Performance comparison of PC LAN file servers — an experimental study

Gee-Swee Poo and Tiow-Seng Tan evaluate the file server capability of a number of PC LANs

This paper considers a comparative performance measurement study which was carried out to evaluate the file server capability of a number of PC LANs, including the 3Com 3Plus network, Novell Advanced Netware 286 system and IBM PC network, all using IBM PCAT as a server. In addition, a special server, 3Com 3Server, is also measured for comparison. In the measurement, the server is subjected to simultaneous access from a number of stations performing read/write operations with varying message sizes. The response time and throughput characteristics of the networks are determined and analysed. A comprehensive comparison is made to discuss the strength and weakness of the various network implementations. The results give impartial performance ratings from an independent source, which are useful to network consumers and implementers. The authors conclude that the performance of file server is insensitive to network technology, but depends mainly on the implementation of file server software which is apparently lagging behind the hardware advancement.

Keywords: local area networks, read/write operations, network, performance measurement, file server

The onslaught of personal computers on offices and within organizations has brought with it the phenomenal growth of PC LANs. A major family of these LANs supports MSDOS 3.1-NETBIOS software interface for IBM and compatible PCs, which is now a *de facto* standard[1]. These include 3Com 3Plus system[2, 3], Novell Advanced Netware[4], IBM PC network[5] and IBM Token Ring network[6]. The first two are based on the IEEE 802.3 standard[7], and the third is a proprietary broadband network. The last is based on IEEE 802.5 standard[7].

Department of Information Systems and Computer Science, National University of Singapore, Kent Ridge, Singapore 0511

A LAN is needed to interconnect PCs as well as associated peripherals. A key component is the file server which provides not only the data sharing necessary in a distributed system, but also the focus point for information interchange. These capabilities overcome the inherent shortcomings of isolated PC systems. Thus, the server is the most heavily used resource in a network. Consequently, the performance of the file server is critical to the success of the PC LAN.

The capacity of the file server in a LAN is governed by the hardware and software design of the server, the disc channels and the network capacity. The server design is the dominant factor, whereas the network capacity is the least sensitive. A specially designed server hardware (e.g. 3Server) usually performs better than one built on a general purpose machine (e.g. PCAT). However, an optimized server software can outperform the hardware feature. In order to make an 'apple-to-apple' comparison, a common server hardware, the IBM PCAT is used. Nevertheless, to enhance the comparison a special hardware, the 3Server, is also included in the measurement.

A description of performance measurement of the 3Com 3Plus system, Novell AN286 and the IBM PC network is made, as these were the networks available to the authors. In the experiment the server was subjected to simultaneous access from one up to 15 workstations performing read/write operations with varying blocks of data. A special timing program written in assembler provided accurate timing measurements and automatic recordings. The response time and throughput characteristics of the networks were determined and analysed from measurements, and a comprehensive comparison made.

0140-3664/88/020071-09 $03.00

The strengths and weaknesses of various network implementations are also presented.

NETWORK SYSTEMS

A common feature of PC LANs is the provision of one or more file servers supporting the activities of a group of PC stations. The functions provided include file management, print spooling, electronic mail and software applications. The server provides centralized network and peripheral management functions, whereas the application processing remains the task of an individual station. The implementation of PC LANs varies, depending on the server design and network technology adopted.

Four major configurations are considered in this study, the 3Com 3Plus system, the Novell AN286 and IBM PC networks, all using PCAT as a server, and additionally the 3Com 3Plus system with a special server, called 3Server. Network topologies of the systems are shown in Figure 1. The 3Com and Novell systems follow a bus topology, whereas IBM PC network uses a tree topology. These systems have different server design and network technology, a comparison of which is summarized in Table 1. As shown, the first three configurations use identical server hardware but different server software. The last configuration uses a special server hardware.

Server software development has undergone a significant evolution. The earlier version of disc-server design has become obsolete, and file server design has emerged to be the common goal. This development has been triggered by an important event: the introduction of MS-DOS 3.1, IBM PC LAN program and Microsoft networks. The MS-DOS 3.1 provides a standardized multiuser access to shared files. It has been accepted in industry as a *de facto* standard to which developers write multiuser applications software that will run across a variety of LANs. To be compatible to the standard, LAN vendors must now provide DOS 3.1-compatible file server software. This is to ensure that all PC applications software can run on the DOS 3.1-compatible LAN. For the PC LANs of interest here, the IBM implementation is the PC LAN program which runs on both the PC network and the token ring. The 3Com implementation is the 3Plus file server software, which is a modified version of MS-Net. The Novell implementation is the high performance Advanced Netware. All the server softwares are made to be compatible to DOS 3.1. However, they are not compatible with one another. It is not possible to communicate between 3Plus and PC LAN program or between 3Plus and Advanced Netware.

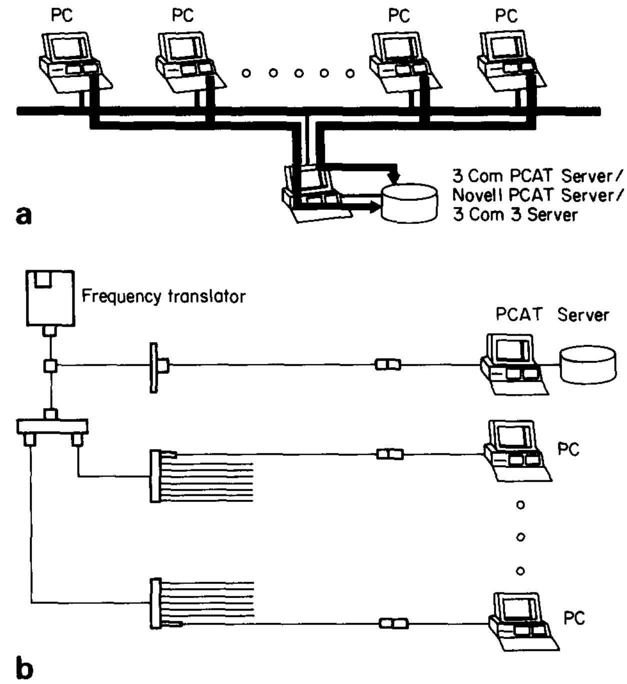

Figure 1. PC LAN topologies (a) 3Com and Novell; (b) IBM

There is no standard on server hardware. The issue is being neglected as the server software, once standardized to DOS 3.1, will provide users across the board functions in a system. All PC LAN systems have server implementations on IBM PCAT or AT-compatible machines. However, the PCAT is not an efficient machine for supporting many stations. As a result, some special machines are produced as servers. These include 3Com 3Server and Novell 286A/286B, which are specially designed and optimized to yield high throughput performance.

Network technology varies. For the four configurations under study, two different types of technology are involved (see Table 1). The implementation of Novell employs the Etherlink adapters produced by 3Com, and hence the same network technology. Only the server software differs. The 3Com network is based on the IEEE 802.3 standard which specifies a 10 Mbyte/s data rate,

Table 1. Feature comparison of PC LANs

	3Com 3Plus	Novell Netware	IBM PC Network	3Com 3Plus
Server software	3Plus	Advanced Netware	PC LAN Program	3Plus
Server hardware	PCAT	PCAT	PCAT	3Server
Network technology	CSMA/CD Bus Baseband 10 Mpbs Coax	CSMA/CD Bus Baseband 10 Mpbs Coax	CSMA/CD Tree Broadband 2 Mpbs Coax	CSMA/CD Bus Baseband 10 Mpbs Coax
Adapter standard	802.3	802.3	Proprietary	802.3

baseband bus topology and the CSMA/CD access method. In terms of ISO-OSI layered structure, the specification covers the physical and link layers which are implemented in the network adapter. This provides a low level communication channel. The higher level protocols are software oriented. The protocols adopted include the Xerox XNS protocol for network and transport layers, the IBM NETBIOS for session layer and the 3Plus server software for the application layer. The adherence of DOS 3.1 gives a common interrupt interface for multiuser access. The adoption of XNS is for the ease of internetworking, which is lacking in SMB protocol.

The IBM PC network is a proprietary broadband LAN designed in a tree-shaped topology (see Figure 1). The network adapter contains a broadband modem which uses frequency shift keying (FSK) to transmit digital signals at 2 Mbyte/s rate. The modem transmits in a 6 MHz channel centred at 50.75 MHz, and receives in a 6 MHz channel centred at 219 MHz. A network translator is used to provide the single channel frequency translation. This unit is supported by a simple 'do-it-yourself' cabling kit which allows connection of up to 72 nodes over a distance of 300 m. The network uses the CSMA/CD channel access method, which is similar to the IEEE 802.3 baseband standard. The higher level protocols are implemented using the proprietary Sytek Localnet/PC protocol for network and transport layers, the NETBIOS for session layer and PC LAN program for the application layer. The file server function is provided by the PC LAN program.

MEASUREMENT

In carrying out the performance measurement it was necessary to fix the operating environment[8, 9]. However, it is difficult to choose a 'typical' one without inviting much argument. For the interest of comparing how best a system can perform, the authors resorted to the fundamental read/write operations on the server, taking the entire system into consideration. In the measurement, the server was subjected to a repeated sequential read or write operation via the network with varying blocks of data. This gives a heavy load environment where all the workstations are heavily engaged in endless read or write activity.

A suite of programs was written for the measurement. This included an assembly routine that provides accurate timing measurements and automatic recordings. Three blocks of message sizes were used, viz 0.75, 6 and 30 kbyte. In the write test, the block of data was sent repeatedly to the server from each workstation. The write operation was sequential, but each block was written on a different sector of the hard disc to ensure a good average value of the disc access time. On the other hand, in the read test the block of data was stored *a priori* in the server. Repeated retrievals were made by each workstation. The response time of the network as seen by a workstation was recorded at the completion of each block of transmission. An aggregate of 800 readings was taken for each measurement point. To avoid disturbance in the measurements due to initial and final transients, timings were only accepted during the steady state which was assumed to exist.

In the experiment a common server, IBM PCAT and a special server, 3Com 3Server were employed together with up to 15 PC workstations (see Figure 1). The AT machine runs on the Intel 80286 processor with a clock rate of 6 MHz and one wait state. It has 896 kbyte ram and a 20 Mbyte hard disc. DOS 3.1 operating system is used for the networks. The system parameter setting of the server software can affect the server throughput. In the measurement, the parameter settings recommended by the vendor for maximum file server performance were used.

PERFORMANCE ANALYSIS

The measured data was assembled and analysed to yield the delay and throughput characteristics. A cross comparison is made among the four server configurations, and the results are presented in various forms of performance metrics.

Response time

Figures 2 and 3 depict the mean response time characteristics of read and write access respectively for the four server configurations. A number of interesting features are observed.

The mean response time is characterized by a region of low delay when the station number is small, but increases to a region of high delay when the station number is large. The increase is more rapid for the larger message size and for the write access. This characteristic is common among all networks. The delay is a manifestation of the build-up of requests for file service rather than the medium

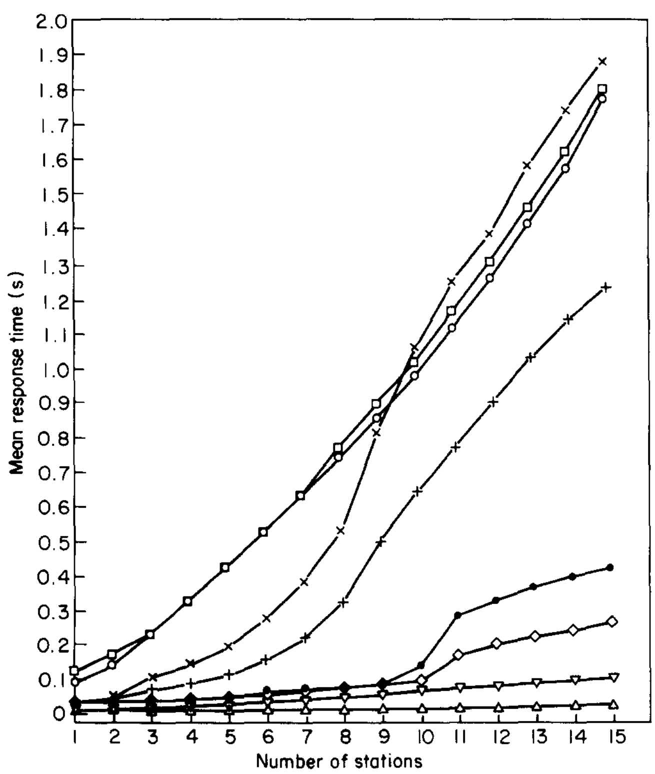

Figure 2. Delay characteristics of various networks for 750 byte read/write operations. □: IBM PC network read; o: IBM PC network write; +: 3Com PCAT server read; ×: 3Com PCAT server write; ◊: 3Com 3Server read; •: 3Com 3Server write; △: Novell AN286 read; ▽: Novell AN286 write

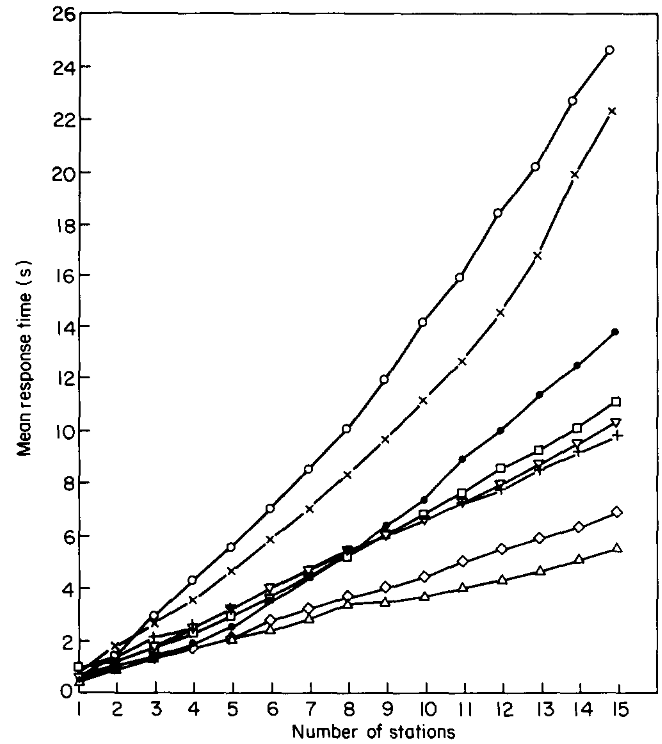

Figure 3. Delay characteristics of various networks for 30 kbyte read/write operations. □*: IBM PC network read;* o*: IBM PC network write;* +*: 3Com PCAT server read;* ×*: 3Com PCAT server write;* ◇*: 3Com 3Server read;* •*: 3Com 3Server write;* △*: Novell AN286 read;* ▽*: Novell AN286 write*

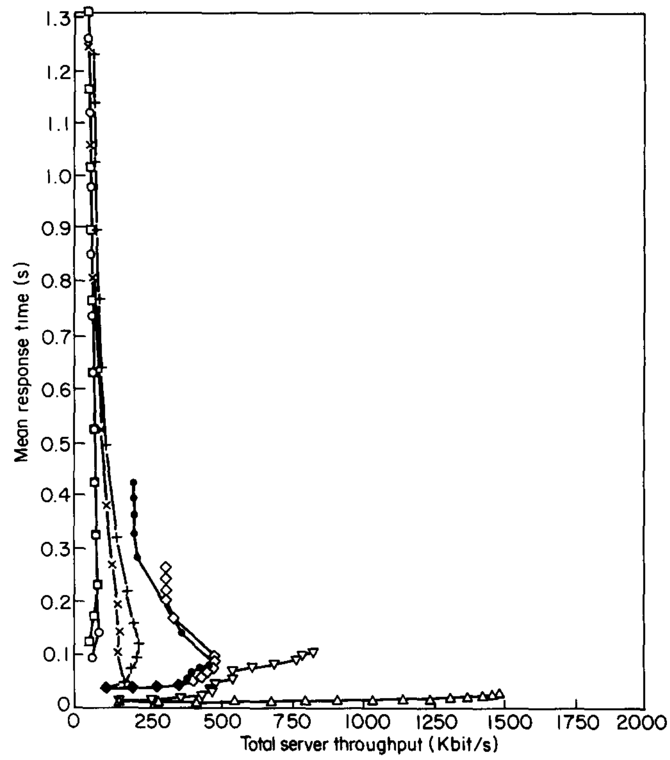

Figure 4. Delay-throughput characteristics of various networks for 750 byte read/write operations. □*: IBM PC network read;* o*: IBM PC network write;* +*: 3Com PCAT server read;* ×*: 3Com PCAT server write;* ◇*: 3Com 3Server read;* •*: 3Com 3Server write;* △*: Novell AN286 read;* ▽*: Novell AN286 write*

contention, the degree of build up being governed by the steepness of the gradient of the curves. It becomes obvious that the file server is the bottleneck of the system.

The read access takes a much shorter time to accomplish the task than the write access. The discrepancy becomes larger with increasing message size and station number. This result can be attributed to the disc access mechanism which has an inherent advantage towards read access. A simple test on the standalone server without the network confirms this tendency.

By and large, among the four configurations tested, the Novell AN286 gives the best response time, whereas the IBM PC network yields the worst delay. The 3Com 3Plus network falls in between, with 3Server yielding a better response than the PCAT server.

A detailed inspection shows that in the region of small message size (see Figure 2), the Novell AN286 gives extremely good response, especially for the read access. This can be attributed to the efficient implementation of the disc caching mechanism in the file server software, in which a 4K cache buffer is assigned to each workstation involved. This is apparently more efficient than the common cache buffer pool approach as implemented in the 3Com as well as IBM PC LAN program. On the other hand, in the region of large message size (see Figure 3), two distinct clusters appear for read as well as write access, with the 3Com 3+ 3Server and Novell AN286 lying in a low delay group, and 3Com 3+ PCAT server and IBM PC network in a higher delay group. The response of 3Server is comparable to that of Novell in this region. In fact, in some workstation ranges, it does outperform the Novell. This confirms the general speculation that the

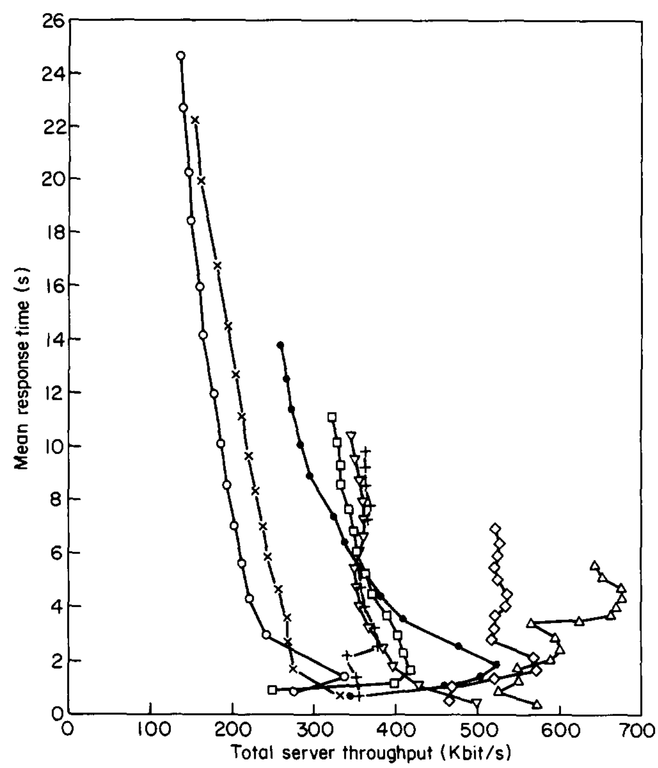

Figure 5. Delay-throughput characteristics of various networks for 30 kbyte read/write operations. □*: IBM PC network read;* o*: IBM PC network write;* +*: 3Com PCAT server read;* ×*: 3Com PCAT server write;* ◇*: 3Com 3Server read;* •*: 3Com 3Server write;* △*: Novell AN286 read;* ▽*: Novell AN286 write*

Novell's approach is superior for the small message size, whereas the 3Com's design is better for large message size. The caching mechanism is insensitive in the region of large message size.

Throughput

The total effective server throughput, S, is defined as the total amount of messages being processed by the server in unit time. It can be deduced experimentally by accumulating the aggregate of messages processed in all workstations during the measurement period, T, yielding $S = N * M/T$ (N: number of workstations; M: aggregate of messages in bits per workstation in period T; T: length of measurement period).

The response-throughput characteristics for read and write access for all configurations are shown in Figures 4 and 5 respectively. It is observed that all curves exhibit similar shapes and trends. Typically the throughput shows an initial rise, with little increase in delay until a peak value is reached. Thereafter it decreases rapidly accompanying a sharp rise in response time. Eventually it reduces to an asymptotic value with the delay growing without bound. The initial rise in throughput signifies the ample capacity of the server. As the traffic load increases to a threshold value the throughput reaches a maximum, indicating that the input rate matches the maximum capacity of the server. Beyond that, the request builds up rapidly, and soon becomes excessive, resulting in high delay and decreased throughput. In some cases the initial rise is missing. The throughput of the single station assumes the maximum value.

The maximum throughput indicates a state of optimum performance of the server. It varies with the types of configuration and modes of operation. For ease of comparison the maximum throughputs of all networks are extracted and displayed in Figure 6. Of all the message sizes, Novell AN286 scores the highest maximum throughput value of 2.4 Mbit/sec for 6 kbyte read access. It also scores a large maximum throughput value of 1.5 Mbit/sec at 750 byte read access. Compared with others, the high throughput at 750 byte read access is about three times higher than the best throughput of 3Server, seven times that of PCAT Server, and 21 times that of PC network. It is obvious that in the region of small message size, Novell AN286 dominates.

However, as message size increases beyond 6 kbyte the maximum throughput drops rapidly. At the large message size of 30 kbyte it assumes a value equivalent to that of 3Com 3+ 3Server. As for the 3Com networks, the distribution of maximum throughputs is rather consistent throughout the whole range of message sizes. The behaviour of IBM PC network is somewhat different: maximum throughput increases with message size. It gives the best throughput at the large message size, matching with that of 3Com 3+ PCAT server. A distinct display of two clusters is shown at the message size of 30 kbyte.

Another informative display is the variation of throughput as a function of the number of stations (shown in Figure 7 for write access only: the behaviour for read access is similar and thus omitted). For the Novell and 3Com networks, the curves at small message sizes are characterized by a broad high throughput region spanning a number of stations. However, as the message size becomes large the peak shrinks in size and width to a bare hump at one or two stations. It indicates that more stations can be supported in the region of small message size.

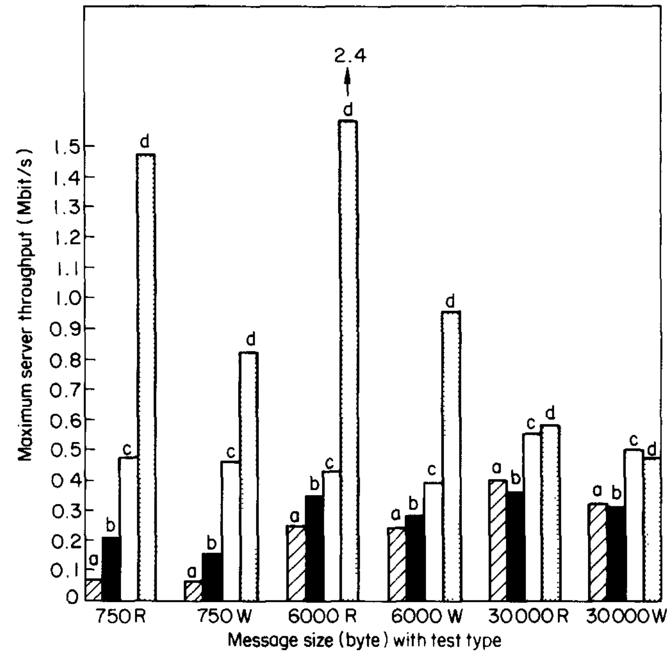

Figure 6. Maximum server throughput for various network configurations. R: read; W: write. (a) IBM PC network; (b) 3Com PCAT server; (c) 3Com 3Server; (d) Novell AN286

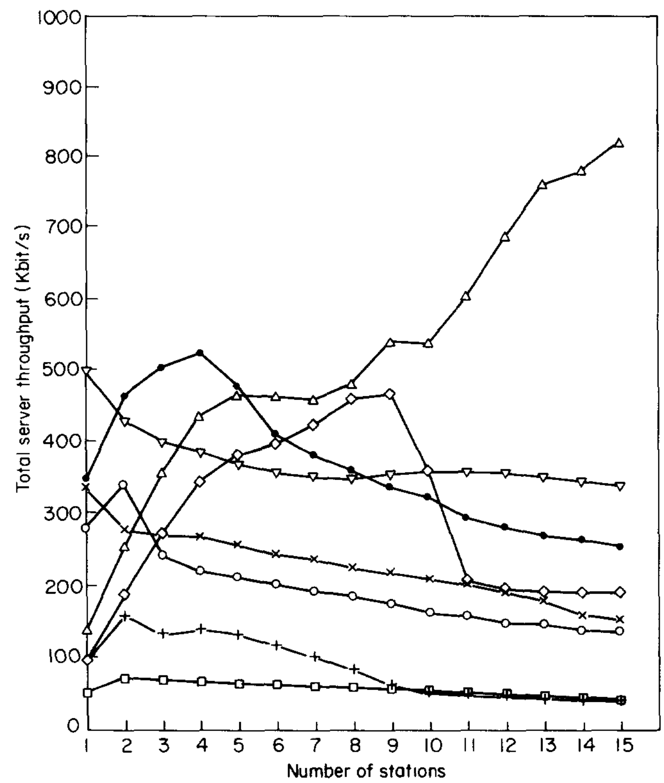

Figure 7. Throughput-station characteristics of various networks for 750 and 30 kbyte write operations. □: IBM PC network 750 byte; ○: IBM PC network 30 kbyte; +: 3Com PCAT server 750 byte; ×: 3Com PCAT server 30 kbyte; ◇: 3Com 3Server 750 byte; •: 3Com 3Server 30 kbyte; △: Novell AN286 750 byte; △: Novell AN286 30 kbyte

Floppy disc equivalent

The response time of read/write access of a local floppy disc on a standalone IBM PC is familiar to many. It can be used as a standard reference to derive the number of effective stations to be supported in a network, each of which yields a response equivalent to that of a local floppy disc. Obviously, the higher the number of stations supported the better the network.

The response time of read/write access of a local floppy disc can easily be measured. By comparing the response with that of a network an equivalent number of stations supported can be deduced. A linear extrapolation is made for the case where the number of effective stations exceeds 15. This may cause some uncertainty in the estimates. The results are displayed in Figure 8. Not surprisingly Novell AN286 supports the highest number of stations, from 104 to six stations for read access and from 58 to four stations for write access. The 3Com 3+ 3Server supports about 17 to five stations. The IBM PC network supports the least, only five to three stations. The merits of the networks are clearly demonstrated.

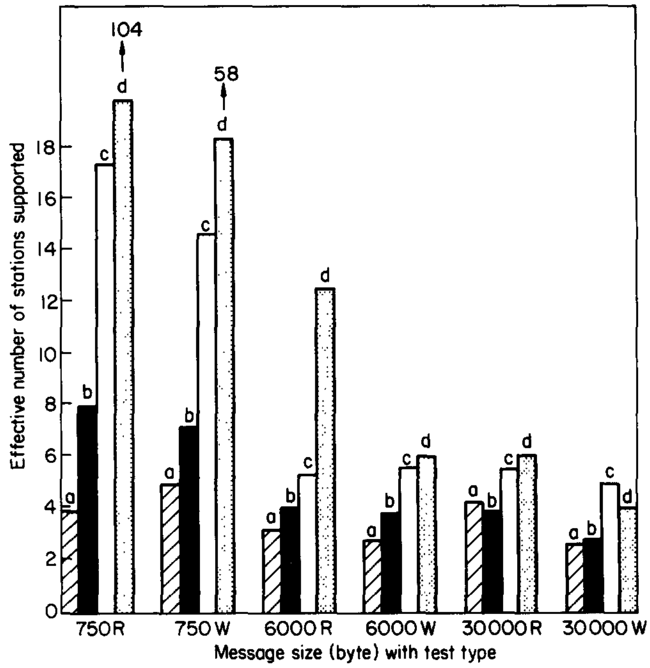

Figure 8. Effective number of stations supported as derived from the floppy disc equivalent for various network configurations. R: read; W: write. (a) IBM PC network; (b) 3Com PCAT server; (c) 3Com 3Server; (d) Novell AN286

Performance penalty

An alternative measure is the performance penalty of a system from which one can derive the relative strength of the number of stations that a system can support. The performance penalty is defined as the increment of delay time obtained when an extra user is added to the network. It gives a constant value if the delay variation is linear. However, for nonlinear variation in general, a mean performance penalty will be a good indicator of the merit of the system.

Using the measured data the mean performance penalty can be calculated from the response time increment. The results for the configurations tested are shown in Figure 9. They show that Novell AN286 has the smallest mean performance penalty and IBM PC network has the largest. It also shows that the penalty for read access is less than that for write access. There is a general trend that the performance penalty rises with the message size.

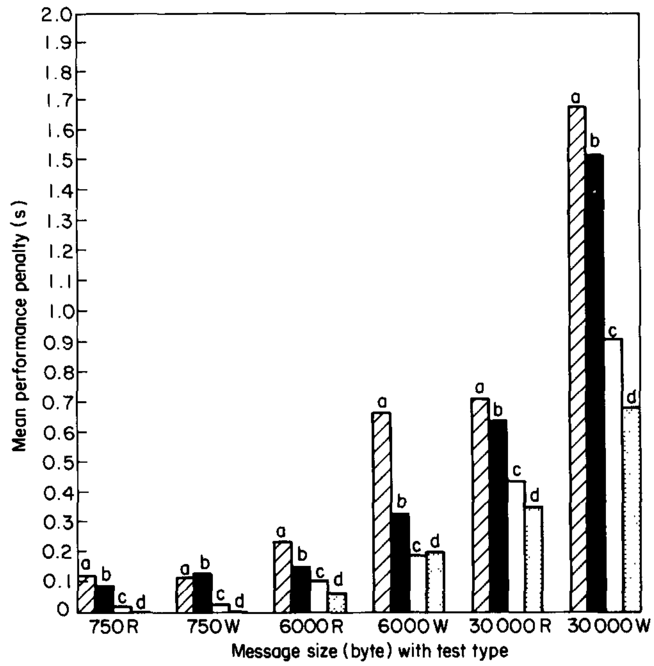

Figure 9. Mean performance penalty per extra user for various network configurations. R: read; W: write. (a) IBM PC network; (b) 3Com PCAT server; (c) 3Com 3Server; (d) Novell AN286

The relative strength of the number of stations that each server can support can be derived from Figure 9 by taking the inverse ratio of the mean performance penalty. The ratio varies with server type and operation mode. For simplicity the read/write modes of operation are grouped together, and only an average value is taken. Two distinct message sizes are considered as they are different. Thus, the following ratios are derived:

Novell AN286 : 3Com 3Server : 3Com PCAT : IBM PC network
= 68 : 51 : 10 : 1 for small message size
= : 2 : 1 : 1 for large message size

It can be deduced that for small message size, Novell AN286 can support 1.3 times as many stations as the 3Server, 6.8 times that of PCAT server and 68 times that of IBM PC network. Similarly, the 3Com 3Server is 5.1 times better than the PCAT server and 51 times better than the PC network. On the other hand, for large message size the networks can be grouped into two distinct classes, with Novell AN286 and 3Com 3Server in one class and 3Com PCAT and IBM PC network in another. Performance in the same class is comparable, but between classes the former performs twice as well on the latter. This is consistent with previous conclusions. It should be remembered that this is a mean value estimate with 15 stations under heavy load conditions. It does, nevertheless, give an indication of the relative merits of the various systems.

NETWORK STABILITY

Network testing under heavy load conditions also serves as a means of examining system stability. Surprisingly, of all the systems tested, the IBM PC network is the most stable, despite the fact that it gives the poorest performance. A number of problems have been encountered with the other networks.

For the 3Com network read/write large message size, the server might be too busy to respond immediately to station requests. The following message would appear at the station:

Network retrying.... Retry can be stopped by pressing ⟨CTRL⟩ ⟨BREAK⟩

The retry continues with the message appearing every so often. This is harmless for PCAT server as the read/write requests are still being processed. However, it was a catastrophe for 3Server: the congestion caused the 3Server to freeze with an error message 'ETH TRAP 370B' displaying on its LCD screen, thereby terminating all read/write access. This problem was later overcome by an upgrade of 3Server motherboard to a higher version together with a reconfiguration of server parameters.

For the PCAT server, 'Error reading drive e: Abort, Ignore or Retry ?' would appear during the read/write access of large message size. This is a timeout problem. When a station waited too long to communicate with the server, the message would pop up and terminate the read/write operation. Fortunately the error occurred only intermittently without causing much damage, and it was then resolved by manual intervention.

A more serious problem for the PCAT server is the abnormal behaviour of the 1 500 byte read access. An alarmingly high response time was recorded. This is shown in Figure 10, where the response time at 1.5 kbyte is higher than that of 39 kbyte. A further test was made by varying message size from 1 350 to 1 650 byte with one and 15 stations. The result confirms the excessive delay. It is difficult to offer any logical explanation other than attributing to the bug of server software for its poor handling of the 1.5 kbyte message size which is on the boundary of fitting into the length of an Ethernet packet.

In comparison, Novell AN286 is the least stable network, though it gives the best performance. The response time distribution among stations shows large fluctuation at small message size. It indicates that the server is not fair to all stations. Some stations are being served rapidly, while others have to queue for a long time. As an illustration, a plot of the normalized mean response times for various stations is shown in Figure 11. The variations are represented by the normalized standard deviations, as indicated by vertical bars, and the normalized maximum and minimum response times, as shown by symbols. The former is calculated as a range from the highest point to the lowest point given by

$1 \pm$ (standard deviation/mean response time).

The latter is evaluated as the ratio of the maximum (minimum) response time among the measured stations to the mean response time of all measured stations. As shown in Figure 11, the variations are as high as $\pm$ 40% at some stations. This behaviour can be attributed to the disc caching mechanism implemented in the server software. In the implementation, the disc read process reads information from the disc in large sections well beyond the data requested by a station, in anticipation of further read requests. A station will experience a long delay if its read request results in a physical disc access. However, once the data is available in the cache memory, the response will be very fast. Similarly, for the write access the cache buffer accumulates multiple write requests until it is full or timeout expired, which then triggers off a physical disc write access. A quick response results if the data is deposited in the cache buffer: otherwise, the delay will be long if a physical disc access is

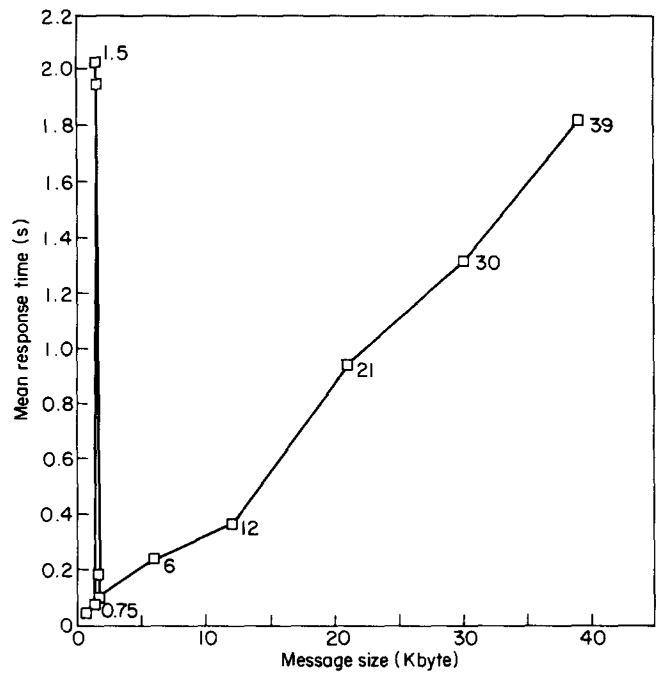

Figure 10. Abnormal behaviour of 3Com PCAT server at 1.5 Kkyte read access with one station

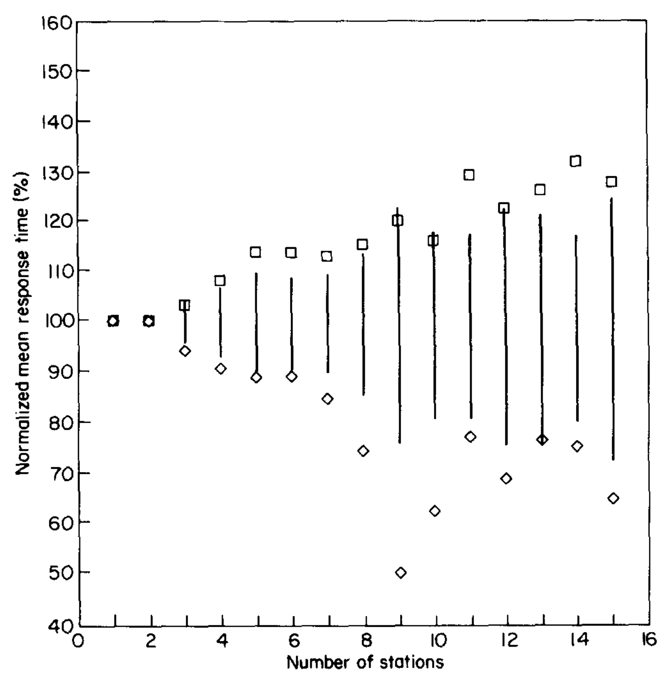

Figure 11. Variation of normalized mean response time/station for 750 byte write operation. □: maximum response time; ◇: minimum response time; (———): standard deviation

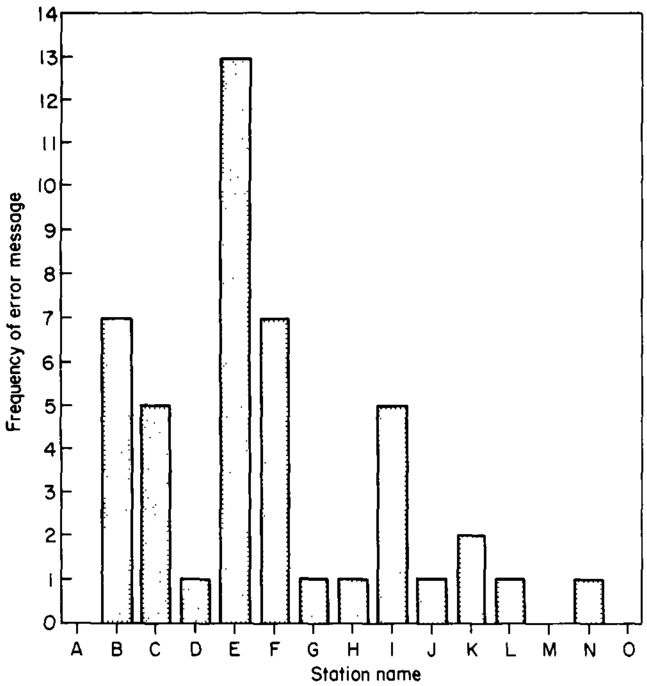

Figure 12. Frequency of error messages for individual station

activated. The uneven response causes a large fluctuation in delay time. The behaviour is sensitive to small message size, but not to the large one, because the cache buffer is always full for large message size, and each request results in a number of physical disc accesses.

A more disruptive problem for Novell AN286 is the abrupt termination of read/write process due to some unknown network error, yielding the message:

Network Error on Server SERVER1:
Error writing to network, Abort or Retry?

The error sets in frequently beyond seven stations and becomes excessive when the message size is large. Figure 12 shows the frequency of occurrence at individual stations for the entire measurement involving 0.75, 6 and 30 kbyte. It is interesting to note that station E has an unusually large number of error messages, whereas stations A, M and O have none. Obviously the error distribution is bias. It is probably a characteristic of the server software to give preferential treatment to some stations but not others. When the error occurred there was no indication on the server's monitor, which only showed that the affected stations were not reading/ writing. The cause of the error remains unknown: it is probably a software bug. Nevertheless, its occurrence causes much disruption to the measurement. To circumvent the problem, a retry routine is incorporated into the timing program so that the process can be reactivated whenever the error sets in. This is a fix rather than a solution.

CONCLUSION

By and large, the comparative performance evaluation carried out in this study produces the following performance ranking: Novell AN286 best; 3Com 3Server and PCAT server intermediate; and IBM PC network poorest. However, with regard to the network stability, the ranking is the reverse with the IBM PC network the most stable, and the Novell AN286 the least.

The comparative features can be specifically outlined. Mean response time increases monotonically with the number of stations and message size. The total server throughput generally exhibits an initial rise to a peak value before the decline. The maximum throughput is large for small message sizes for Novell AN286 and 3Server, though it is high for large message sizes for PCAT server and IBM PC network. The read access takes a much shorter time to accomplish the task than the write access, resulting in a higher throughput than the write access. The special server can perform better than the general purpose server. This is true when comparing the 3Server with the PCAT server for the 3Com network.

Finally, it can be concluded that the performance of the file server is insensitive to the network topology (be it the bus or tree), transmission technology (be it baseband or broadband) and raw data rate (be it 10 Mbyte/s or 2 Mbyte/s). It is, however, very sensitive to the implementation of the file server software. The fact that the server software of Novell running on PCAT can outperform the special 3Server of 3Com in the region of small message size signifies the importance of server software design. This conclusion is derived from networks which are of CSMA/CD type. It would be interesting to know if the result is extendable to cover other networks based on token ring or token bus technology.

By and large, the efficiency of server software depends on the file serving (disc access) mechanism, as well as network communication protocols. There exist a number of techniques for performance enhancement of file servers[10]. Some of the techniques have been implemented in the networks under study. A commonly implemented technique is the disc, caching, which performs prefetching and postwriting operations, thus reducing the numbers of physical disc access. Nevertheless, different emphasis has been made in the implementation. The caching in Netware favours random disc access, whereas that in 3Plus and PC LAN program favours sequential disc access. In addition, Netware, unlike the other two, does not use DOS for file serving functions. It replaces DOS and incorporates directory hashing and elevator seeking techniques[11] in the software. The former enables a fast retrieval of directory. The latter allows the ordering of requests according to the part of disc accessed, thus minimizing disc head travel. The measures account for the superior performance of Netware over the others, as shown in the experimental data.

The development of LANs follows the rapid advancement of chip technology. The arrival of Intel 80386 will undoubtedly trigger the production of more and more powerful server hardware. However, if the server software is incapable of matching the power of the hardware, it is unlikely that it will achieve a quantum leap in the performance of file serving operation.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the assistance of Miss Bee Sim Neo, Miss Poh Tiang Sim and Miss Guay Lee Teo for carrying out some of the measurements. They also wish to acknowledge the helpful comments of the referees.

REFERENCES

1 **Stieglitz, M** 'IBM provides industry with a versatile local network standard' *Data Communications,* McGraw Hill, USA (June 1985) pp 195-206

2 **3Com Corporation** *Theory of operations for the 3Plus product family* (1986)

3 **3Com Corporation** *3Server Administrative's Manual* (1986)

4 **Novell Corporation** *Advanced Netware User's Reference Manual* (1986)

5 **Sunshine, C A and Ennis, G** 'Broadband personal computer LANs' *IEEE J. on Selected Areas in Commun.* Vol SAC-3, No 3 (May 1985) pp 408-415

6 **Dixon, R C, Strole, N C and Markov, J D** 'A token-ring network for local data communications' *IBM Syst. J.* Vol 22 No 1/2 (1983) pp 47-62

7 **IEEE Standards for Local Area Networks** *802.3: CSMA/CD Access Method; 802.5: Token Ring Access Method* IEEE Press, USA (1985)

8 **Noorda, R** 'Transportable operating system promotes network benchmarks' *Mini-Micro Syst.* (May 1984) pp 161-172

9 **Cowart, R and Feldmann, P** 'Benchmarks for network ratings' *PC Magazine* (February 1985) pp 148-156

10 **Tanenbaum, A S** *Operating Systems* Prentice-Hall, USA (1987)

11 **Hurwicz, M** 'Evaluating LAN operating systems for microcomputers' *Data Commun.,* McGraw Hill (November 1986) pp 203-215