Analysis of Disk Workloads in Network File Server Environments

John R. Heath*
Department of Computer Science
University of Southern Maine
Portland, Maine 04103

Stephen A.R. Houser
Academic Computing Services
University of Southern Maine
Portland, Maine 04103

Workloads of network file server disk IO subsystems have very different characteristics than observed in time sharing or local IO systems described in the literature. In this study we provide a detailed analysis of disk workload traces collected from network file servers. Our results characterize file server disk traffic and give insights into file server disk access patterns. Measurements and statistics presented in this paper will aid designers and managers in developing and tuning disk subsystems for network file servers. Moreover, our results can be used by analysts to parameterize synthetic workloads for server subsystem studies.

1 Introduction

LAN-based distributed operating systems often employ a client-server design in which one or more network hosts, called network file servers, provide a global file system that is shared by client workstations. Workstations submit file service requests to a file server which processes them and replies to the workstations over the LAN via a simple network protocol. Client workstations may be "diskless", have no local disks, in which case a file server provides the workstations' only file system, or workstations may have local disks. Workstations with local disks may have local file systems and only access the server's file system for some application programs or to access shared data files.

To be effective, a file server must provide performance and reliability levels comparable to that of local file systems. Adequate performance levels are achieved by devoting a large portion of server memory to file caching. Recently accessed files or parts of recently accessed files are stored in the server's memory. This memory region is referred to as a *file cache*. Read IO requests for records stored in the cache are serviced directly from the file cache without accessing the server's disk storage, thereby substantially reducing response time. A request for data not in the server's file cache requires that the data be transferred from disk storage to the server's memory before being transferred to the client workstation.

Considerable performance gain is achieved if requested file data is stored in the server's cache, but performance degrades with frequent accesses to the server's substantially slower disk storage. The adverse effects on server performance of frequent disk accesses can be alleviated by improving disk subsystem performance. To improve disk subsystem performance, file caches are often placed in IO subsystem adapters, disk controllers, or both. A variety of read-ahead and write-behind algorithms have been implemented or proposed for managing IO caches. The most cost-effective choice of caching algorithms, cache parameter settings, cache size and location depends on IO traffic characteristics which differ in different environments.

To design effective data management algorithms requires a thorough understanding of file server access patterns. Knowledge of disk access patterns is also necessary for tuning file cache parameters to achieve optimal performance. In this study, we provide a detailed analysis of disk workload traces collected from two network file servers. Our purpose is to provide insight into disk access patterns in network file server environments. We intend that insights gained from this study will be useful in designing and tuning file server disk subsystems. Also, study results can be used to develop server disk workload models which are useful in designing and parameterizing analytic models. Workload models can also be used to create synthetic workloads for simulation and benchmarking. Finally, the traces provide real workloads for trace-driven simulations and for measuring and testing actual systems.

Traces of disk workloads described in the open literature focus primarily on file systems in time sharing environments or local file systems in distributed environments. File IO traces and disk access patterns of local file systems in VAX/VMS cluster environments are described in [BISW90, RAMA92]. These traces are used in [BISW93] to drive simulation models to assess the performance of non-volatile write caches for a variety of write-behind strategies, and are also used in [KARE94] to evaluate cache replacement algorithms. Disk workload traces of a BSD UNIX-based time sharing system are described in [OUST85]. These traces

* This work was supported in part by Digital Equipment Corporation, Shrewsbury, MA.

are used in a simulation model to evaluate the performance of main memory caches in diskless workstations [NELS88]. Smith [SMIT85] characterizes file system traces of three large IBM mainframe systems and uses the traces to investigate a number of disk cache design issues. In [BODN91], workload characteristics of client workstation requests to a file server are discussed, but IO subsystem workload characteristics are not presented. Because of a file server's large main memory file cache, distributed operating system, and large number of diverse server users, server disk workloads are different from those characterized in other studies and may very well suggest different solutions to improving and tuning file server subsystem performance.

2 Trace Methodology

2.1 Trace Environments

Disk workload traces analyzed in this study were collected in April 1994 over a period of several days from two file servers on different LANs. One server provided the file system for seventy diskless workstations connected to a general access LAN used primarily by university students from all academic disciplines. The most frequently used server applications were word processing, accounting for 60% of all connect time, spreadsheet programs, accounting for 13% of connect time, and communications services, 10% of connect time. Other applications included programming, MS-DOS, database, and courseware applications. The LAN is a hub-configured Ethernet with workstations interconnected by twisted-pair cable. The LAN's single file server runs Novell NetWare and has 16 megabytes (MB) of RAM, of which, we estimate, 75% is used for file cache. The server uses two SCSI disk drives to support its file system. A 1.2 gigabyte (GB) drive [FUJI93] primarily stores application programs, temporary files, and shared data files. The drive employs a 256 kilobyte (KB) read-ahead cache. The second drive [IBM93] stores operating system files. It has a 320MB capacity and a 128KB read-ahead cache.

The second set of traces was obtained from a network file server used by administrative staff in several university departments. Client workstations have local file systems and use the server for shared applications, primarily word processing, and access to shared data. During working hours, the number of workstations connected to the server varies from 50 to 75. The administration LAN is also hub-configured Ethernet and the file server has essentially the same configuration as the student server.

We note that the two networks are quite similar except that the student LAN workstations are diskless while the administration LAN workstations have local disks. Consequently, comparison of the servers' disk workloads and access patterns should provide insight into the effects of using diskless workstations as compared to workstations having local disks.

2.2 Trace Tools

A SCSI bus monitor [PEER93], installed in an MS-DOS PC, was used to collect workload traces. The monitor was attached to the server's SCSI bus and recorded all commands transmitted on the bus. Specifically, for each transmission, the monitor recorded the target device, the logical block number (LBN) of the beginning of the request, the length of the request in 512-byte blocks, the requested IO operation, and times of beginning and end of transmission. File data were not recorded. Records were stored in the monitor's 3MB memory. During periods that the bus was idle, trace records were transferred to the host for storage. Trace data was interpreted and analyzed after the tracing was completed. The monitor was passive and did not affect disk subsystem performance.

At each site, traces were collected on consecutive days of a typical work week. The specific hours traced varied somewhat from day to day, but included hours that the LANs were actively used. For comparison purposes, we selected for analysis trace segments of nine continuous hours during which the servers were most active. Specifically, traces used in this study were collected on the administrative server from 8:00am until 5:00pm for 5 days, Monday through Friday[1]. Student server traces are for the period from 9:00am until 6:00pm, Monday through Friday. The results presented in this paper, unless otherwise stated, were derived for each server by combining the server's daily traces.

3 Disk Workload Characteristics

In this section, we present file server disk workload measurements characterizing server disk traffic. We expect our results will give insights into server disk access patterns and aid designers and managers in developing and tuning IO subsystems for file servers. Also, the results can be used to build and parameterize analytic models of file server IO subsystems.

3.1 Traffic and Operation Characteristics

Recall that each of the measured server file systems includes two disk drives. On each system, one disk is used primarily to store application programs, shared data, and user files. These disks will be referred to as application disks. The second disk on each system is used to store system files; we refer to these as the system disks.

Both servers' system disks exhibited more activity than did the application disks with the greatest disparity occurring in the diskless workstation system. Nearly 80% of the student server's IO traffic was to its system

[1] Due to a system failure, the Tuesday trace did not begin until 9:30am

disk. Traffic was more balanced in the administration LAN where sixty percent of disk accesses were to system files on the system disk. We attribute this difference in system disk workloads to the fact that the administrative file server serves workstations with local disks that can be used for storing temporary files, paging, and storing some system binaries. The student workstations are diskless and must access the server for all file activity, requiring frequent system file accesses.

Most, 72%, of the IO operations issued by the student server were writes. Of the operations issued by the administration server, 50% were writes. Most write activity, by both servers, was to the system disks. Eighty-nine percent of student server writes were to the system disk; at the administration server, 74% of writes were to the system disk. Read requests were more evenly balanced between the two drives; 53% of student server reads were to the system disk, while 46% of the administrative servers reads were to the system disk. IO activity to the servers' disks is summarized in Tables 1 and 2.

Operation	System	Application	Total
write	64%	8%	72%
read	15%	13%	28%
total	79%	21%	100%

Table 1: Percentages of read and write IOs submitted by the student network file server.

Operation	System	Application	Total
write	37%	13%	50%
read	23%	27%	50%
total	60%	40%	100%

Table 2: Percentages of read and write IOs submitted by the administration network file server.

Read-to-write ratios for the servers' disks are given in Table 3. We observe in Table 3 that more read requests than write requests were submitted to the application disks and more writes than reads submitted to the system disks. Given the types of files stored on these disks, this is to be expected. Table 3 also shows the ratio of blocks read to blocks written. Because read requests were much larger than write requests, we observe that, for all disks, more blocks are read than written. The ratio of blocks read to written is 8:1 for both application disks and, somewhat lower, 3:1 and 2:1, for the system disks.

For each disk drive, we derived read-write ratio fluctuations during each day of the week. Specifically, for each week day, at 10 minute intervals, we computed read-write ratio for the past 30 minutes. A plot of the derived values provides a measure of read-write ratio changes that occur during the day. For each drive, we obtained five plots, one for each week day.

Server	Application		System	
	requests	blocks	requests	blocks
Student	8:5	8:1	1:4	2:1
Admin	2:1	8:1	5:8	3:1

Table 3: Read-to-write ratios, IO requests and blocks transferred.

Fig. 1 shows the daily read-write ratio curves for the student server system drive, the most active of the traced drives. Read-write ratio curves for the other drives are given in [HEAT95]. As noted above, this drive exhibits more write than read activity. Most days, between 5:00 and 5:30, we observe an increase in read activity. We attribute this to the fact that student-LAN users do not store their permanent files on the server's file system. Consequently, we speculate that increased read activity is caused by diskless workstation users copying files to their floppy disks and closing applications as they prepare to leave for the dinner hour. At times, other than this half-hour period, computed read-write ratios range between 1:2 and 1:6 with an occasional further increase in write activity, reaching a maximum read-write ratio of 1:12.

The daily read-write ratio curves for the student server's application disk requests ranged, for most half-hour periods, between 4:1 and 1:4 with most ratios above 1. There were infrequent periods with a high percentage of write IOs, reaching a maximum of 16 writes for each read, and periods with a high percentage of read IOs, reaching as many as ten reads for each write.

Read-write ratios in administration system drive requests ranged from 3:2 to 1:5. Except for two spikes, one reaching a ratio of 1:20 and the other reaching 9:1, read-write ratios at the administration application drive varied between 5:1 and 1:5, similar to those observed on the student server application drive.

File server read requests are in multiples of 4KB (eight 512-byte disk blocks). In fact, all disk read requests submitted by the administration server were for eight 512-byte blocks. Read IOs submitted by the student server were of three lengths, 8, 16 or 24 blocks. Ninety percent of reads submitted by the student server to its system disk were for 24 blocks. Forty-seven percent of all reads to the application disk were for 24 blocks; 46% were for 16 blocks.

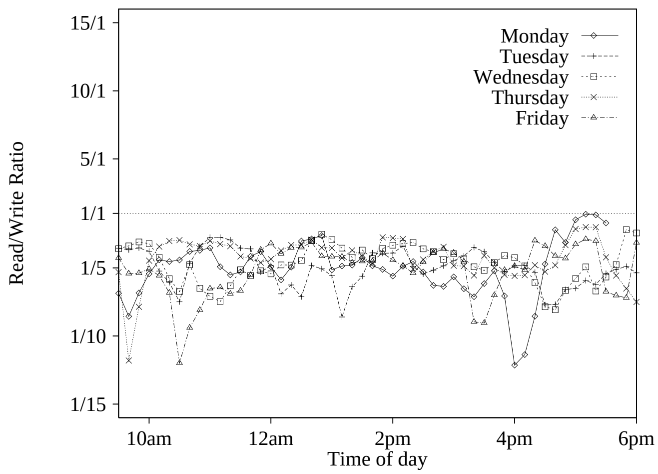

Figure 1: Graphs of read-write ratios as they fluctuate throughout each day, M-F, student server system drive.

Most write requests were for one block. Sixty-seven percent of write requests submitted by the student server were for one block, 91% of which were to the system disk. Of the writes submitted by the administration server, 78% were for one block and, of these, 75% were to the system disk.

Request length distributions and operation frequencies should be helpful in designing storage system caches; particularly so, in determining cache size, optimal cache line size, and the efficacy of partitioning cache into read and write caches. Complete IO length distribution histograms are given in [HEAT95].

3.2 Locality of Access

Figs. 2 and 3 show the frequency with which each disk cylinder was accessed. In the figures, disk cylinders are indicated along the x-axis; for each cylinder, the height of the associated bar is the frequency with which the cylinder was accessed. We observe that, generally, a high percentage of the disk accesses are to a small number of cylinders. Fig. 2a shows the cylinder access frequency of writes to the student server application disk. We observe that 36% of all writes access one cylinder and 83% of all writes are to only one percent of the cylinders. Forty-one percent of writes to the student server's system disk are to two consecutive cylinders; see Fig. 2b. Furthermore, 99.8% of writes access the first one-third of the disk's cylinders. The administration server's disks also exhibit a high frequency of writes to a small number of cylinders, Fig. 3. Thirty-five percent of writes to the application disk are to one percent of the cylinders. Thirty percent of write accesses to the system disk are to four consecutive cylinders and 54% are to only one percent of the cylinders.

Read requests are more widely distributed over the disk cylinders; read frequency histograms are not shown. System disks, in particular, exhibit read activity over the entire disk. However, in both environments, much of the read activity was limited to a small number of cylinders. Thirty-one percent of reads to the student server's application disk were to one percent of the cylinders; 19% of reads to the student system disk were to one percent of the cylinders. At the administration disk, 44% percent of application disk reads accessed only one percent of the cylinders and 18.5% of reads to its system disk accessed one percent of the cylinders.

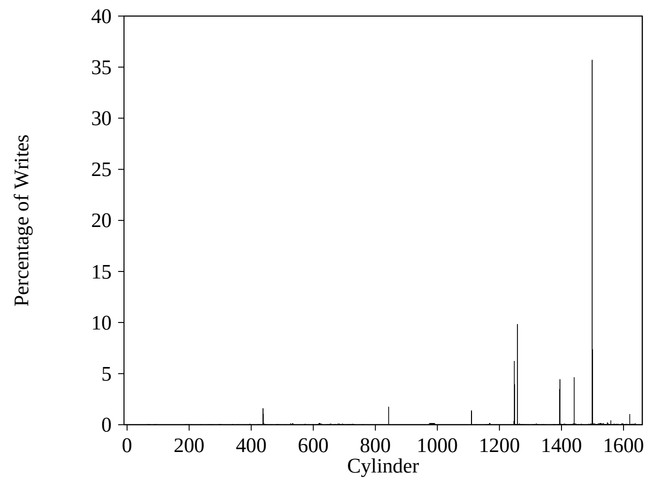

(a) application drive

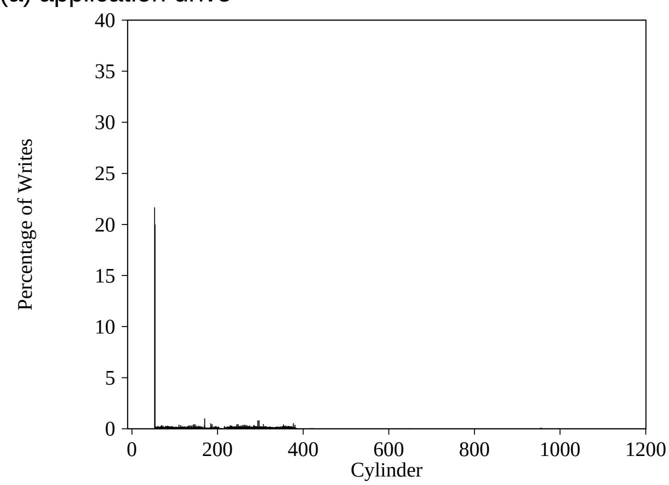

(b) system drive

Figure 2: Frequency of cylinder accesses by write operations, student server drives.

We observed that a high percentage of requests required no seek. At the student server, 57% of accesses to its application disk required no seeking and 50% of system disk accesses did not require seeking. At the administration disk, 46% of application disk IOs and 36% of system IOs required no seeks. See [HEAT95] for seek distance distributions.

Both servers employ a *lazy-write* algorithm whereby disk write requests are delayed by the server before being written to the disk. The algorithm's purpose is

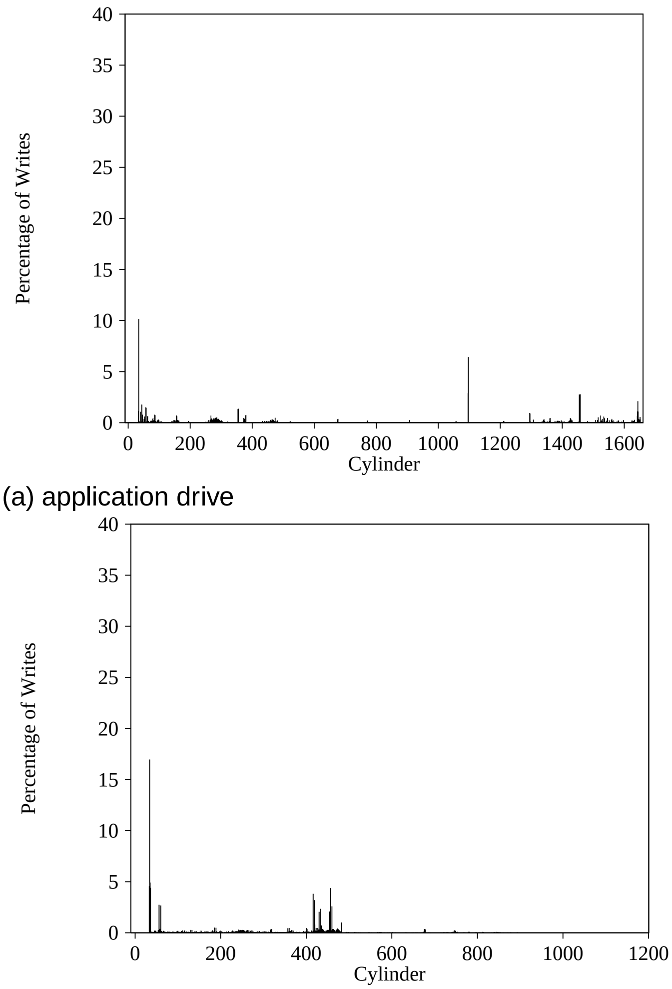

(a) application drive

(b) system drive

Figure 3: Frequency of cylinder accesses by write operations, administration server drives.

to reduce disk writes by processing multiple writes to cached data without writing to disk, and by collecting several writes to consecutive blocks into a single disk write. During the trace periods, this delay was set to the default, 3.3 seconds for file records and 0.5 seconds for directory writes. Our measurement tools do not measure the reduction in disk writes resulting from this algorithm; however, our disk traffic measurements show a high percentage of one-block writes to the same cylinder. A possible explanation of this behavior is that the lazy-write algorithm was ineffective in reducing writes to the servers' disks. Further study will be needed to determine conclusively the effectiveness of the lazy-write algorithm. However, given the high degree of locality and the large amount of write traffic, it seems that increasing write delay in file cache would improve write IO performance, at the risk of increased data loss should the system fail.

Daily Throughputs				
Weekday	Student		Administration	
	IOs/sec	KB/sec	IOs/sec	KB/sec
Monday	11.9	42.0	2.7	5.8
Tuesday	9.3	35.6	3.9	9.3
Wednesday	13.4	53.2	2.7	6.2
Thursday	11.3	48.3	2.3	5.3
Friday	5.5	20.4	3.0	7.3
Mean	10.3	39.9	2.9	6.8

Table 4: Daily throughputs for the two servers.

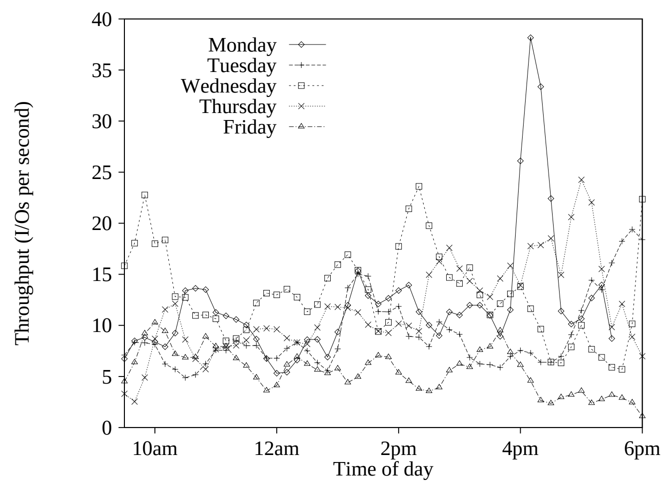

Figure 4: Daily throughput fluctuation, student server.

3.3 Throughput

Daily throughputs for each server are shown in Table 4. Throughputs shown are for nine hour periods during which the servers were most active, from 9:00am to 6:00pm for the student server and from 8:00am to 5:00pm for the administration server. The results indicate much lower activity at the administration server than at the student server. The student server's mean daily throughput, excluding Friday which has relatively low activity, is 11.5 IOs/sec and, including Friday, is 10.3 IOs/sec. Of the total student server throughput, the system disk contributed about 80%. The much less active administration server's mean daily throughput was 2.9 IOs/sec.

Fig. 4 shows the fluctuation in student server throughput for each day. Each point plotted on the curves shows measured throughput for the past 30 minutes and is calculated every ten minutes. Most plotted throughput values, for all days, except Friday, range between 5 and 25 IOs/sec, although Monday's throughput for one 30 minute period reached 38 IOs/sec. On Friday throughputs are generally lower, ranging from 5 to 10 IOs/sec until late afternoon when throughput remains below 5 IOs/sec.

Table 5 shows daily throughputs for each of the stu-

Student Server				
Weekday	Application		System	
	IOs/sec	KB/sec	IOs/sec	KB/sec
Monday	2.3	12.6	9.6	29.4
Tuesday	2.1	13.0	7.2	22.6
Wednesday	2.8	16.3	10.6	37.0
Thursday	2.2	17.0	9.1	31.3
Friday	1.3	8.1	4.1	12.0
Mean	2.1	13.4	8.1	26.4

Table 5: Daily throughputs for student server disks.

Administration Server				
Weekday	Application		System	
	IOs/sec	KB/sec	IOs/sec	KB/sec
Monday	0.8	2.1	1.9	3.7
Tuesday	1.3	4.0	2.6	5.3
Wednesday	0.9	2.9	1.9	3.8
Thursday	1.0	2.8	1.3	2.5
Friday	1.2	3.8	1.8	3.5
Mean	1.0	3.1	1.9	3.8

Table 6: Daily throughputs for administration server disks.

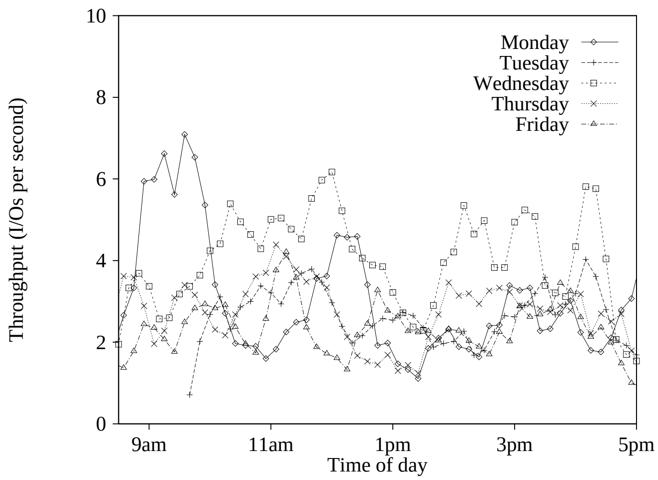

Figure 5: Daily throughput fluctuation, administration server.

dent server's drives. The lower activity of the student application drive is apparent. Application drive throughput, measured during the day at 10 minute intervals as described above, generally remains below 4 IOs/sec with only an occasional increase above this rate; although, for one half-hour period on one day, application disk throughput reached 23 IOs/sec.

Fluctuations in daily student server system disk activity are difficult to characterize as there are no clear patterns. On three days, maximum half-hour throughput exceeded 20 IOs/sec; the times of day that high throughput periods occurred were different each day. Times at which the server was busiest one day were periods of low activity on other days. On some days throughput increased more or less monotonically throughout the day, whereas on other days maximum throughput was reached by mid-morning.

Administration server throughputs for each day are shown in Fig. 5. Not surprisingly, low activity levels generally occur early in the morning, at lunch time, and late afternoon; at these times, throughput is often as low as 1 IO/sec or less. Administration server half-hour throughput levels reach a maximum of about 7 IOs/sec with throughput for most periods ranging from 2 to 4 IOs/sec, except during the periods noted above. Table 6 shows the daily throughputs for each of the administration server's disk drives. We note that the system drive contributes about 65% of the server's total throughput.

3.4 Response Time

Response time distributions for read and write operations to the four disks are shown in Figs. 6-9. The distributions are for the most active nine hour periods, 9:00am to 6:00pm and 8:00am to 5:00pm for the student and administration servers, respectively. We observe in Figs. 6 and 7 that, in both environments, a large percentage of read IOs have relatively low response times. The low response times are indicative of disk controller cache hits and demonstrate the effectiveness of disk cache in reducing read response time. At the administration server's application disk, over 45% of read IOs appear to be cache hits and, at the student server's application disk, nearly 40% of read IOs are apparently read hits; see Fig. 6. Similarly, in Fig. 7, we observe a large percentage of low read response times further demonstrating the effectiveness of disk controller cache.

Write response time distributions are shown in Figs. 8 and 9. Mean write response times to the two application disks are nearly identical, 16 ms with 7 ms standard deviation. Mean write response times to the administration system disk and the student system disk are 18 ms and 15 ms, respectively, Fig. 9.

The mean response time for all write operations issued by the student server is 15 ms and the mean read IO response time is slightly more, 18.5 ms. At the administrative server mean response times for read and write operations are 14 ms and 18 ms, respectively. Mean response time for all operations issued by the student server is 16.2 ms with standard deviation of 7.9 ms. For the administration server, mean response time for all operations is 16.0 ms with 8.7 ms standard deviation.

3.5 Heavy Traffic Characteristics

IO performance during heavy loading conditions is of particular interest to system designers. To acquire bet-

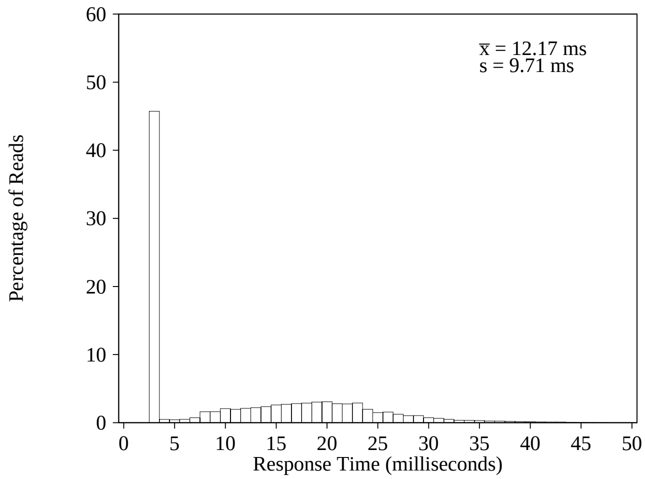

(a) administration server, application drive

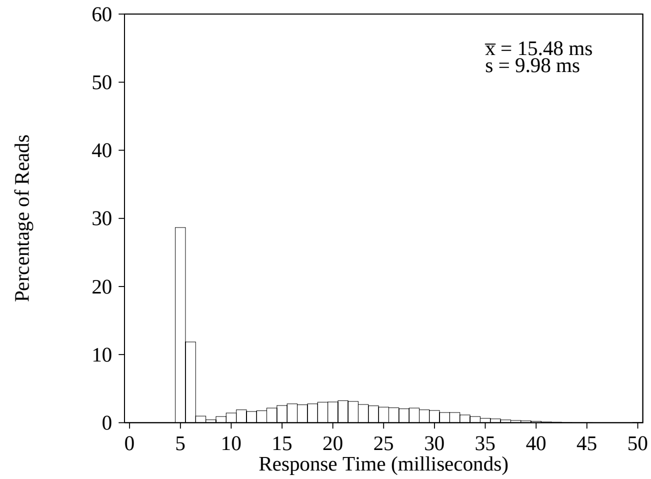

(a) administration server, system drive

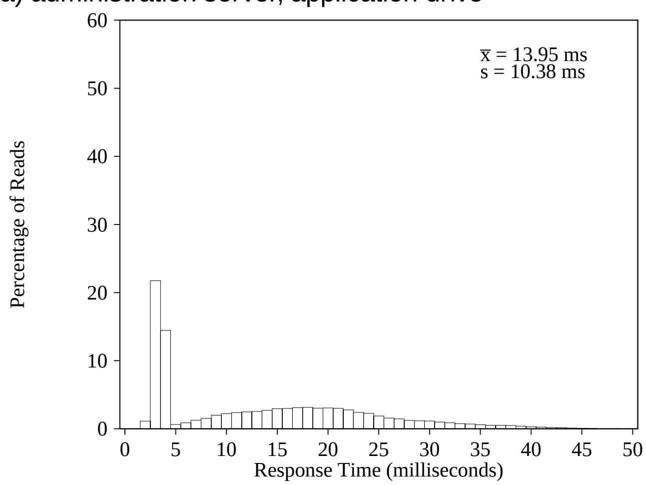

(b) student server, application drive

Figure 6: Read response times, application drives.

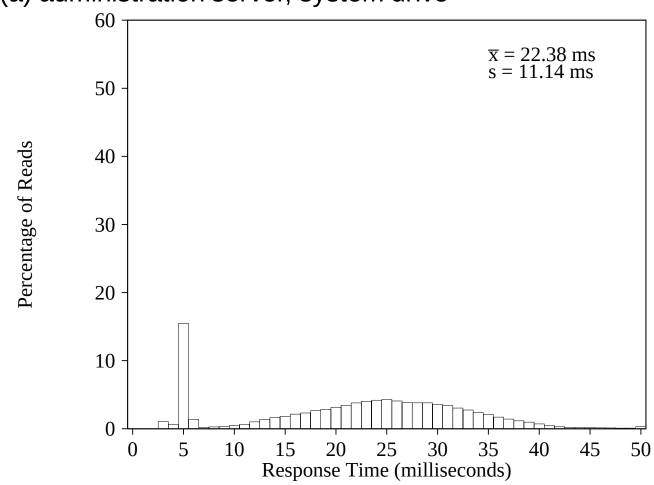

(b) student server, system drive

Figure 7: Read response times, system drives.

ter understanding of disk IO workloads with heavy traffic, we identified five time periods in the traces as having high disk activity and collected performance statistics for these periods. We selected as our heavy load periods trace segments during which 30 minute throughputs exceeded 20 IOs/sec. Then we further restrict high activity periods to be the time from the point when throughput first exceeds 20 IOs/sec until throughput falls below 20 IOs/sec for the the final time during the selected busy segment. The selected busy periods and performance measurements for the periods are described in Table 7. All but one of these high activity periods were observed on the student server system disk; one heavy workload trace, labelled in Table 7 as Wed-A, was collected on the student application disk.

Disk throughput during these busy periods is shown in Figs. 10-15. Note, the figures include throughput for short periods before and after the busy segment. Points on the figures represent throughput for the past minute and are computed each minute. On all but one figure, throughput reaches a maximum of about 40 IOs/sec; on Monday, however, throughput reaches 60 IOs/sec and remains between 50 and 60 IOs/sec for 15 minutes. We see from the figures that the heavy trace periods display dissimilar patterns of disk activity. During some periods, high throughput is sustained during the busy time, while during other periods throughput rapidly fluctuates between high and low levels during the busy period. The Monday and Tuesday curves represent the two extremes.

In the last column Table 7, we note that during heavy load periods the percentage of disk accesses not requiring a seek ranges from about 50% to nearly 80% which is generally higher than we observed over the entire trace; in the overall trace, 57% of all accesses to the student server system disk required no seeking. Further examination of these high activity trace segments shows that, as observed in the overall traces, much of the traffic was restricted to only two consecutive cylinders. Specifically, 35% to 50% of disk IOs during these

Measurement Period			Mean Response Time(ms)	Throughput IOs/sec	Write/Read Ratio	%-Zero Seek Access
Day	Start	End				
Mon	3:48	4:08	14.9	47.9	13.9	75.4
Tues	4:50	5:57	16.5	14.8	5.6	56.9
Wed-A	5:46	6:01	14.5	31.7	15.0	79.3
Wed-B	9:21	10:10	17.2	17.8	3.9	53.7
Wed-C	1:42	2:31	17.8	18.3	3.4	47.1
Thurs	4:02	4:57	17.2	19.8	2.9	52.8

Table 7: Performance statistics, heavy traffic periods.

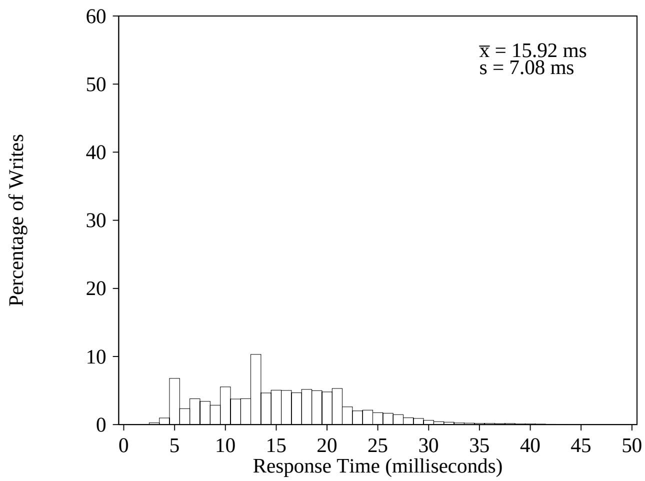

(a) administration server, application drive

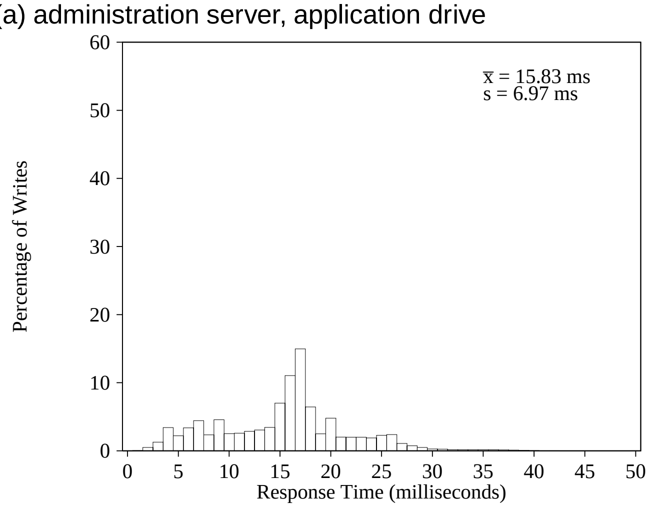

(b) student server, application drive

Figure 8: Write response times, application drives.

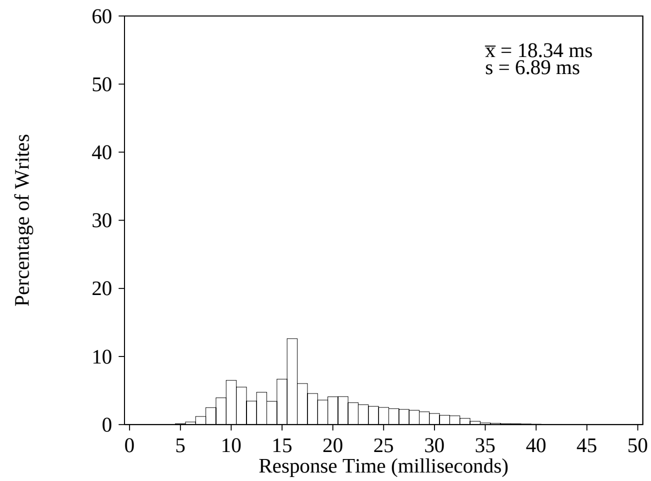

(a) administration server, system drive

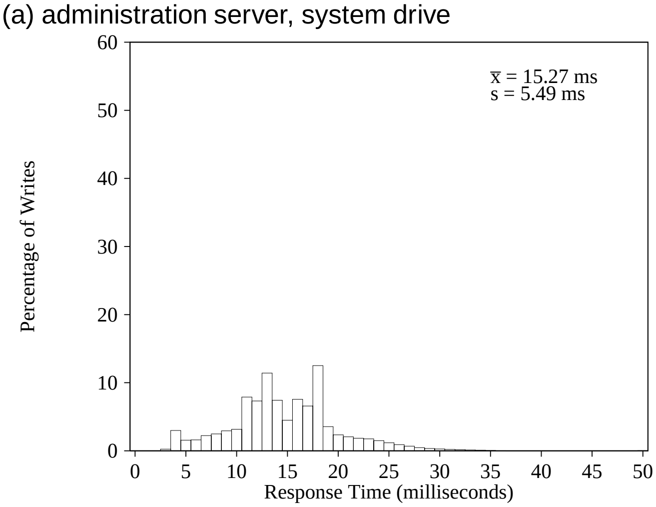

(b) student server, system drive

Figure 9: Write response times, system drives.

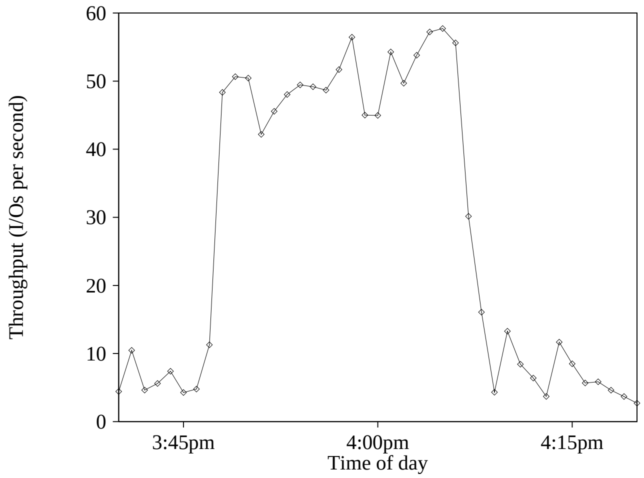

Figure 10: Throughput computed at one minute intervals for the busy period labelled MON.

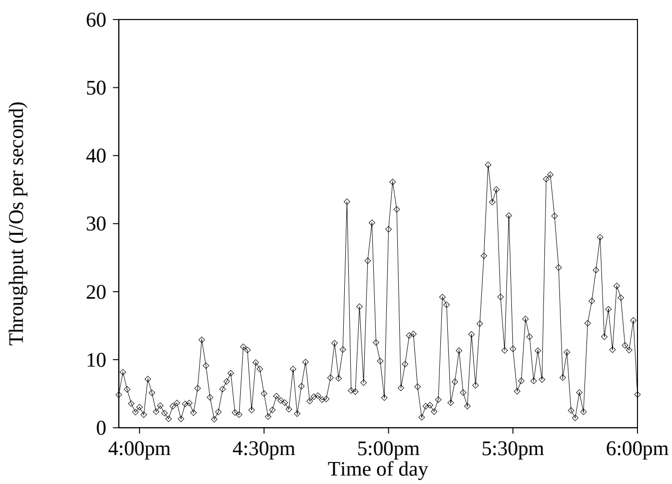

Figure 11: Throughput computed at one minute intervals for the busy period labelled TUES.

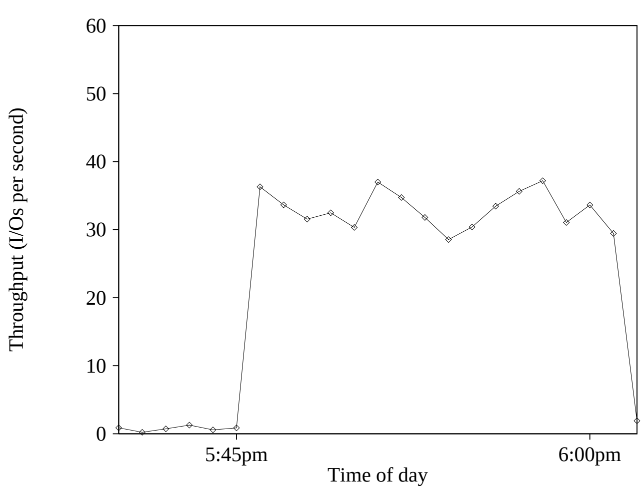

Figure 12: Throughput computed at one minute intervals for the busy period labelled WED-A.

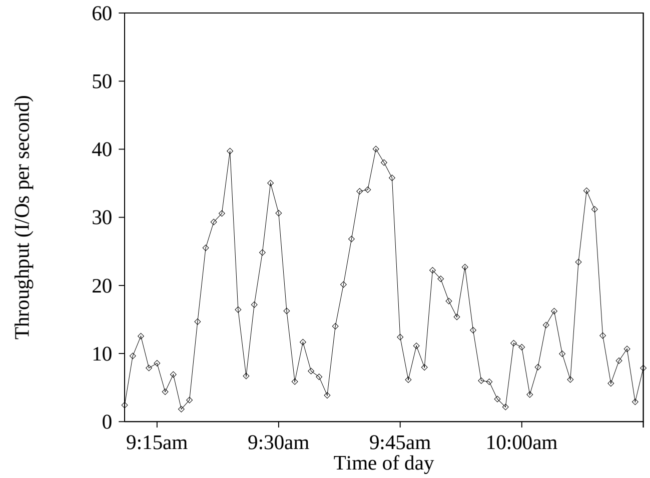

Figure 13: Throughput computed at one minute intervals for the busy period labelled WED-B.

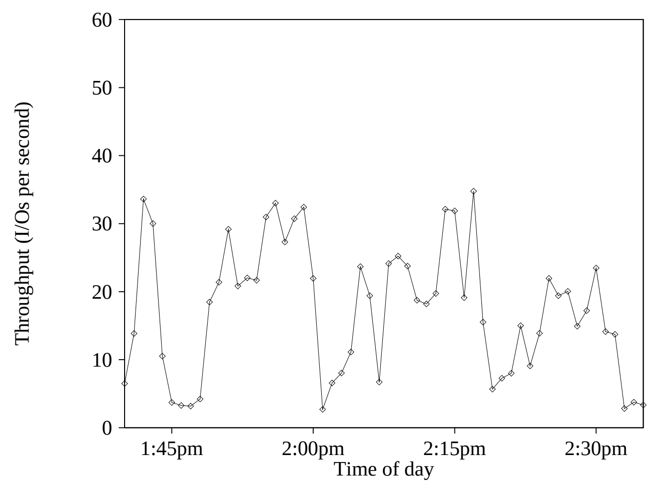

Figure 14: Throughput computed at one minute intervals for the busy period labelled WED-C.

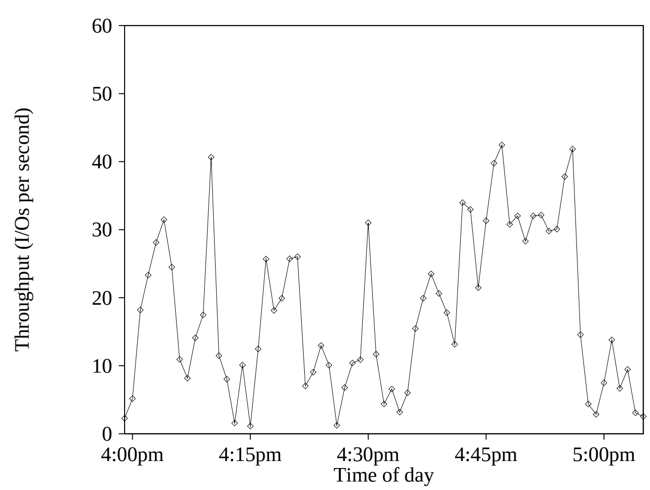

Figure 15: Throughput computed at one minute intervals for the busy period labelled THURS.

periods were to two cylinders.

During busy periods, the ratio of write operations to read operations ranges from 3:1 to 15:1, as shown in Table 7, with most ratios for system disk traffic similar to the overall trace ratio of 4:1, shown in Table 3.

4 Concluding Remarks

We have analyzed disk workload traces for two network file servers in different LAN environments. Our measurements and derived statistics provide qualitative as well as quantitative information on server disk traffic. Furthermore, the results can be used to parameterize synthetic workloads for server disk studies. We present statistics of request length, seek distance, and cylinder access frequency, and we characterize fluctuations in read-write ratio and throughput that occur during the day.

File server read-ahead algorithms are designed to take advantage of the locality of client file references. Specifically, we expect read requests to locations near recently accessed records to be processed by the server file cache, and, as a consequence, we might expect to observe a low degree of locality in disk accesses. Nevertheless, our seek distance and cylinder access frequency histograms imply significant locality in disk accesses as well. The observed high level of locality and measured read response time distributions demonstrate that disk controller caching in a file server environment is effective in reducing read response time.

The observed high percentage of single block writes and the high percentage of write operations overall indicate that further subsystem performance gain could be achieved by focusing on improving write performance; specifically, the performance benefit of adding write cache would seem worthy of further study.

The two traced systems have nearly identical configurations, except one system uses diskless workstations while the other has workstations with disks. Comparing the workloads of these two systems gives insight into the effect this difference has on server disk usage. In particular, in the diskless workstation system, a much higher percentage of IOs were directed to system files and a much larger percentage of the IOs were writes. The length of the server disk requests were larger in the diskless workstation LAN. The student server disks were more active than the administration server disks. The ratio of mean daily throughputs for the two servers during the trace periods is 3.5:1.

References

[BODN91] R. Bodnarchuk and R. Bunt. A Synthetic Workload Model for Distributed System File Server. In *Proc. 1991 ACM Sigmetrics*, pages 50–59, May 1991.

[BISW90] P. Biswas and K.K. Ramakrishnan. File Characterizations of VAX/VMS Enviroments. In *Proc. 10th International Conference on Distributed Computing Systems*, pages 227–234, May 1990.

[BISW93] P. Biswas, K.K. Ramakrishnan, and D. Towsley. Trace Driven Analysis of Write Cacheing Policies for Disks. In *Proc. 1993 ACM Sigmetrics & Performance*, pages 13–23, June 1993.

[FUJI93] Fujitsu Computer Products of America. *Fujitsu M2266A 1.2 GB Disk Drive*, 1993.

[HEAT95] John R. Heath and Stephen A. Houser. “Disk Access Patterns in Network File Server Environments”. Technical Report TR 95-5, Dept. of Computer Science, University of Southern Maine, May 1995.

[IBM93] IBM Corporatrion. *IBM OEM Storage Products 0661 Model 467*, 1993.

[KARE94] R. Karedla, J. Love, and B.G. Wherry. Cacheing Strategies to Improve Disk System Performance. *IEEE Computer*, 27(3):38–46, March 1994.

[NELS88] Michael N. Nelson, Brent B. Welch, and John K. Ousterhout. Cacheing in the sprite network file system. *ACM Trans. on Computer Systems*, 6(1):134–154, Feb. 1988.

[OUST85] John K. Ousterhout and et.al. A trace driven analysis of the UNIX 4.2 BSD file system. *Proc. Tenth Sympos. on Op. Sys. Princ.*, pages 15–24, Dec. 1985.

[PEER93] Peer Protocols Inc. *Peer Protocol SCSI Analyzer*, 1993.

[RAMA92] K.K. Ramakrishnan, P. Biswas, and R. Karedla. Analysis of File I/O Traces in Commercial Computing Enviroments. In *Proc. 1992 ACM Sigmetrics & Performance*, pages 78–90, June 1992.

[SMIT85] Alan J. Smith. Disk-cache – miss ratio analysis and design considerations. *ACM Trans. on Comp. Sys.*, pages 161–203, August 1985.