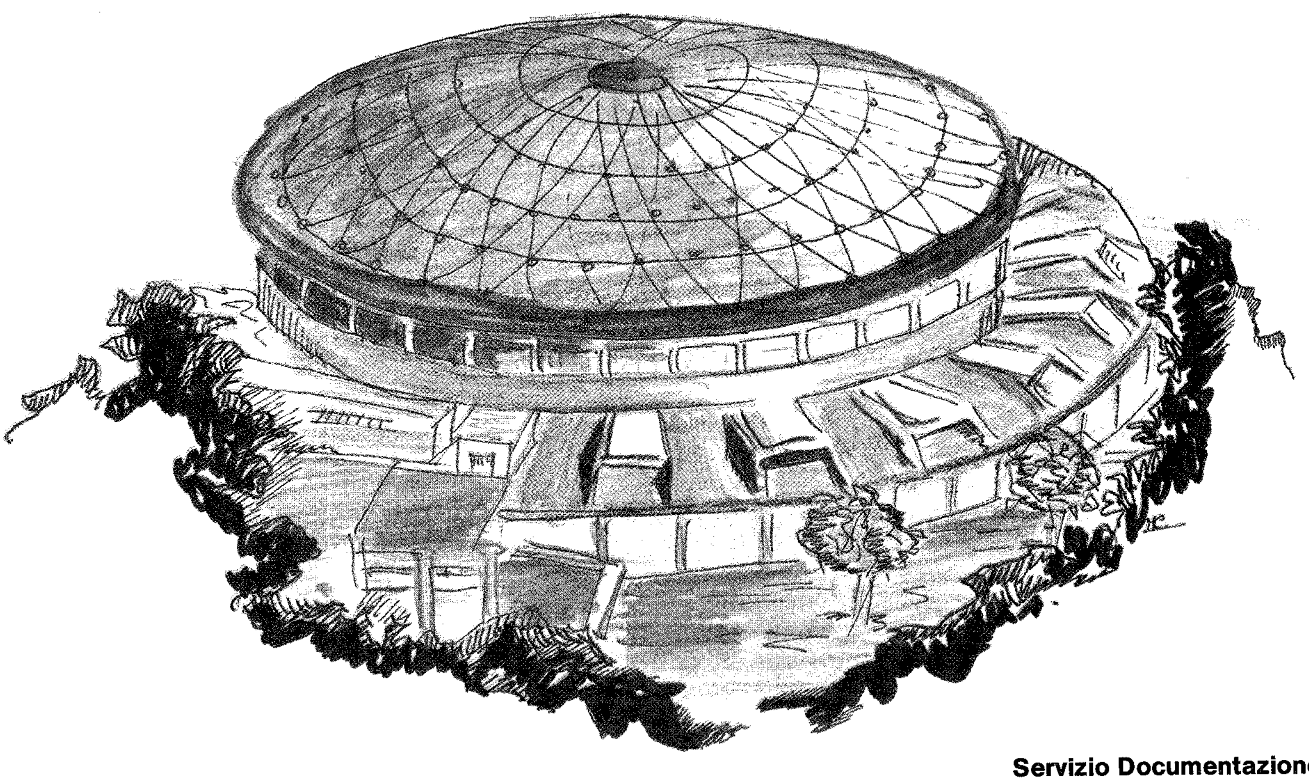

Servizio Documentazione
dei Laboratori Nazionali di Frascati
P.O. Box, 13 - 00044 Frascati (Italy)

INFN - Laboratori Nazionali di Frascati
Servizio Documentazione

LNF-88/31(P)
3 Giugno 1988

ETHERFAR: A LONG DISTANCE FIBER OPTIC ETHERNET CONNECTION

G. Di Pirro, L. Trasatti
INFN - Laboratori Nazionali di Frascati, P.O.Box 13, 00044 Frascati, (Italy)

Abstract

We have connected two ETHERNET networks at a distance of 10 Km using two DEC Bridges and an optical link with single mode optical fibers. Tests on packet losses have been carried out using two traffic generators and a traffic analyzer, without attempting to use DECNET software. Although the installation violates ETHERNET distance limits, the losses seem to be relatively small.

1- Introduction

One of the limitations in the use of the ETHERNET local area network is the maximum distance spanned: DEC limits the distance between two nodes of the same network to 2.8 Km, and the distance covered by an extended LAN (more than one LAN segment interconnected by remote bridges) to 22 Km. However, there is an additional limitation: two bridges (DEC fiber optic LAN Bridge 100) cannot be further away than 2 Km. The theoretical maximum distance is based on the use of 7 bridges.

In a score of possible applications, the need to have powered apparatus scattered every 2 Km is a definite drawback. For instance, in the Laboratori Nazionali del Gran Sasso, it would be very convenient to have ETHERNET links between the experiments under the gallery and the outside laboratories and counting rooms, but the distance is 8 Km, and it is not too easy to install and maintain powered LAN bridges scattered under the tunnel.

The 2 Km limitation in the distance between two bridges is due to more than one reason:

- The use of multimode optical fibers limits the optical transmission distance to 2 Km without repeaters:
- Even if monomode fibers and appropriate Tx-Rx pairs are employed, the distance is still limited to less than 5 Km because the two bridges employ on the optical link an ETHERNET protocol, and, although the fiber optic connection is full-duplex (two fibers, one per direction of transmission), still the bridges are not capable to work in a full-duplex configuration, due to limitations of the LANCE IC's used to interface to ETHERNET. The LANCE chip set is not capable to receive while it is transmitting (which would normally be useless in an ETHERNET cable segment). Therefore, a bridge which has sent a packet on a long optical line may not be able to get in time the announcement of a collision, and may believe erroneously that his packet has been delivered without problems.

On the other end, the ETHERNET protocol itself only guarantees that a packet will pass under the nose (cable interface) of a station, and not that it will be received by the station: if the buffers are full, the packet is anyway lost. Moreover, a bridge does not guarantee that a packet will be delivered across it: again, if the traffic is too heavy the packet can be lost due to buffer filling.

The task to recover these errors is laid on the upper level layers of the protocol. Every application must therefore be capable of recovering lost packets (For instance, the VMS COPY utility sends three 1500 byte packets during file transfer, and than awaits for an acknowledge from the receiving station).

Against the advantages of a longer network one must weigh the disadvantages of an additional packet loss: the difference is quantitative and not qualitative.

Therefore we decided it would be interesting to try out a network using two DEC LAN Bridge 100 and a 10 Km fiber optic link and to monitor the loss of packets under different conditions of network traffic and packet length.

2- Experimental set-up

A diagram of the set up is shown in Fig. 1.

- LNF ETHERNET network is the DEC network of the Laboratori Nazionali di Frascati
- DELNI is the DEC ETHERNET concentrator, which on the upper side acts as a coupler to the rest of the network, while on the lower side simulates an ETHERNET segment.
- VICE (VME Intelligent Controller for ETHERNET) is a home-built, multiuser intelligent VME board interface for ETHERNET, which in this case acts as a traffic generator (Technical description to be published as LNF Internal Report).
- DEC LAN Bridge 100 is a local DEC ETHERNET bridge, without fiber optic interface. It connects to the optical system through the standard ETHERNET tranceiver connector (ECL signals).
- Fiber optic Tx-Rx is a system from General Optronic, installed by Pirelli, modified by us. The system was originally designed to operate in a point to point configuration, and therefore assumed that a carrier was always present. The amplification was regulated by an AGC circuit. This is no good

for an ETHERNET carrierless network, because when no packet is being received, the gain goes up to maximum, thus amplifying the noise on the network up to detectable levels. We removed the AGC circuit and replaced it with a stable control

The optical fibers, supplied by Pirelli SpA, are single mode, with an attenuation better than .4 dB/Km.

MVME 330 is a MOTOROLA VME-ETHERNET interface board, again used here as traffic generator.

- Lanalyzer is an ETHERNET analyzer by Excelan.

Comparison measurements were taken substituting the fiber optic link with a short cable between the bridges, and also eliminating one of the two bridges.

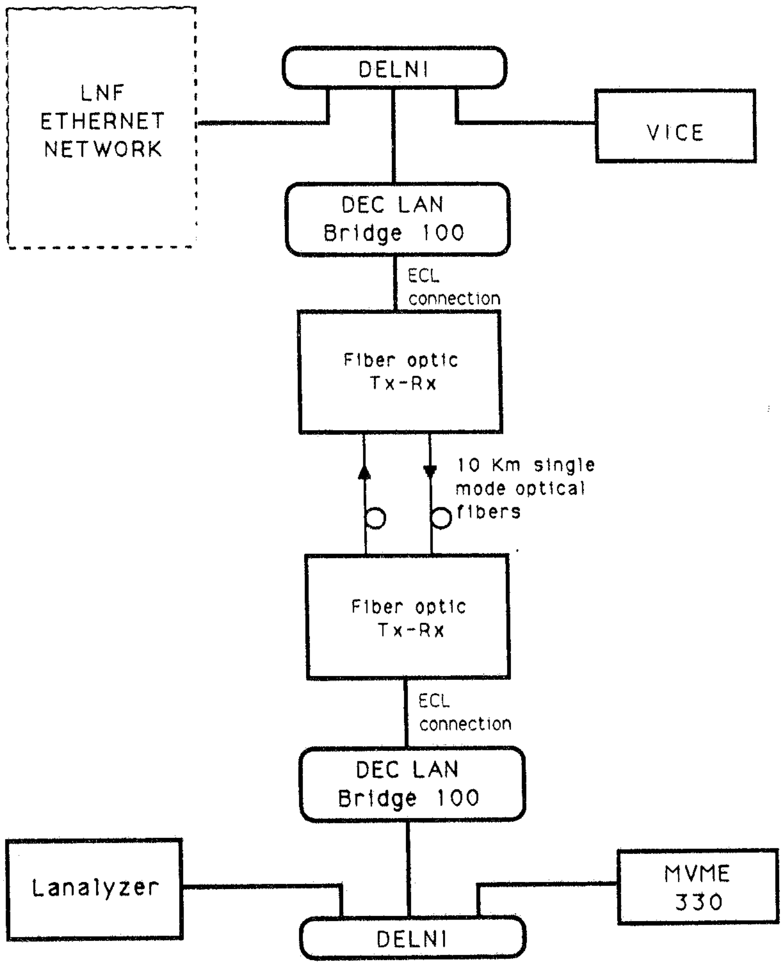

FIG. 1 - ETHERFAR experimental set-up.

3 - Measurements

Measurements were taken sending a constant stream of traffic from the MVME330, then sending from the VICE a known number of packets at different rates and measuring with the Lanalyzer the number of VICE packets delivered across the bridge and the total traffic. Two packet sizes were used: 1500 Bytes and 64 bytes.

Fig. 2a,b,c shows the percentage of packets lost versus "theoretical traffic" generated by the VICE for different values of traffic generated by the MVME330 and 1500 byte packets. "Theoretical traffic" means the percentage of the bandwidth (10 Mbit/s) which the generator would have occupied without opposite traffic.

Fig. 3a,b,c shows the same measurements, with the ascissa corresponding now to 'experimental traffic", i.e. what was experimentally measured by the traffic analyzer.

Fig. 4a,b,c, links together the two preceding figures showing experimental traffic vs theoretical traffic.

Note that the measurements were extended up to quite unrealistic traffic figures (about 90% of the bandwidth).

The first thing which appears from the figures is that, although the losses are relatively small, the graphs have a highly irregular shape. This is due to the non statistical properties of the process, where two clocked generators interact constructively or destructively. To study this behavior, we looked for the worst possible combination of the two traffic streams and performed a very fine scanning. The results are plotted in Fig. 5: the two graphs show the same scanning with different horizontal scale. Note on the lower graph that a very complicated structure (3 peaks) can be seen in a traffic interval ranging only from 27% to 28.5%.
Also note on the upper graph that, changing the MVME330 traffic by only 1.1%, the losses go from 80% to practically 0. All of the measurements reported did not suffer from statistical errors and where highly repeatable. The conclusion is that, although particular situations may occur where losses are very heavy, their relevance in a general average is not very high.

Fig. 6 shows the data in a very condensed form, plotting the percentage of packets lost vs the sum of the theoretical traffic in both directions. Again, for reasonable traffic figures, the average loss is of the order of 2%.

Some of the measurements were repeated substituting for the fiber optic link a very short electric connection. They show that the losses in this case are substantially lower than for the 10 Km case, but not completely absent. Fig. 7 shows the results for three cases.

Fig. 8 shows the measurements for 64 byte packets in the same conditions. Here the number of packets is plotted instead of the percent bandwidth utilization, because this is the relevant parameter when the packets are very small.

4 - CONCLUSIONS

Although the losses are much higher with a 10 Km link than with a short one, the absolute value is small, and the link seems serviceable.

Further measurements should be carried out on the performance of standard DEC software products to confirm what we have found here. We feel, however, that these measurements are a necessary prerequisite for further tests.

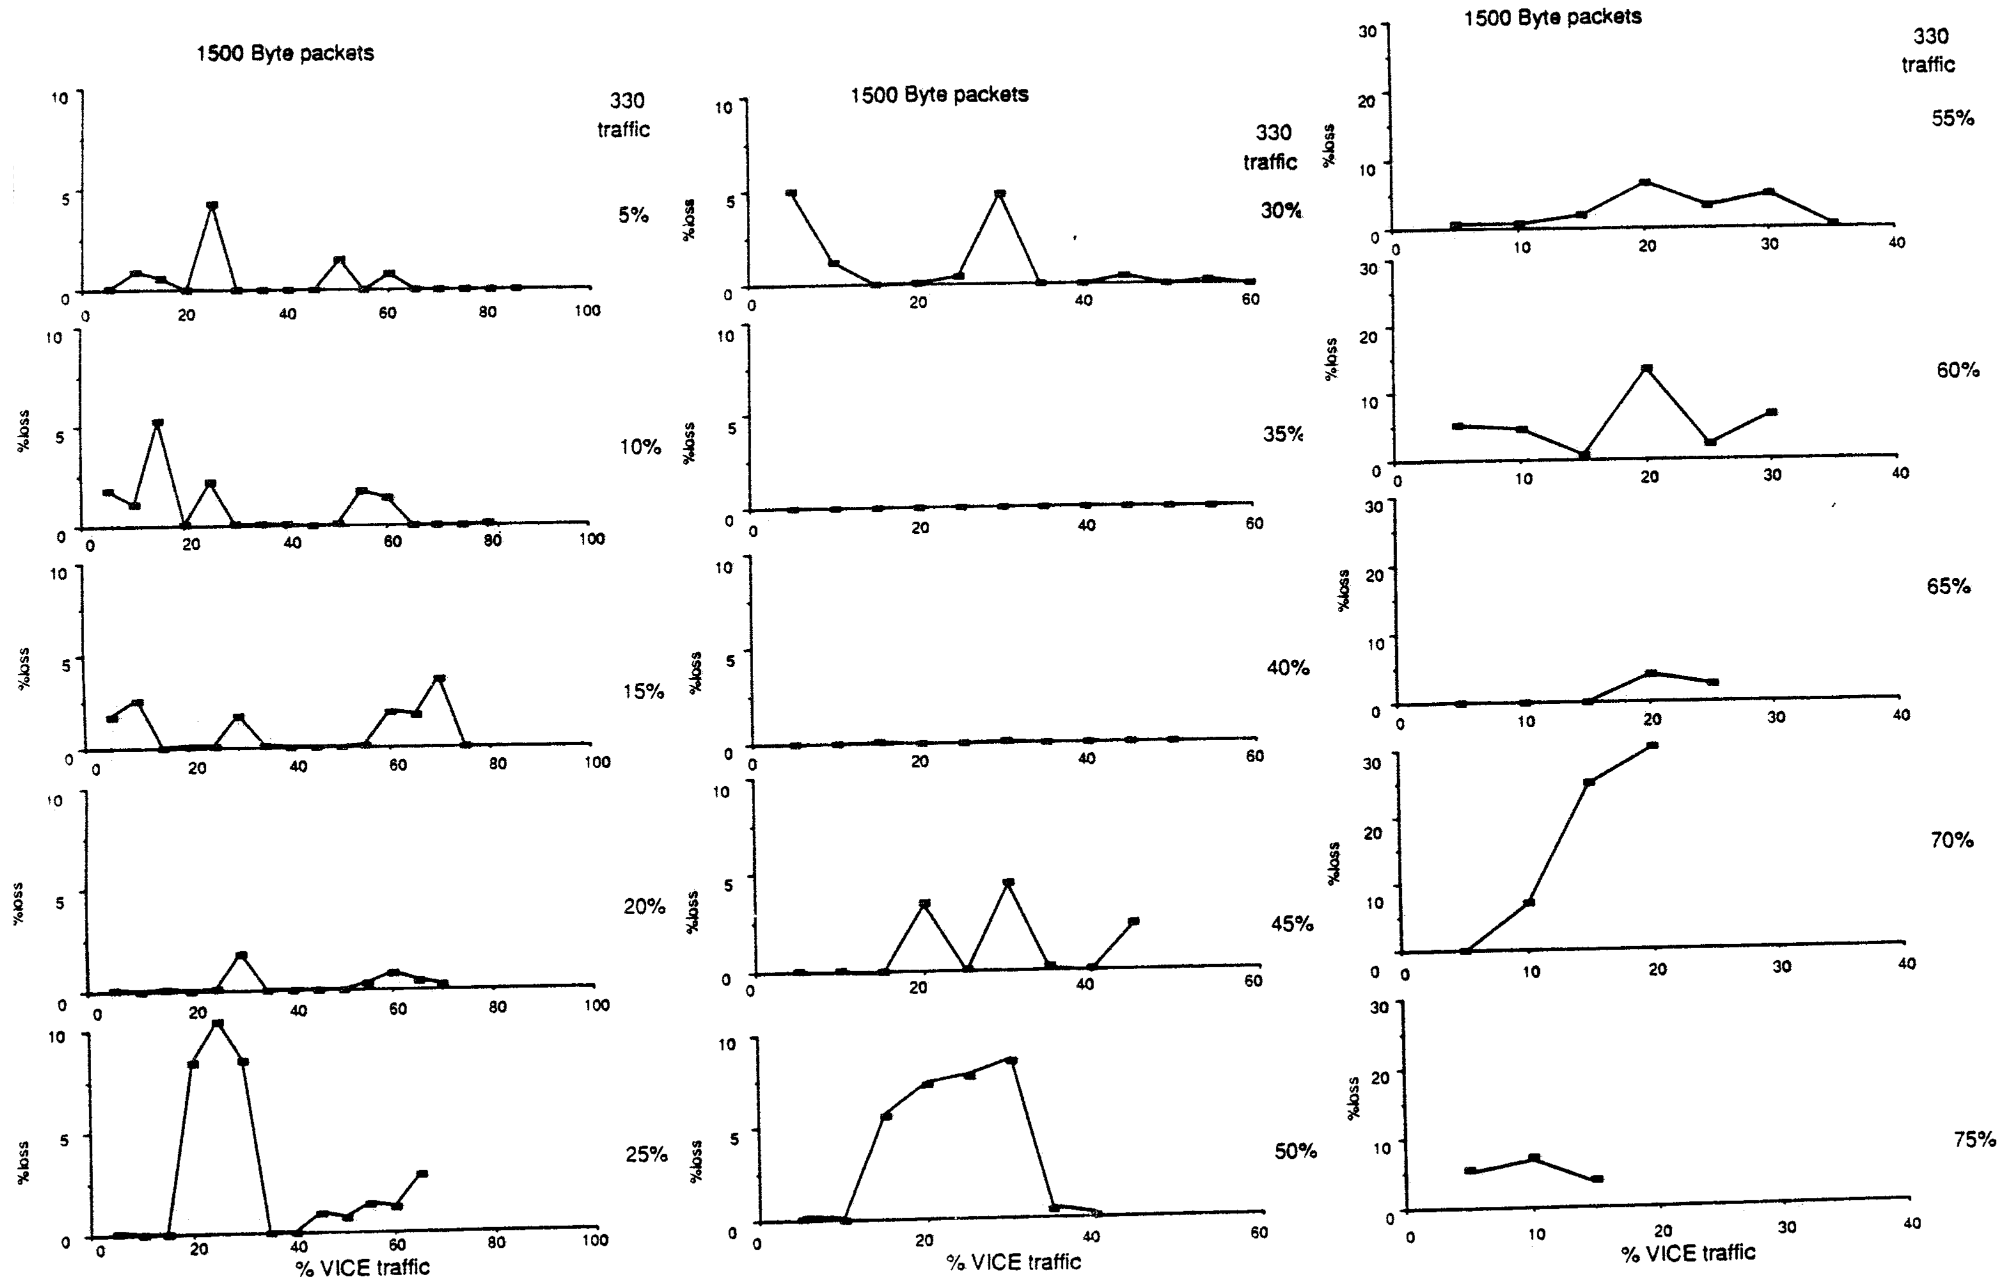

FIG. 2 - % packet loss vs theoretical traffic (VICE) as a function of opposite traffic (330).

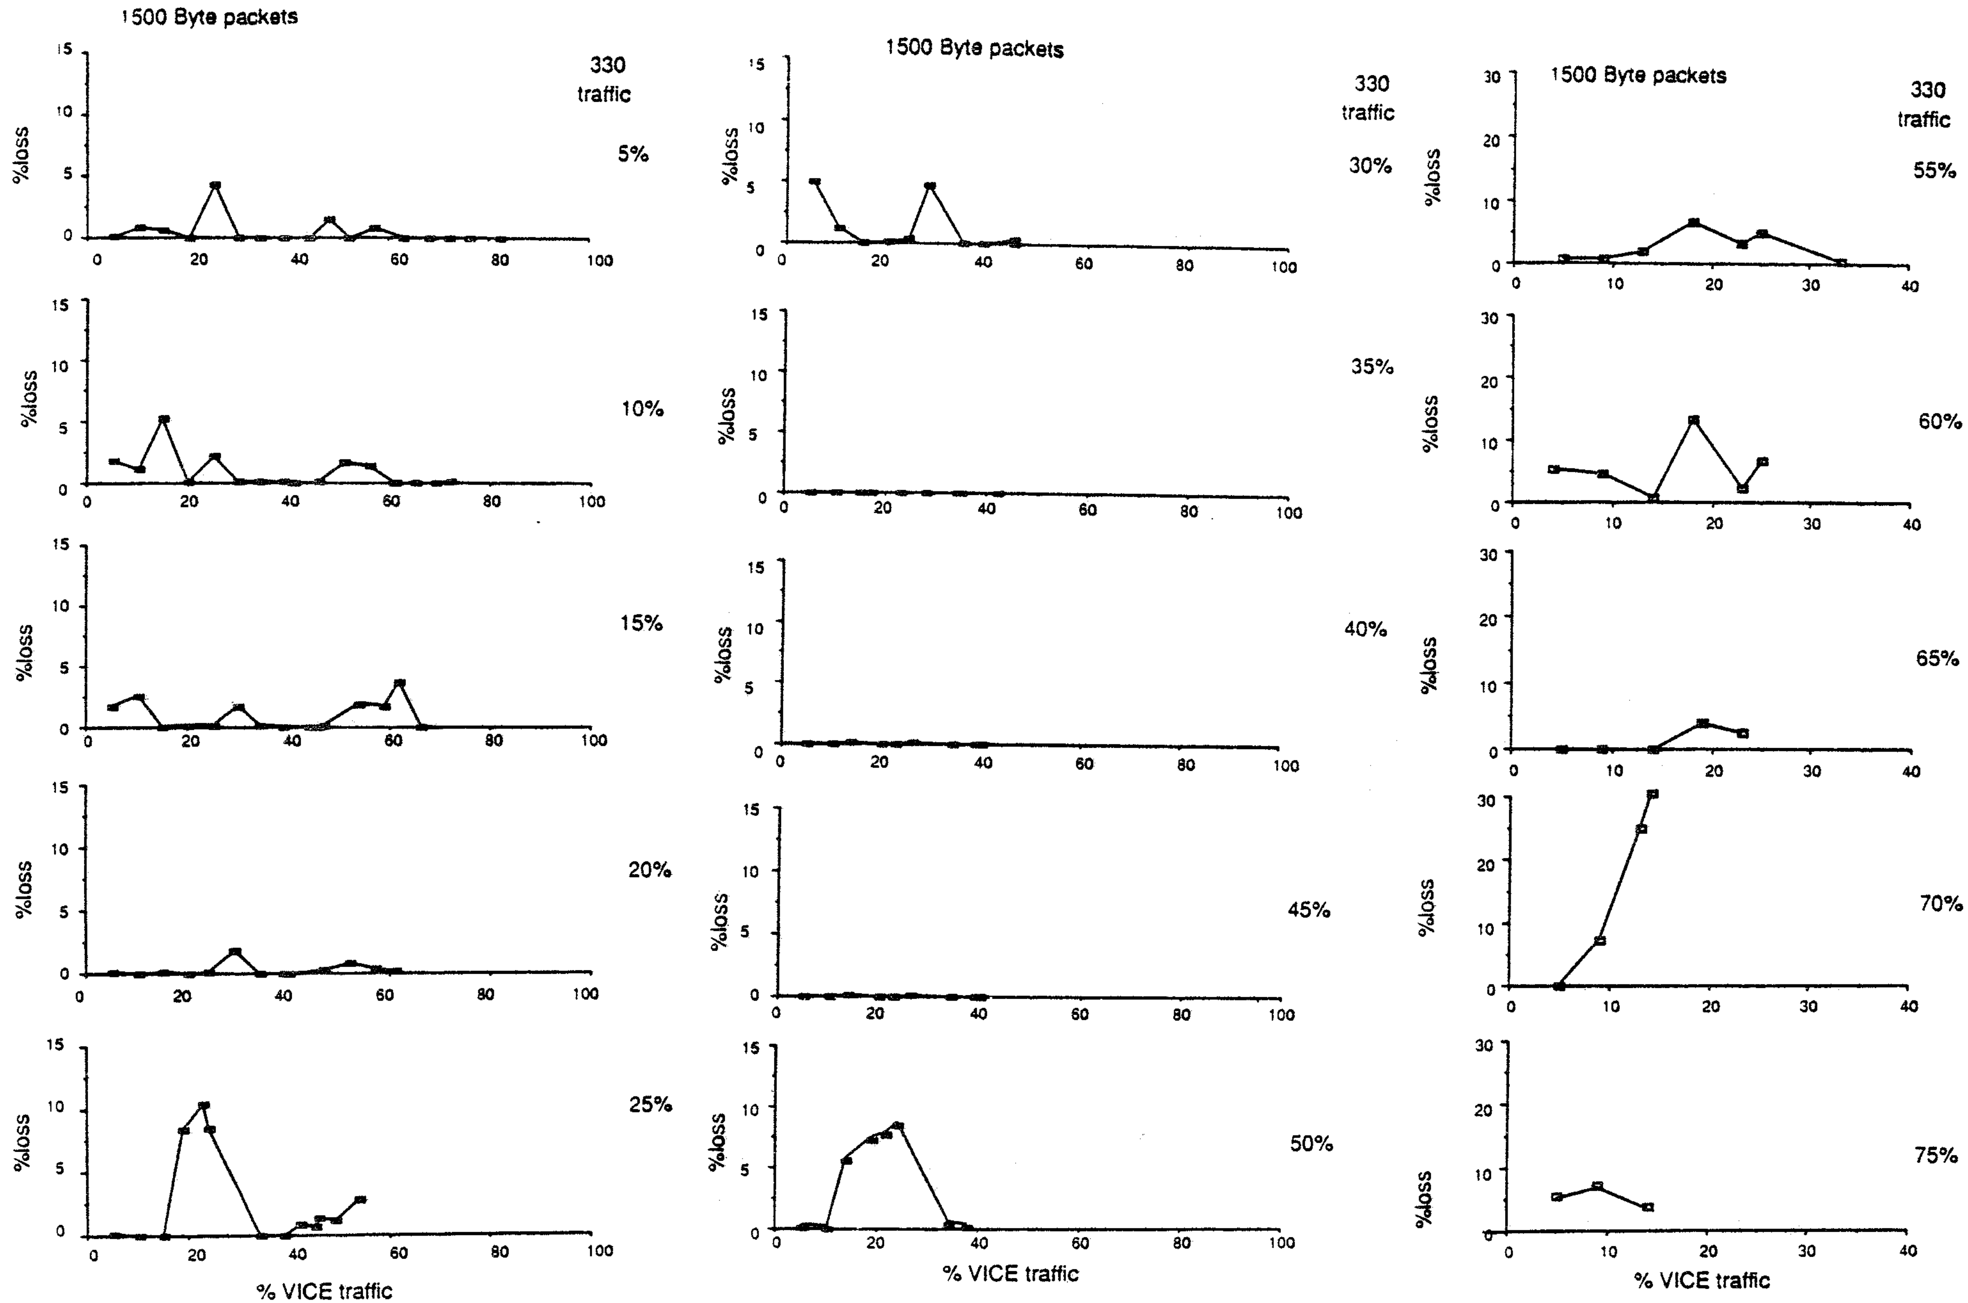

FIG. 3 - % packet loss vs experimental traffic (VICE) as a function of opposite traffic (330).

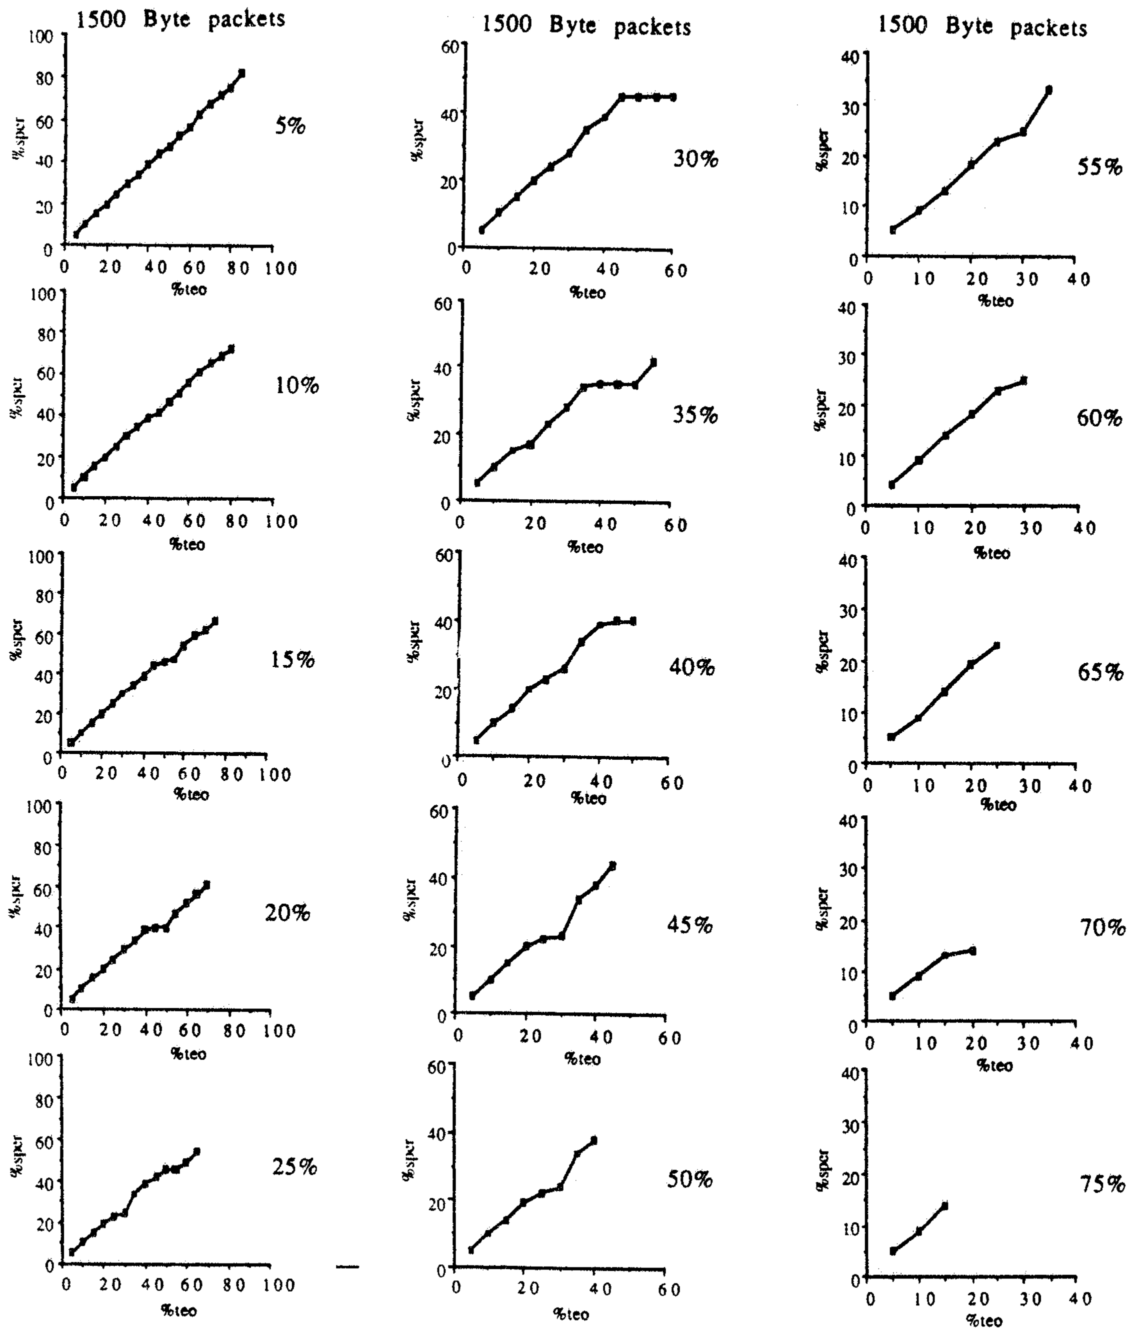

FIG. 4 - Measured VICE traffic vs theoretical VICE traffic as a function of opposite traffic.

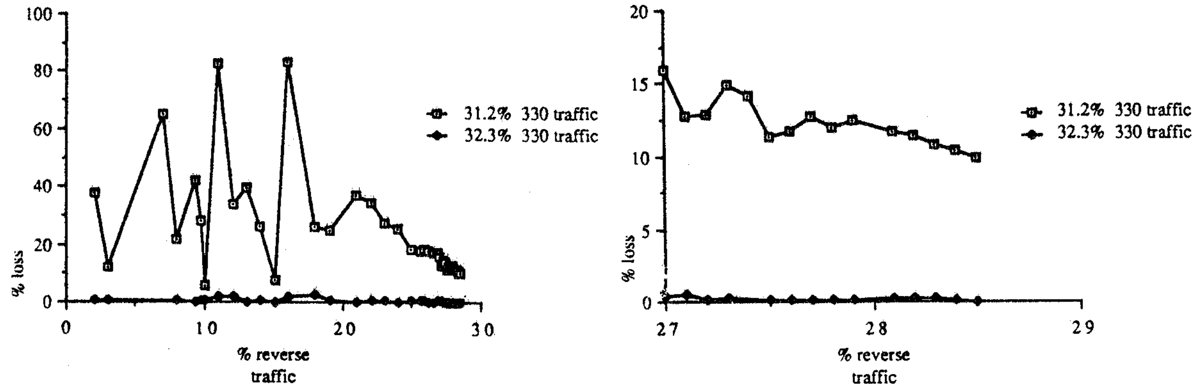

FIG. 5 - % packet loss vs theoretical VICE traffic: detailed scanning of a critical region.

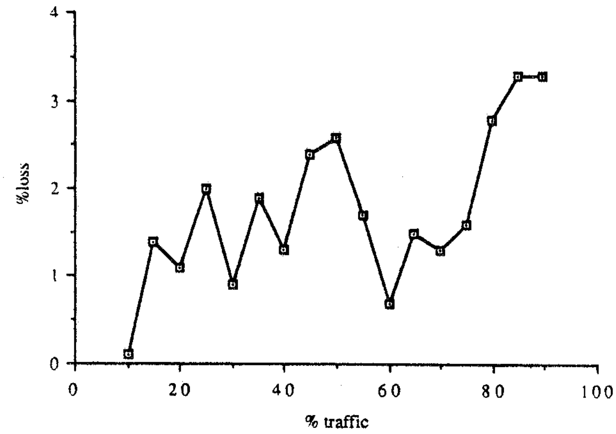

FIG. 6 - % loss (VICE) vs sum of traffic in the two directions.

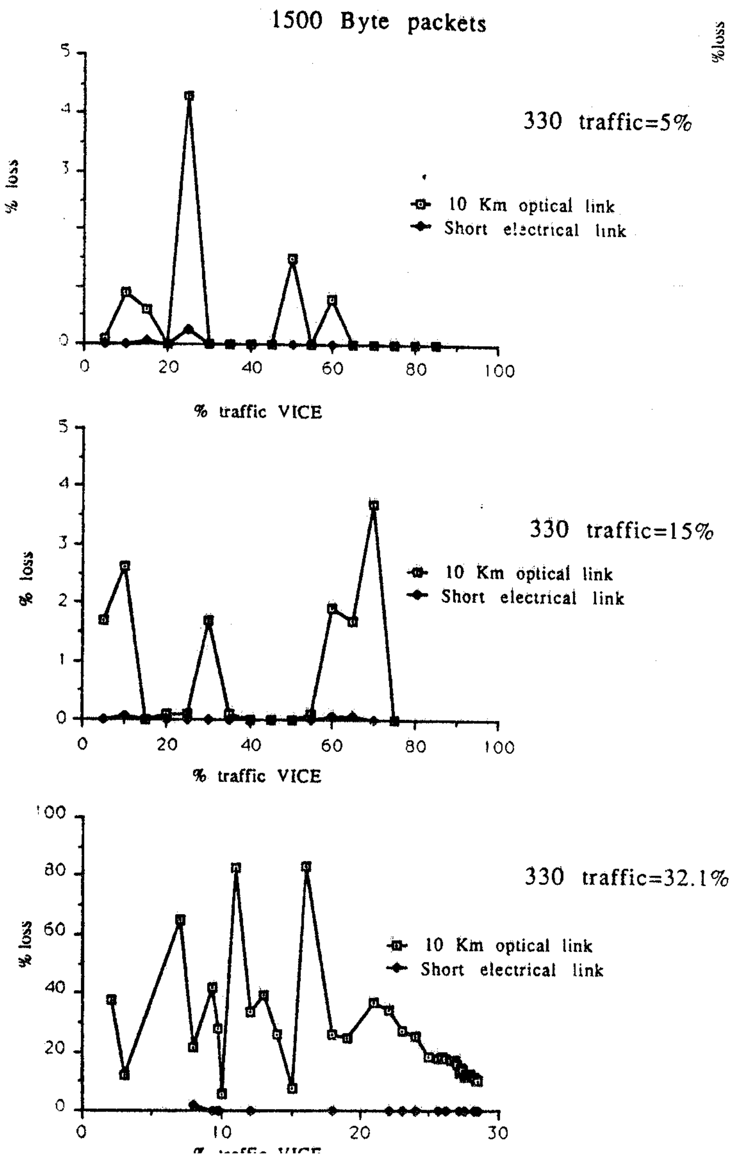

FIG. 7 - Comparison between measurements with 10 Km fiber optic link and short electric link between the two bridges.

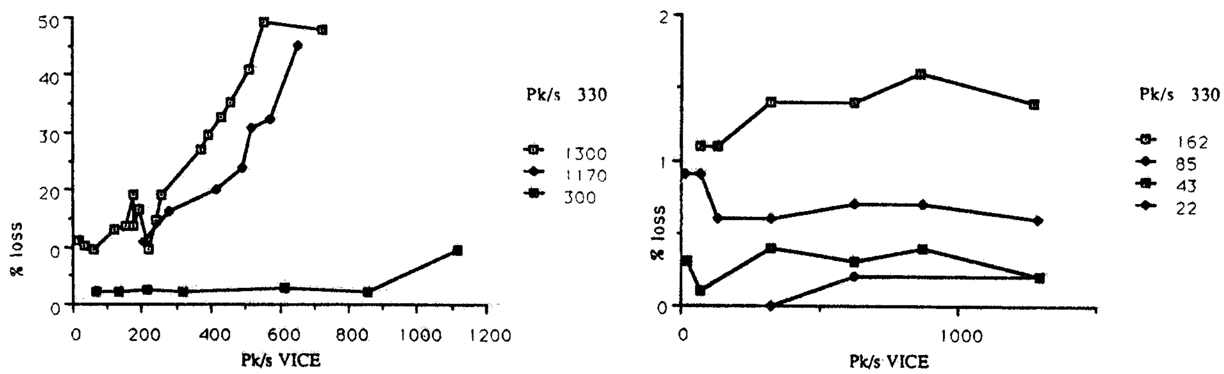

FIG. 8 - % loss vs VICE traffic as a function of 330 traffic (64 Byte packets).

We wish to thank DEC Italia who supplied the DEC LAN Bridges at very favorable conditions, and A. Mannarino for his help in understanding their setup configuration; G. Stabelini, who designed and built the VICE interface and its traffic generator software; O. Ciaffoni, M.L. Ferrer, G. Mirabelli and E. Valente for stimulating discussions.