



Fermi National Accelerator Laboratory

FERMILAB-TM-1773

Network Load of X Terminals at CDF

F. Abar, C. O'Reilly and E. Wicklund

*Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510*

March 1992

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Network Load of X Terminals at CDF

F. Abar, C. O'Reilly, E. Wicklund

A report on the measurement of the network
bandwidth used by X terminals at Fermilab.

Abstract:

As an aid to planning future usage of X terminals, the authors measured the network bandwidth used by X terminals on a single Ethernet in a typical application environment at Fermilab. The objectives were twofold. The first objective was to identify the impact of adding one more X terminal onto the CDF Ethernet. The second objective was to determine the maximum number of X terminals allowable on the CDF network. The findings reveal that an X terminal in this environment used, on average, 431 +/- 45 bytes/second. This is quite small compared to the utilizable bandwidth of the thinwire Ethernet network (375 Kilobytes/second, equivalent to 30% of total bandwidth). This amounts to an average load of .034% on the network per X terminal. This paper addresses the network load solely. It does not address the loads imposed by X terminals on the memory and CPU of the host systems.

1.0 Introduction

The number of X terminals at Fermilab has grown from near zero to over 100 in the past year. It is expected that this number will double in 1992. The chief use of X terminals at Fermilab has been as replacements for character cell terminals, although, once placed, they are often viewed as workstation alternatives. While X terminals increase user productivity over character based terminals, like workstations, they add complexity to system and network administration. Quantifying the load X terminals put on the network was the focal point of a Computing Division strategy meeting in April, 1991. Since no clear consensus was reached on network load and due to the intrinsic topology of the CDF environment, obtaining actual network load data to forecast the affects of additional X terminals in the CDF portakamps was deemed valuable.

The Collider Detector Facility (CDF) has over 35 member institutions, 375 physicists with approximately 125 of them working full time at the Fermi National Accelerator Laboratory (Fermilab). If the desktops of these members are examined, a wide range of computing devices are found. There are over 50 DECwindows-based VAXstations, over 25 NCD X terminals, over 10 Silicon Graphics workstations, and many other devices. At present, these are all on one Ethernet (known as CDF_PKMPS), and it is clear that the network utilization of these devices is an important question to examine. The number of desktop devices is rapidly growing in preparation for the 1992 data taking run. Careful planning for this continuing growth is essential.

The usage of these devices can be best described as a software development environment. The CDF members write and debug programs much of their day. Visual display of results of this activity is a quite important, but relatively infrequent, part of the total activity. There are no continuously updating displays or real time video applications. Unlike some applications of X-based displays (such as stock exchanges), there are no applications which require simultaneous updates of many displays in the CDF environment. An individual's usage of the network can best be described as placing bursts of activity on the network. Usage is not at all constant and is not correlated from one display to another.

2.0 The Experiment

Several testing methodologies were considered in designing this experiment. The most thorough way would be to actually add several hundred X terminals onto an Ethernet and view the affect(s) on the network. Neither Fermilab nor other organizations have the financial or labor resources to conduct such an experiment. The most feasible and optimal estimation technique is to simulate the existence of several hundred X terminals. A simpler testing methodology was implemented in this experiment, sampling of a small number of active X terminals. While this is not the most accurate technique, it is sufficient for the purpose of this experiment.

An existing Ethernet in the CDF portakamps (CDF_PKMPS) was chosen for the network traffic measurements. Seven NCD X terminals were chosen at random. They were of two varieties, the 17" color model, the NCD17c, and the 19" monochrome model, the NCD19. A Novell LANalyzer was set up to monitor detailed information on network traffic of the global network as well as the X terminals. The Novell LANalyzer is capable of filtering and recording the traffic activities of up to eight channels of communication. Each channel was set up to monitor traffic that was either sourced from or destined to the Ethernet address of the X terminals under investigation.

The X terminals were monitored during the daytime hours (10:00am to 6:00pm) for a period of one work week (5 days). These hours were chosen since they are the normal working hours of CDF user community. Network traffic statistics such as bit rate, packet rate, and network utilization of each X terminals and the overall network traffic of the CDF Ethernet were recorded at one minute intervals. The users of the seven X terminals were unaware of the experiment and their network traffic was not disturbed or biased by it.

2.1 Caveats

Several caveats should be considered before applying the findings of this report to another environment.

First, the findings mentioned below are specific to the CDF experiment at Fermilab. Other experiments within the lab and unrelated organizations will have different user environments and user profiles from those tested in this experiment and therefore different average loads. It is important to note that the CDF user profile does not include continuous real time displays or downloading of bitmaps, which imply a significantly higher load on the network.

Second, this experiment used X terminals that meet what the Computing Division refers to as "minimum acceptable functionality".¹ The minimum acceptable functional X terminal is defined as an X terminal with features such as a local window manager, local clients and a minimum of 4MB of RAM. These features imply that a potential portion of the network load was off-loaded since the window manager and clients (e.g. local telnet) were run locally on the X terminal as opposed to the host computer. These restrictions may not apply to other organizations.

Third, in determining the maximum number of X terminals allowable on the CDF portakamps Ethernet, the locality of access was not considered. Ethernet bandwidth (and consequently allowable number of X terminals) can be increased by simply partitioning it into smaller Ethernets through the use of bridges and routers. This approach, however, gets harder and harder to implement as the number of hops between the X terminals and application servers increases.

Fourth, one must use care in interpreting the peak Ethernet utilization data calculated in this experiment. The LANalyzer updates its counters as network traffic occurs. However, since a value was requested every 60 seconds from the LANalyzer, it averaged the cumulative data over the 60 second period to attain its reported value. Therefore, the peak values reported were smaller than the actual instantaneous peaks since the peak value reported is an average of the 60 second aggregate.

1. C. O'Reilly and U. Pabrai, *Recommendations for a Computing Division Policy on X Terminals*, Fermilab Computing Division internal document DR0003, July 1991

Fifth, the conclusions drawn from the data in this experiment assume that there is a linear relationship between the number of X terminals and the load on the network. That is, the 100th X terminal puts the same absolute load on the network as the first X terminal. However, as the number of X terminals increase, the probability of collisions also increases; introducing a non-linear relationship between the number of X terminals and their network load. The number of collisions were not monitored in this experiment.

Sixth, peak and average utilization figures are used to achieve the objectives of this report. The judgement of which figures best describes reality is a subjective decision. This report presents both in an objective light, applies them in drawing a conclusion for CDF, and offers a paradigm in choosing between the two.

Lastly, the experiment was not conducted during the peak time for the CDF experiment. When the experiment starts taking data later this year, there will be considerably more users and CPU cycles.

3.0 Interpretation of Data

Figure 1 illustrates the average and peak data for the seven X terminals averaged over five days to depict a typical day as a function of time of day.

Average bandwidth values were calculated by the iterative process of averaging the data over five minute intervals, then averaging the results over the seven X terminals, and lastly averaging the results over five days. This created one dataset of average values for every five minutes between 10am to 6pm.

Peak bandwidth values were calculated in a similar iterative process. The maximum value within each X terminal dataset was carried across a five minute period, then across the seven X terminals, and lastly across the five days. This created one dataset of peak values for every five minutes between 10am to 6pm.

Figure 2 illustrates the average and peak data for all Ethernet traffic over the five day period to depict a typical day as a function of time of day.

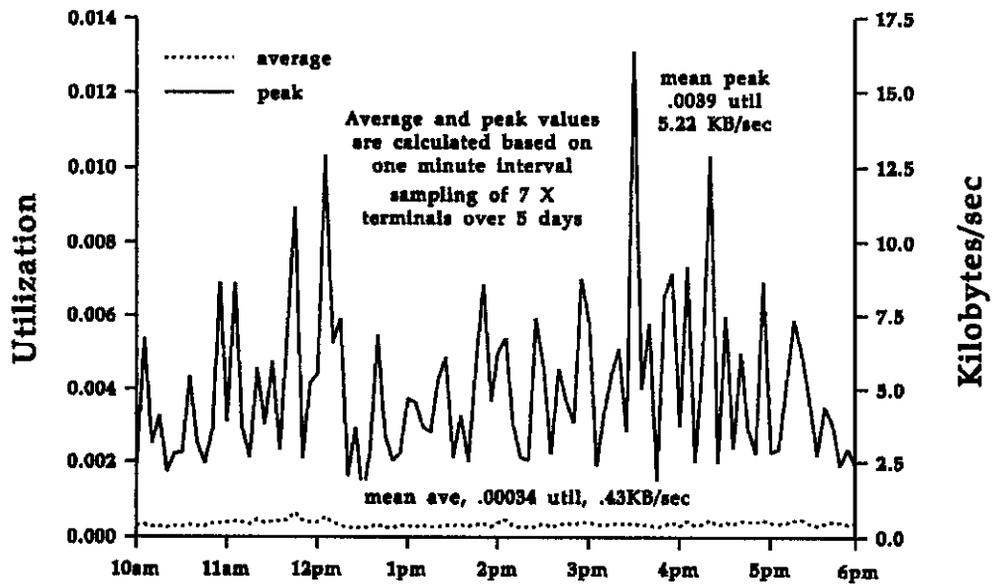
The mean utilization of the entire CDF portakamp network was 16.4% for the period monitored, while the mean peak utilization was 30.8%. The total number of X terminals allowable before degradation in network performance occurs was calculated by dividing the average and peak values ($X_{terminal_{Mean}}$ and $X_{terminal_{Peak}}$) into the difference between the maximum utilizable bandwidth ($Util_{max}$) and the total average/peak utilization ($Util_{GlobalMean} / Util_{GlobalPeak}$). The following formulas illustrate this calculation.

$$NumberofXterminals_{ave} = \frac{(Util_{max}(0.30) - Util_{GlobalMean}(0.164))}{Util_{XterminalMean}(0.00034)}$$

$$NumberofXterminals_{peak} = \frac{(Util_{max}(0.30) - Util_{GlobalPeak}(0.308))}{Util_{XterminalPeak}(0.0039)}$$

FIGURE 1.

Average and peak network load of X terminals on the CDF portakamp Ethernet.



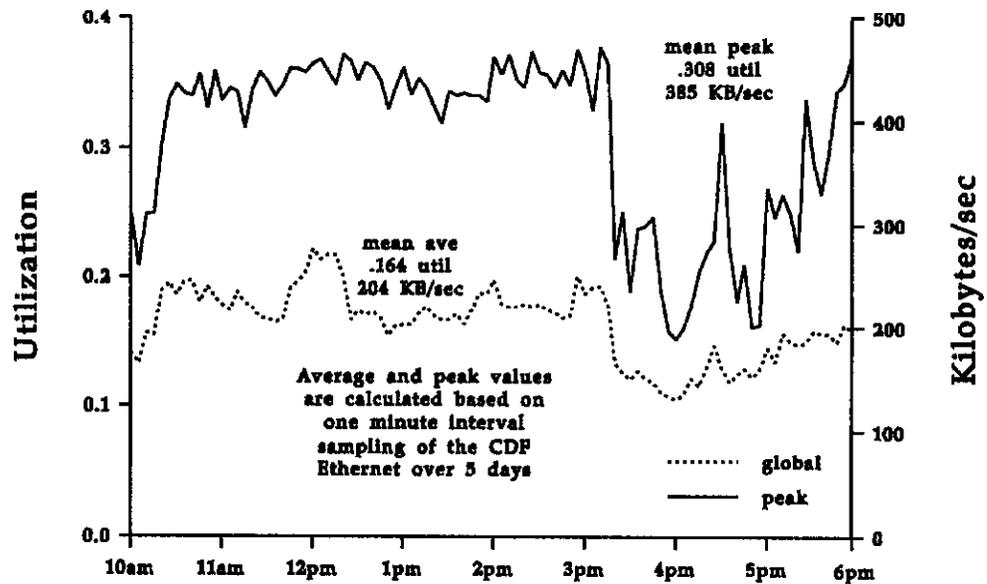
4.0 CDF Specific Results

The conclusions drawn from the analysis that specifically pertain to the CDF experiment are fivefold. First, an average bandwidth utilization of one X terminal is 431 +/- 45 bytes/second. The systematic error on the measurement was estimated by varying the number of days and the number of X terminals. The average peak bandwidth utilization of one X terminal is 5,223 +/- 567 bytes/second.

This measurement was much lower than expected. It implies that a single X terminal at the CDF portakamps will add an average of .034% load on the network. The average peak load is .39%. A large part of the difference between peak and average is the recognition that the users do not continuously change the displays on their X terminals, but pause for various reasons (e.g. thought) between redraws.

FIGURE 2.

Average and peak load on the CDF Ethernet



Second, regardless of the method of calculating the incremental load of an X terminal, the CDF portakamps Ethernet can accommodate several more X terminals before there is degradation in network performance. CDF has projected approximately 30 additional X terminals over the next year. Using the formulas described in Section 3, the maximum number of X terminals allowable, based on the average utilization figures, is 400. Given the peak utilization figures, since the mean peak utilization already exceeds the maximum utilizable bandwidth (30.8% v. 30%), this Ethernet is already too heavily loaded to add more X terminals. The fact that these numbers are so close is coincidental. In this experiment, the mean peak values tended to be limited by the capabilities of the host system, Ethernet controller and X terminal processor. If the efficiency of host systems and Ethernet controllers were optimal, one could expect the mean peak utilization to be significantly higher and perhaps approximate the maximum utilizable bandwidth.

Third, since the average and peak values for this Ethernet are very high, steps have been taken to reduce them. Specifically, this Ethernet has been partitioned into smaller Ethernets via bridges.

Fourth, after studying the data, it was evident that since the individual users are independent of each other, the peak bandwidth utilization is also independent. Clearly, peak utilization cannot be scaled by a simple multiplicative factor to estimate the peak utilization of several X terminals. It is highly improbable that all users will request a display simultaneously. The highest recorded peak utilization of any one X terminal in a 60 second period was 17 Kilobytes/second, which corresponds to 1.4% utilization.

Lastly, the one minute averaged Ethernet utilization sample points were gathered to calculate a probability distribution of an X terminal bandwidth utilization within the CDF Ethernet. Figure 3 illustrates the cumulative probability of the random variable μ , percent utilization, of a CDF X terminal. The sample of 17,213 points was taken from seven X terminals over five working days. As is shown, a CDF X terminal is idle (0% utilization) about 30% of time, and there is a 90% chance that a CDF X terminal will have less than a 0.1% load on the network. As mentioned earlier, this mode of computing is specific to the CDF Ethernet which closely parallels a software development environment.

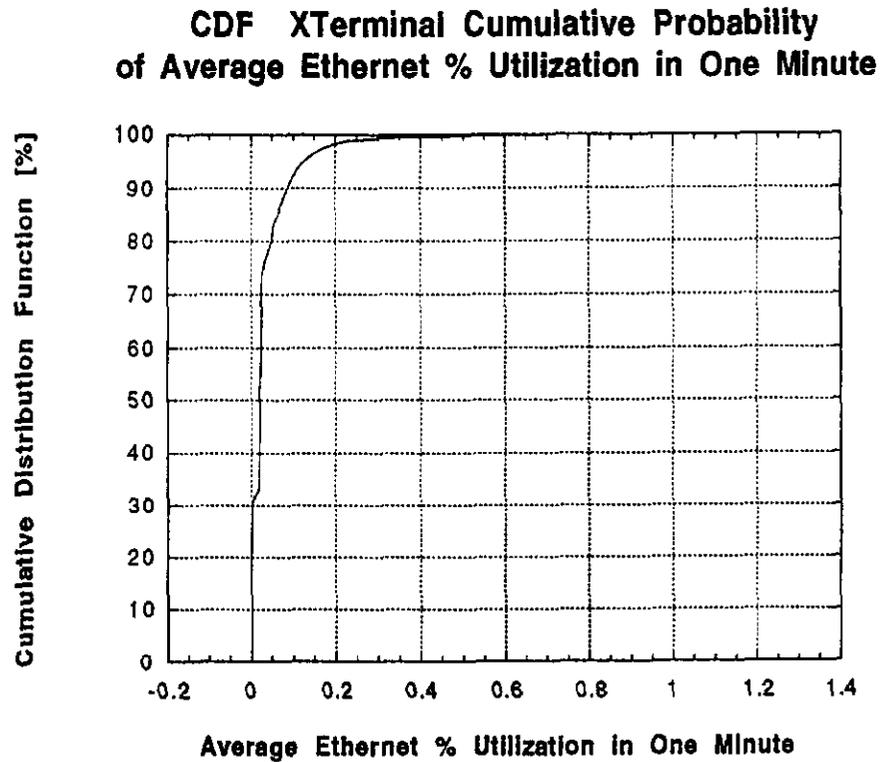
5.0 General Observations and Conclusions

In this experiment, much was learned about the dynamics of the CDF portakamps network architecture and the load that X terminals put on that Ethernet. The discussion above dealt with specific conclusions for the CDF experiment. The following discussion presents more general conclusions that will apply to all experiments at Fermilab as well as to the central network in Wilson Hall.

First and foremost, this experiment uncovered the fact that since the network load of X terminals for non-graphical applications is minuscule, other potential bottlenecks deserve more serious consideration. As the number of X terminals and other X display devices (e.g. workstations and PCs) are added to the network, the host machine's CPU, memory and Ethernet controller can get saturated. While the first two are of concern at CDF, the host Ethernet controller is not an issue. These three potential bottlenecks deserve more serious attention than the network bandwidth utilization. The network load is unlikely to be the limiting factor.

Second, network architecture is of primary importance in managing network load. Network load of any X display device can be readily controlled by adding bridges and routers. However, this will increase the response time.

FIGURE 3.



Third, for a thumbnail sketch, designed for network planners, for determining the number of X display devices that can be placed on a given Ethernet, the following paradigm is offered. It is based on the organization's attitude toward risk (e.g. risk of long response time, increased network collisions and subsequently lower productivity) as applied to the allowable number of X terminals based on average and peak numbers. For example, if the allowable number of X terminals based on average and peak numbers are 400 and 50 respectively (assuming the Ethernet is not already fully utilized based on peaks as in the case above) and your organization is risk-adverse, then the appropriate number of X terminals would be closer to 50. If your organization is not risk-adverse (i.e. risk-loving), a more appropriate number of X terminals would be closer to 400. Placing 400 X terminals on an Ethernet may be feasible from a network load standpoint, but is not recommended by the authors.

This model is intended to be used only as a thumbnail estimation. This estimation does not consider other factors in determining the appropriate number of X terminals such as host machine CPU and memory. As is evident from the magnitude of the numbers based on averages, organizations will need to assess these bottlenecks before implementing a large number of X terminals.

Fourth, the role of collisions and network contention was not addressed. When host CPU and memory are allocated effectively, determining a more accurate number of X terminals is close at hand. The most effective way to encompass the affect of collisions and, in general, model the network architecture is through simulation. In the simulation, as X terminals are added, the experimenter can investigate the network load, collision rate and the delay time experienced on the network.

5.1 Acknowledgments

The authors wish to thank the HEPNET organization for the loan of the LANalyzer used in this experiment. Additional thanks are in order to Phil DeMar, Mark Fischler, Mark Kaletka, Jeff Kallenbach, Bill Lidinsky, Jack Pfister, Donna Reid, JJ Schmidt and Joe Stith for their useful discussions during this experiment and review of this paper. This work was conducted by the Fermilab Computing Division.