LBL-8277 c.2
UC-32

RECEIVED
LAWRENCE
BERKELEY LABORATORY

MAY 31 1979

LIBRARY AND
DOCUMENTS SECTION

LONG TERM FILE MIGRATION - PART I: FILE REFERENCE PATTERNS

Alan Jay Smith

August 1978

Prepared for the U. S. Department of Energy
under Contract W-7405-ENG-48

Thi py
whi weeks.
For all
Tec

LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA
BERKELEY

LBL-8277 c.2

LEGAL NOTICE

This report was prepared as an account of work sponsored by the United States Government. Neither the United States nor the Department of Energy, nor any of their employees, nor any of their contractors, subcontractors, or 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.

LONG TERM FILE MIGRATION - PART I:

FILE REFERENCE PATTERNS*

Alan Jay Smith

University of California, Berkeley**
Lawrence Berkeley Laboratory
Berkeley, California 94720

August, 1978

CR Categories: 4.35, 4.33, 4.41

Keywords: File Migration, Replacement Algorithm, Time Series Analysis, Memory Hierarchies, Mass Storage.

ABSTRACT

In most large computer installations, files are moved between online disk and mass storage (tape, integrated mass storage device) either automatically by the system or specifically at the direction of the user. This is the first of two papers which study the selection of algorithms for the automatic migration of files between mass storage and disk. In this first part, we examine the use of the text editor data sets at the Stanford Linear Accelerator Center (SLAC) computer installation through the analysis of thirteen months of file reference data. We find that most files are used very few times. Of those that are used sufficiently frequently that their reference patterns may be examined, we find that (a) about a third show declining rates of reference during their lifetime, (b) of the remainder, very few (about 5%) show

correlated interreference intervals, and (c) interreference intervals (in days) appear to be more skewed than would occur with the Bernoulli process. Thus about two-thirds of all sufficiently active files appear to be referenced as a renewal process with a skewed interreference distribution. A large number of other file reference statistics (file lifetimes, interference distributions, moments, means, number of uses/file, file sizes, file rates of reference, etc.) are computed and presented. The results of our analysis of file reference patterns are applied in the following paper to the development and comparative evaluation of file migration algorithms.

---

*Partial support for this research has been provided by the National Science Foundation under grants MCS75-06768 and MCS77-28429 and by the Department of Energy under contract W-7405-ENG-48 to the Lawrence Berkeley Laboratory. Computer time has been provided by the Department of Energy under contract EY-76-03-0515 to SLAC.

**Computer Science Division, EECS Department, University of California, Berkeley, Ca. 94720. The author is also a staff member of the Lawrence Berkeley Laboratory (University of California) and a visitor at the Stanford Linear Accelerator Center.

I. INTRODUCTION

Almost all computer installations (excluding hobby computers) employ a memory hierarchy much like that in figure 1. Each of the levels of storage from cache to mass storage is successively larger, slower and less expensive per bit. By

Figure 1

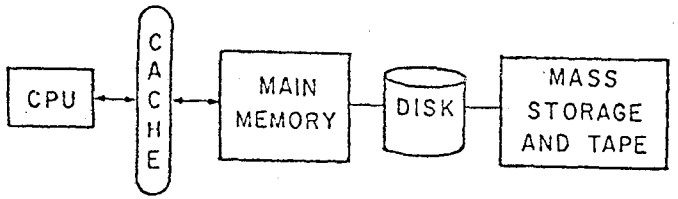

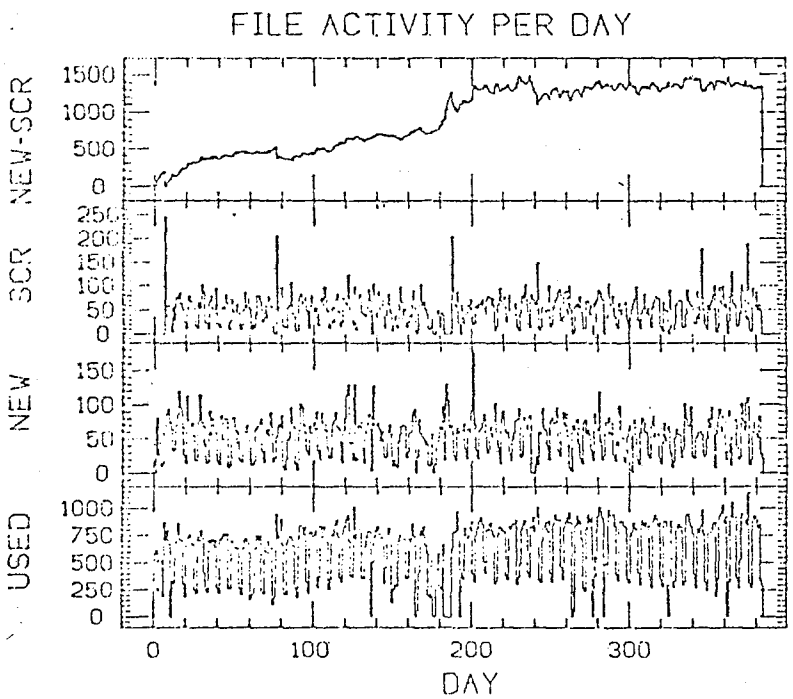

Figure 2

Figure 3

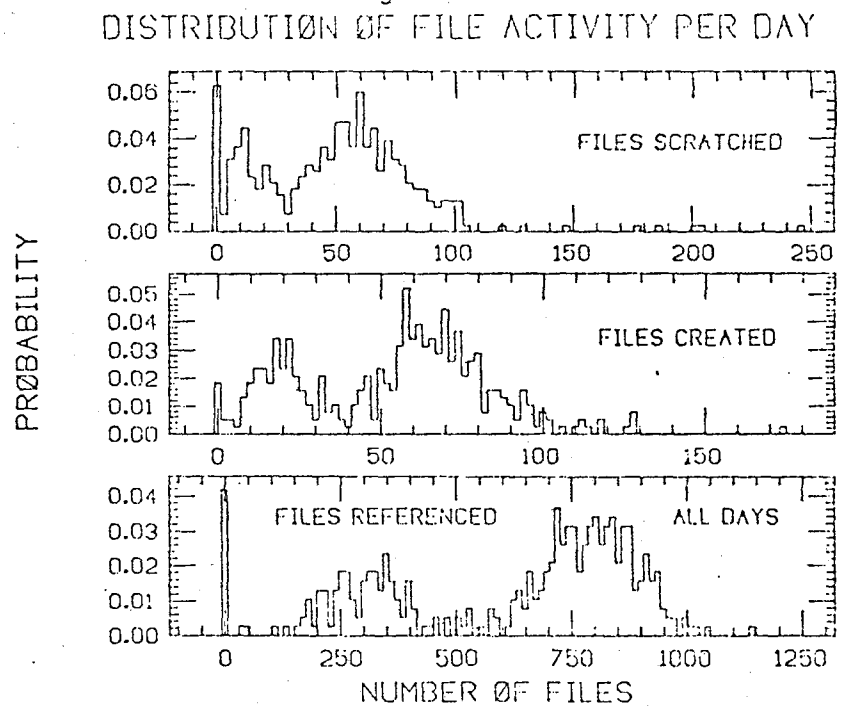

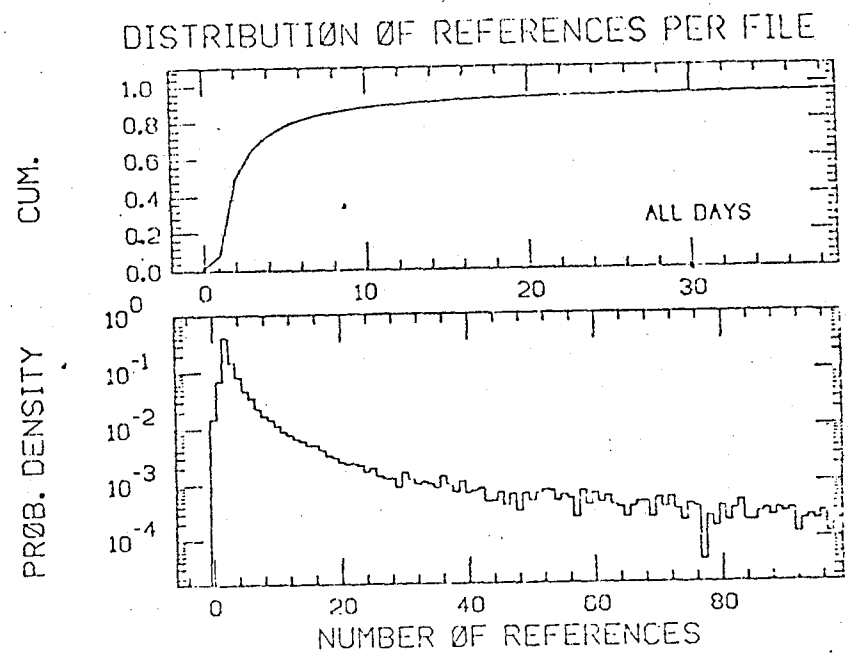

Figure 4

dynamically moving information between levels of the hierarchy, the system can usually arrange to have each level capture a vastly larger fraction of all memory references than the levels below it. In effect, the user sees a system in which the total storage capacity is the combined capacity of all levels of the memory, while the average access time is very close to that of the fastest. The success of such dynamic information movement derives from the empirically observed "principle of locality" (Denning, 1972) which essentially states that (a) information in recent use is likely to be reused and (b) information logically adjacent to recently used information is likely to be referenced soon.

An important factor in the effectiveness of this dynamic movement among the levels of the hierarchy is the selection of algorithms for when/where to move the information. Considerable study has been devoted to the movement of information to/from cache memories (Conti, 1969, Kaplan and Winder, 1974, Smith, 1978a, 1977a, 1976a) and main memories (Denning, 1970, Smith, 1976b, 1978b). The problem of the transfer of files between mass storage and disk has been largely neglected, however, almost certainly due to the lack of suitable data. The only previous useful study of this problem was by Stritter (1977) who studied the same data as is analyzed here. In a companion paper to this (Smith, 1978c), we derive and evaluate a number of file replacement algorithms. A more complete review of the literature and discussion of the problem is also provided in that article.

In this paper, we study and analyze data on the use of user text editor data sets at the Stanford Linear Accelerator Center over a period of thirteen months. We look at questions

of how and when files are used. Some statistical analysis is performed on file reference patterns in order to determine what stochastic process models might be appropriate to represent the individual or aggregated file reference patterns. A large variety of statistics are gathered and presented, but particular attention is paid to those aspects of file reference behavior that relate to construction of effective file migration algorithms. Throughout, our emphasis is on "data derived" models and algorithms. The intent is to give the reader a good overall understanding of the file reference process, both as an aid in the development of file migration algorithms and for a general understanding of how users use the file system.

The next section of this paper describes in some detail the nature of our data and its limitations. The actual tabulation, charting and analysis of this information is presented in section III, which comprises the bulk of this paper. We provide an overview of our findings and discuss some of the implications of our measurements in the conclusions section.

II. DATA DESCRIPTION

It is very important to understand both the nature of the data that we analyze and the type of system from which it was collected. As will be evident below, our information does not include some parameters of interest; thus we are unable to make some studies. The utility of our results will be constrained by the similarity between the system we discuss and the system to which the results are to be applied. In this section (II), a relatively detailed description is given

of the nature of the data and the system from which it was taken.

The data that we analyze in this paper was collected at the Stanford Linear Accelerator Center by Dr. Edward P. Stritter. His analysis of this data appears in his doctoral dissertation (Stritter, 1977); our analysis here goes well beyond that already presented and reaches different conclusions in some cases. The measurement period was 384 days beginning August 2, 1974. Our discussion below, although phrased in the present tense, is descriptive of the system at that time; there have been changes since.

SLAC has a large computer system, consisting of two 370/168-Is running VS2 release 1.6 and an IBM 360/91 running MVT release 21.8. These three machines are loosely coupled using ASP and share most of the I/O devices. The 360/91 has been in service since 1968; the two 370's were added in the early Spring of 1974. The user interface has been essentially stable for several years preceding our measurements. Most of the computing activity is concerned with the study of high energy physics and the dominant programming language is Fortran. Most of the programmers are themselves either physicists or scientists in some closely related discipline.

The vast majority of user interaction with the computer system is accomplished through Wylbur (Fajman and Borgelt, 1973), which is an interactive text editor system with the capability to submit jobs into the batch job queue and later fetch the output. Each user with a Wylbur account is allocated a fixed number of 2314 Disk tracks (@7Kbytes/track) which can be used to store Wylbur data sets. Most users have allocations ranging from 10 to 200 tracks (70K to 1.4M bytes).

(For a standard of reference, a 2500 line Fortran program occupies 12 tracks.) Our data is limited to these Wylbur data sets, and excludes two important classes of files: system data sets (paging data sets, VTOC's, catalogs, etc.) and user files on tape. The size of most user space allocations constrains users to keep mostly computer programs in their Wylbur data sets; input to these programs (which often consists of huge volumes of data from accelerator experiments) is usually kept on tape. Further, we have no way of even following a file which has been moved offline, and if that file is later reestablished on disk, we are forced to treat it as a new file. Several scratch disks are used for temporary and staged (by the user) tape data sets; these disks are cleared every night and no record is available of their use.

Users are not charged for their disk space, but they are not allowed to exceed their allocation. The user is therefore inclined to ignore inactive data sets until the space is needed for some other purpose, at which time the file may or may not be copied to tape before being scratched. This particular accounting and allocation system undoubtedly affects user reference patterns, but without data taken from installations with some other scheme, we can only guess at the effect. There is no automatic file migration on the SLAC system, and therefore the users have no incentive to generate spurious file references in order to prevent automatic file migration programs from moving their files to mass storage. This is different from another system that has been measured (Revelle, 1975), and thus we believe that in that respect our data is "uncontaminated".

Information has been recorded for every Wylbur data set

which was seen to exist over the 384 day measurement period indicated above. For each file, one bit is recorded for each day indicating whether that file was used on that day. No indication is available as to how many times the file was used each day, nor whether the file was read or written. The date on which the file was created is available, as is the date on which the file was scratched, if it ever was. The name of the file, the user account ID and the file size (in tracks) was also recorded. Files can be used by other than the user that creates or owns them, but no record is available as to who used which file, nor is a record available to show whether the access was by an interactive user or batch job.

We have grouped files into three "classes". When a user session is suddenly interrupted (automatic logout, system crash), the system creates a file for the user containing his currently active Wylbur data set and saves it under the name "ACTIVE". Files with this name were considered to be one class.

Wylbur data sets can exist physically in one of two forms: a standard OS sequential data set or a partitioned data set (PDS). A PDS is usually used to hold a collection of small (less than one track) data sets, since one track is the quantum size for file allocation. Partitioned data sets therefore contain several individual and usually unrelated data sets, but our data shows only the access to the PDS as a whole. Further, our data does not directly distinguish PDS's from standard files. Fortunately, most users name their first and largest PDS "LIB" (for "library"). Those files whose first three characters were "LIB" were considered to be a class; clearly some unknown fraction of all PDS's were not

recognized. It is also possible (but unlikely) that some files which were not PDS's were considered to be such.

All files which were not placed in the class "active" or the class "library" were placed in the class "other" or "other files". As will be noted below, the three file classes have rather different reference patterns.

Files were also grouped according to their size. The size class is denoted "LSIZE" or "LOGSIZE" on the illustrations and is calculated from the logarithm base 2 of the file size in tracks. Thus file size class 0 consists of files of size 1 track, class 1 of files of sizes 2 and 3 tracks, etc., up to size class 6, which is all files of at least 64 tracks. Users can be expected to treat files of different sizes rather differently because large files occupy such a large fraction of a user's allocated space.

Because SLAC is a scientific shop located in the midst of an academic community, there is considerable user activity at night and on weekends and holidays. Never-the-less, activity is still a great deal lighter at those times than during first shift on weekdays. We have therefore studied file reference patterns and file migration algorithms both for the entire period of observation ("all days") and then for only the working days during that period ("working days"). In the case of "working day" analysis, all file activity during weekends or holidays was mapped onto the next following working day. Thus use of a file on both Saturday and Sunday would be treated as a use of the file on the following Monday. The logic behind this is the following: (a) file migration would most reasonably occur only on working days; i.e. a file migration program would run late at night on work nights and

move files off line as necessary, (b) by avoiding holidays and those days with missing data, some aspects of the analysis are simplified and/or improved and (c) data taken only on working days should be less affected by day of the week periodicities. This paper presents data for both cases (but with the stress on "all days" rather than "working days"). The companion paper considers working days only, since file migration would most reasonably occur only on working days.

III. DATA PRESENTATION AND ANALYSIS

A. Definitions and Symbols

In order to be clear about the meaning of our data presentation and analysis, in this subsection we define some terms and symbols. Let:

Nfiles be the number of files observed (24,898).

Nref be the number of file-day references. I.e. the sum over all files of the number of days that file was referenced (238,871).

Nref(i) be the number of days on which file i is referenced. (Reference to a file on a specific day will henceforth be called a "reference" to that file, although the file may have been referenced many times on that day).

I(i,j), i=1...Nfiles, j=1...384 is the indicator function for file i. That is, if file i is referenced on day j of the measurement period, I(i,j)=1, otherwise I(i,j)=0.

F is the date of the first day of measurement.

F(i) is the first day (within and relative to the measurement period) on which file i exists.

L is the date of the last day of measurement.

L(i) is the last day (within and relative to the

measurement period) on which file i exists.

B(i) is the date on which file i is created (birth).

D(i) is the date on which file i is scratched (death). If the file is never scratched, D(i) is undefined, except that it is known to be greater than L.

C(i,j), i=1...Nfiles, j=1...Nref(i)-1 is the sequence of interreference intervals for file i. That is, if file i is referenced on a total of Nref(i) days, C(i,*) will be the sequence of the number of days between uses. (Use on successive days yields an interreference interval of 1).

A(i,j), i=1...Nfiles, j=1...D(i)-B(i)+1 is equal to 1 if file i is referenced on day j-1 after its creation (valid only during the period of measurement) and is zero otherwise. A(i,1) = 1, A(i,D(i)-B(i)+1)=1.

B. Basic Numbers

Some of the basic numbers relating to our data are collected in table I; we also discuss them throughout this section. A total of 24,898 different files existed over the period of observation, and they belonged to 710 different accounts. (Not all of the accounts were associated with individual users, however). The average number of accounts showing activity per day was 183.5, the mean number was 214 and the maximum was 291. For an account showing activity on a given day, a mean of 3.41 files owned by that account were referenced. A total of 238,871 file-day-uses (references) $(= \sum_{i,j} I(i,j))$ took place.

Table II shows the file activity distributed by day of the week and holidays. We note that the level of activity was about 2 1/2 times higher during the week than on weekends. This fact will be relevant to some of our later discussion.

Table I

Basic Statistics

Total Days: 384
Working Days: 256
Number of files: 24,898
Number of accounts: 710
Number of file/day/uses: 238,871
Files on line initially: 3977
Files on line at end: 5320
Ave. number of files used/day: 622
Ave. number of files created/day: 54.48
Ave. number of files scratched/day: 50.98
Ave. number of tracks referenced/day: 11,727
Ave. number of tracks allocated/day: 351.7
Ave. number of tracks scratched/day: 343.5
Ave. volume of online files: 44,489 tracks
Ave. number of references/file: 10.6
Median number of referenced/file: 2
Average number of users/day: 183.5
Median number of users/day: 214
Max. number of users/day: 291
Ave. number of files used by logged on user per day: 3.41

Table II

Day	Total Files Used	Total Files Created	Total Files Scratched
Monday	40458	3643	3614
Tuesday	42532	3629	3254
Wednesday	43632	3843	3554
Thursday	40767	3694	3418
Friday	35714	3044	3009
Saturday	17061	1454	1304
Sunday	17004	1483	1322
Holidays	1703	131	103
Total	238871	20921	19578

Table III

Serial Correlation Coeficients

	All Days			Working Days		
	Used	New	Scr	Used	New	Scr
s(1)	.392	.269	.106	.315	-.122	-.111
s(2)	-.092	-.114	-.108	.295	-.114	-.056
s(3)	-.252	-.230	-.112	.224	-.122	-.120
sp(2)	-.289	-.199	-.120	.218	-.132	-.070
sp(3)	-.095	-.144	-.088	.087	-.159	-.137

Table IV

Number of References Per File

Number	Fraction	Cumulative
0	.015	.015
1	.069	.084
2	.411	.495
3	.152	.647
4	.084	.731
5	.047	.778
6	.035	.812
7	.023	.835
...	...	...
10	.011	.878
20	.0025	.930
50	.00048	.962
100	.00012	.980

Table V

Unweighted File Size Distribution

Size	Fraction
1	.391
2	.162
3	.078
4	.053
5	.053
6	.030
7	.021
...	....

Table VI

Mean/Median File Size

Weighted by	Mean	Median
Unweighted	7.08	1.7
Lifetime	8.99	2.3
Use	18.85	8.3

Table VII

File Size Range	Average File Size Weighted By Lifetime	Use	Unweighted
1	1	1	1
2-3	2.34	2.35	2.33
4-7	5.12	5.21	5.12
8-15	10.74	11.22	10.52
16-31	22.09	22.30	21.86
32-63	44.57	45.22	44.76
>=64	117.70	118.16	114.95
Total	9.02	18.85	7.08

C. Activity Over Time

Let: $Used(i) = \sum_k I(k,i)$ be the number of files referenced on day i,

New(i) be the number of files created on day i,

Scr(i) be the number of files scratched on day i and

NS(i) = New(i)-Scr(i) be the excess of new files over scratched files.

The four functions defined immediately above are plotted in figure 2. The mean values are given in table I; day of the week averages can be computed from table II. Immediately visible from figure 2 is the weekly periodicity; activity is generally high for five days and then low for two. The empirical distributions for Used(i), New(i) and Scr(i) appear in figure 3, where we see that the distributions each have three modes. The rightmost one reflects workday activity, the middle one weekend/holiday activity and the spike at 0 the days when the system was down or no data was collected (15 days). Plotting this same data for working days only (not shown) leaves only the single rightmost mode.

The number of tracks referenced is generally proportional to the number of files accessed. In table I we note that the average number of tracks referenced per day is 11,727, the average number of tracks in files allocated per day is 352 and the number of tracks in files scratched is 343.

As noted, the uppermost curve in figure 2 is the net accumulation of files. Over the period of measurement, there is an increase of 1343 files occupying an additional 3160 tracks. It appears from these numbers and the shape of the curve that an additional disk spindle was added during the middle of the measurement period. Disk space allocations for

users seem to have gone up gradually during this time. Comparing these figures to the 41,349 tracks occupied at the start of the measurement period (mean of 44,489 over the measurements), we see an increase of 7.6% in the total occupied file space. We note that in this system, increases in the number of files and occupied disk tracks should only occur when the capacity of the system is increased. This is in contrast to the system at IBM Research, San Jose (Revelle, 1975) where the file system is open ended and automatic file migration occurs. In that system a fairly steady accumulation of new files was noted, the rate being considerably greater than that observed here.

A factor affecting some uses of our data for some purposes is whether or not the level of activity on the system is stationary over time. We've already seen that the number of files and allocated tracks increased significantly over this period. We also choose to look at the time series Used(i), New(i) and Scr(i). We do this as follows:

First, we define the chi-square test for goodness of fit; this test will also be referred to later, so we make our definition general. (See Bickel and Doksum, 1977, for further information). Let y(i) be the empirical distribution of interest and let z(i) be the distribution against which it is being compared for fit. (Both y(i) and z(i) are discrete, i=1...k). Let n be the total number of (independent) samples used to form the empirical distribution. Then we compute:

$$\chi^2 = \sum_{i=1}^{k} \frac{n\left(y(i)-z(i)\right)^2}{z(i)} \qquad (1)$$

χ^2 should be distributed as the chi-square distribution with

k-1 degrees of freedom if y(i) is indeed drawn from the distribution z(i).

We let x(i) represent the time series of interest. (One of the three noted above either for all days, as shown in figure 2, or just for working days. Thus we have six time series.) We divide the total measurement period T into ten sections of duration T/10, and count the number of events y(j) of interest in each section. I.e:

$$y(j) = \sum_{i=(j-1)T/10+1}^{jT/10} x(i) \ , \quad j = 1 \cdots 10 \qquad (2)$$

y(j) was computed for all three series of events, for both all days and working days. (The all days test for stationarity is relatively worthless because of the nonuniform occurrence of holidays and days with missing data.) If the file activity were actually stationary, all values of y(j), j=1...10 should be about the same; specifically y(j) should be a good fit to the uniform distribution. The y(j) distribution was tested against the uniform distribution using the chi-square test for goodness of fit, and the hypothesis that y(j) is the uniform distribution was rejected in every case with better than 99.9% probability. This statistical test simply confirms casual observation - from the numbers it is clear that there is a generally higher level of activity during the second half of the measurement period. The change in the rate of activity is only moderate, however, and amounts to 10% to 20%. We therefore don't believe that this is likely to have a major effect on most of the remainder of our analysis and no effort has been made to correct the data for changes in the level of overall activity.

It is also interesting to determine to what extent user or system wide activity is correlated from day to day. That is, if the level of activity is high on day i, is it likely to be high on day i+1? The usual procedure for determining this from a nonstationary time series (such as we have) is to remove the trend from the data and then compute the serial correlation coefficient, since correlation coefficients are statistically meaningful only in the case of a stationary time series. Detrending the data is a laborious process which did not seem worthwhile, considering that this particular question is peripheral to the major thrust of this paper (but see Lewis and Shedler, 1976 for an example of such a procedure.) We did compute the serial correlation coefficient for lags of 1, 2 and 3 and also the partial correlation coefficient of orders 2 and 3 using the original nonstationary data. This was done for general information only and we make no claims for the statistical validity or meaning of the numbers obtained.

The expressions used to calculate these values are as follows: Let s(i) be the serial (or auto-) correlation of order i for a time series x(t), t=1...n, with mean $\bar{X}$. Let sp(i) be the partial autocorrelation coefficient of order i; that is, sp(i) is the correlation between elements x(t) and x(t+i) of the time series, having removed the effect of all correlations between x(t) and x(t+j) and x(t+j) and x(t+i), for 1<=j<i. Then

$$s(i) = \frac{\frac{1}{n-i}\sum_{j=1}^{n-i}\left(x(j)-\bar{x}\right)\left(x(i+j)-\bar{x}\right)}{\frac{1}{n}\sum_{j=1}^{n}\left(x(j)-\bar{x}\right)^2} \qquad (3)$$

$$sp(2) = \frac{s(2)-s(1)^2}{1-s(1)^2} \qquad (4)$$

$$sp(3) = \frac{s(1)^3 - 2s(1)s(2) + s(1)s(2)^2 + s(3)\left(1 - s(1)^2\right)}{s(2)\left(s(1)^2 - s(2)\right) - s(1)\left(s(1) - s(1)s(2)\right) + \left(1 - s(1)^2\right)} \qquad (5)$$

These expressions are drawn from Box and Jenkins (1976) and the reader is referred to that text for further discussion. We also consider the problem of estimating s(i) further in section III.F.2.b.

The values obtained from these expressions are presented in table III. Interpretation of these results should be done with care, since nonstationary data tends to yield spurious positive serial correlations. That is, if there are trends in the data, then consecutive samples are likely to both be in either a high or low trend and thus similar to each other relative to the mean. Therefore, we shall largely ignore positive serial correlations in this case; negative ones are probably meaningful. Looking at the working days results in table III, it is difficult to interpret the high positive correlations for the number of files used for the reason mentioned, although it makes sense to believe that if the system is busy today, it will be busy tomorrow. Periods of high system usage generally continue for several days for a given user and this should be reflected in these positive serial correlations. The negative correlations for New and Scr activity are much more meaningful, since they should not be produced by long term trends in the data. The author believes that this indicates a basically steady underlying rate of file creation or destruction. If a large number of files are created today, this implies that very few will be created tomorrow. Similarly for scratching files. The

correlations for the all days data is largely due to day of the week effects (e.g. days 1 appart are both likely to be either working or non working days; days 3 appart are not, etc.) and therefore aren't very useful.

D. Individual Files

Figure 4 shows the empirical distribution of the number of times each file was used during the measurement period. For some values, this distribution is tabulated in table IV. The average file is used 2 or fewer times but the average number of references per file is 10.6 (see table I). The distribution is thus highly skewed; most files are used very little and a few are accessed a large number of times. A note of explaination is needed for this data, however. It is detected that a file has been scratched by looking for it late at night (since it existed 24 hours earlier) and not finding it. This implies that the file is always shown as not having been referenced on the day on which it is scratched. We have assumed (by setting I(i,L(i))=1) that when a file is scratched, it must be referenced; that is, the user probably copied it to tape or at least made a listing of it. Thus any file which was both created and scratched during the measurement period must be referenced at least twice. This does not necessarily always happen, but our feeling is that setting I(i,L(i))=1 is closer to being accurate than not doing so.

The small number of references to most files was rather surprising to the author, but there appears to be a reasonable explaination (which will also be important later). A user's perception of his file activity tends to be heavily weighted towards those files which are actively used, but in fact most

files seem to be created as temporaries - e.g. to hold a modified version of a program. They are created, left unreferenced until the user needs the file space, and then destroyed. The number of files which contain information in regular use or which follow the intuitively appealing pattern of intensive use during development followed by occasional use later is relatively small.

Measurements of file size distributions are given in figure 5 and in tables V and VI. It is possible to define the distribution of file sizes in three ways. Let SU(i) be the unweighted empirical file size distribution, SWU(i) be the file size distribution as weighted by use and SWL(i) be the file size distribution as weighted by lifetime. Then:

$$SU(i) = \sum_{j=1}^{Nfiles} (1 \text{ if } S(j)=i;\ 0 \text{ otherwise})/Nfiles \qquad (6)$$

$$SWU(i) = \sum_{j=1}^{Nfiles} (Nref(j) \text{ if } S(j)=i;\ 0 \text{ otherwise})/Nref \qquad (7)$$

$$SWL(i) = \sum_{j=1}^{Nfiles} ((L(j)-F(j)+1) \text{ if } S(j)=i;\ 0 \text{ otherwise})/ \sum_{j=1}^{Nfiles} (L(j)-F(j)+1) \qquad (8)$$

The unweighted distribution is simply that obtained by considering the file sizes of all those files in existance during the period of measurement. The weighted by use distribution is that seen by an observer picking an arbitrary file reference and observing the size of the file referenced. The weighted by lifetime distribution is that which would be observed by selecting a random file (on the disks) at a random time. Table V lists the unweighted file size distribution and

Figure 5

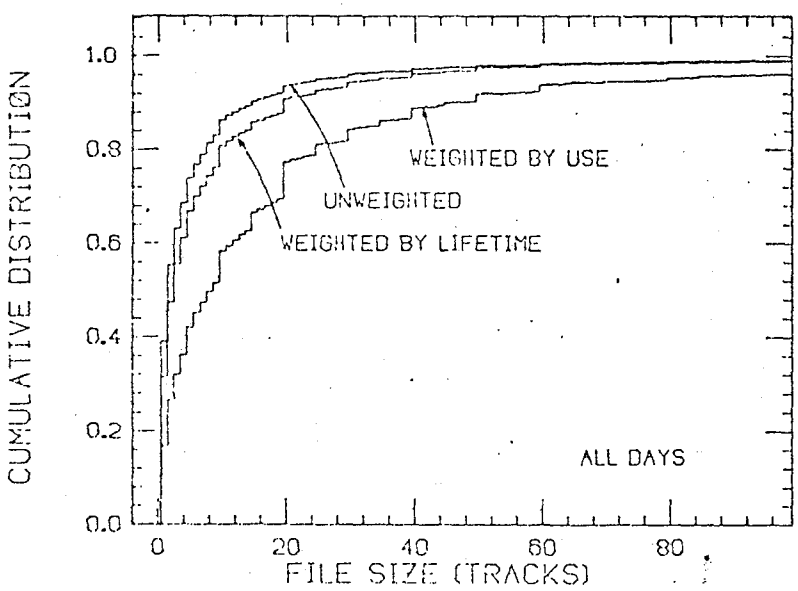

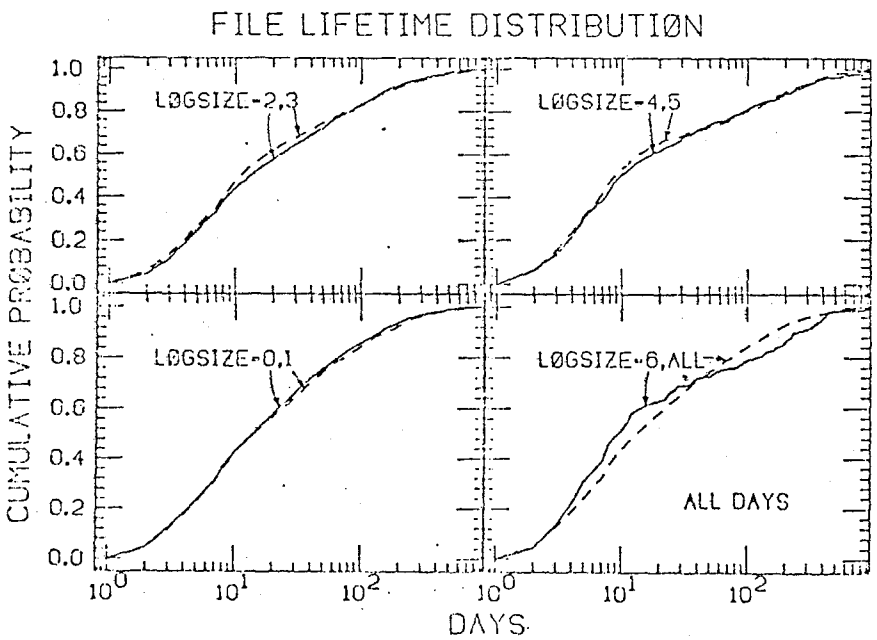

Figure 6

Figure 7

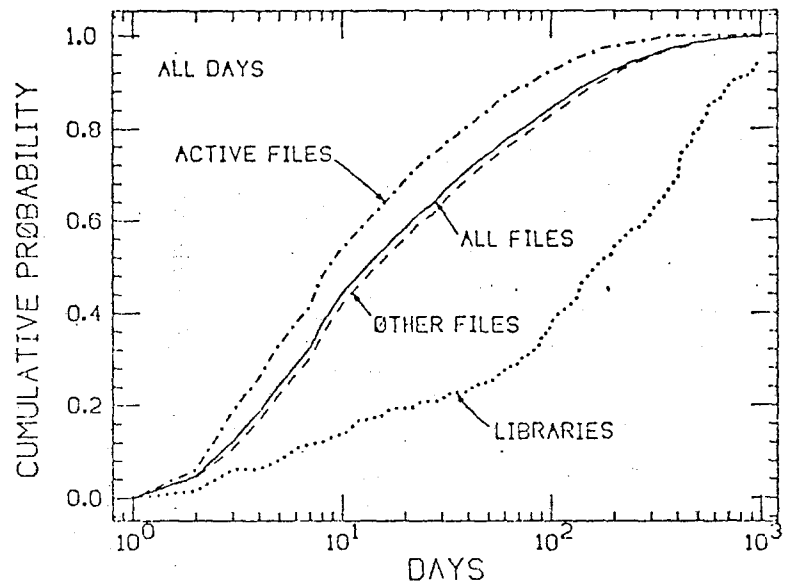

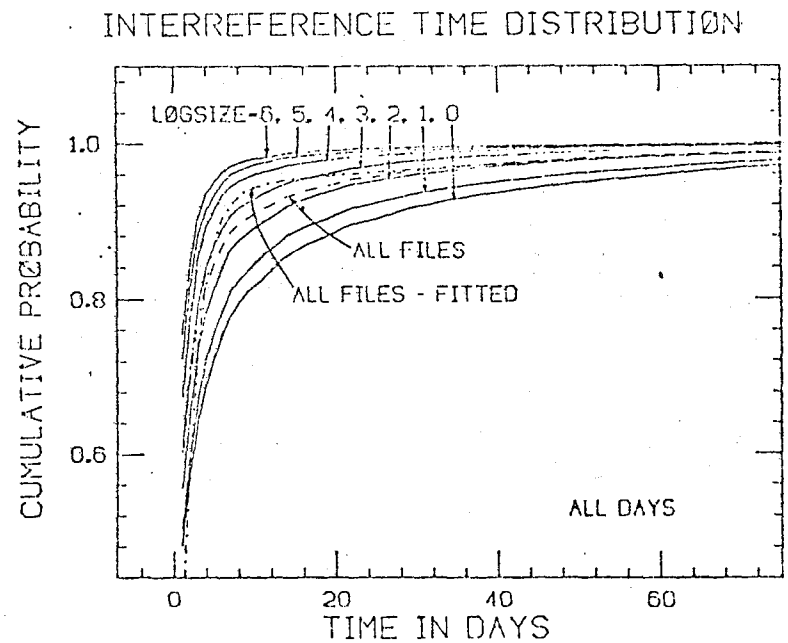

Figure 8

table VI gives the mean and median for all three cases. It is evident that larger files are used more heavily. This observation is intuitively reasonable. Large files are "expensive" to keep since they use a large amount of scarce disk space, and therefore if a file isn't used frequently enough, it is likely to be scratched.

We noted earlier that files were classified both as to size and class. Table VII shows the average file size within each class. Interestingly, the mean file size by the three measurement methods are almost the same for each size class, but very different overall. This suggests (but does not in any sense prove) that within the particular size groupings that we have selected, files are referenced relatively uniformly with respect to size. Additional data arranged by file size and class appears in table VIII, which we discuss below.

E. File Lifetime

An interesting aspect of file behavior is file lifetime; that is, how long the file exists between creation and being scratched. Let LF(i) be the empirical file lifetime distribution. Then:

$$LF(i) = \sum_{j=1}^{Nfiles} (1 \text{ if } D(j)-B(j)+1=i;\ 0 \text{ otherwise})/Nfiles \qquad (9)$$

There is one problem with this definition. D(i) is undefined for files which have not been scratched at the end of the measurement period, and from table I, we see that this is 21% of all files. There appears to be three ways to deal with this difficulty, but none of the three is completely satisfactory. These are: (a) Consider only files which are

Table VIII

All Days

Size (tracks)	Class	Number of Files	Fraction	Number of Refs.	Fraction	Number of Inter-reference Intervals	Mean Inter-reference Interval	Coef. of Variation	Fit Parameters a	b	c	Mean Lifetime	Mean Rate of Reference	Fraction Died
1	Library	12	.00048	256	.00107	244	4.30	3.34	.940	.551	.023	311	.111	.417
2-3		29	.00116	1022	.00428	944	4.65	3.16	.956	.396	.019	186	.167	.310
4-7		120	.00482	4926	.02071	4828	4.13	2.77	.930	.494	.031	319	.175	.242
8-15		249	.01000	17253	.07223	17005	2.99	2.73	.976	.472	.026	330	.277	.265
16-31		232	.00932	24591	.10295	24360	2.26	2.36	.982	.562	.034	347	.382	.272
32-63		67	.00269	9424	.03945	9357	1.77	1.61	.964	.707	.089	86	.534	.313
>=64		14	.00056	2129	.00891	2115	1.86	2.37	.991	.636	.031	422	.477	.286
1	Other	6952	.27922	33249	.13919	26421	9.58	2.56	.851	.349	.021	61	.068	.787
2-3		5180	.20805	33185	.13892	28092	8.51	2.67	.888	.317	.020	63	.078	.755
4-7		3394	.13632	31076	.13010	27733	6.36	2.88	.925	.351	.020	64	.113	.776
8-15		2421	.09724	27894	.11677	25508	5.04	2.91	.946	.373	.022	67	.143	.788
16-31		1087	.04366	18546	.07764	17471	4.05	3.00	.959	.425	.023	73	.187	.751
32-63		456	.01831	12892	.05397	12442	3.05	3.33	.982	.473	.018	81	.261	.717
>=64		385	.01546	12191	.05104	11807	2.14	2.68	.984	.608	.030	74	.374	.805
1	Active	2777	.11154	6629	.02775	3884	14.01	2.24	.831	.202	.017	34	.064	.924
2-3		774	.03109	1838	.00769	1073	12.24	2.19	.808	.256	.021	28	.073	.943
4-7		406	.01631	958	.00401	555	9.98	2.68	.932	.193	.013	23	.087	.931
8-15		183	.00735	412	.00172	231	6.97	3.48	.960	.279	.011	16	.113	.945
16-31		116	.00466	271	.00113	155	5.06	2.93	.890	.639	.030	13	.200	1.00
32-63		43	.00173	105	.00044	62	3.11	1.46	–	–	–	8	.274	.930
>=64		1	.00004	4	.00002	3	11.67	1.23	–	–	–	37	.111	1.00
1	All	9741	.39124	40134	.16802	30549	10.10	2.52	.852	.317	.020	52	.067	.825
2-3	Files.	5983	.24030	36045	.15090	30159	8.52	2.66	.887	.316	.020	58	.079	.777
4-7		3920	.15744	36980	.15481	33116	6.09	2.90	.930	.353	.020	61	.117	.776
8-15		2853	.11459	45559	.19073	42744	4.23	2.97	.959	.403	.022	71	.175	.752
16-31		1435	.05764	43408	.18172	41986	3.02	2.96	.975	.490	.025	84	.263	.693
32-63		566	.02273	22421	.09386	21861	2.50	3.17	.988	.523	.019	74	.332	.686
>=64		400	.01607	14324	.05997	13925	2.10	2.65	.985	.612	.031	79	.387	.788
All	Library	723	.02904	59621	.24959	58903	2.58	2.70	.979	.529	.029	303	.321	.272
	Other	19875	.79826	169033	.70763	149474	6.23	2.94	.922	.376	.021	64	.108	.774
	Active	4300	.17270	10219	.04277	5963	12.70	2.33	.853	.209	.017	30	.070	.931
All	All	24898	1.0	238871	1.0	214340	5.41	3.08	.938	.399	.020	59	.126	.786

actually scratched during the period of observation. The difficulty is that this will give an incorrect estimate unless the system is completely in steady state. This is because if files are being created more quickly than they are scratched, the file lifetime distribution will be overweighted towards short lived files; the long lived files will not have had a chance to die yet. (b) Assume all files are destroyed on the last day of observation. This doesn't appear to make any sense statistically. (c) Assume that since the last day of observation is a random point in time, it happens to exactly bisect the lifetime of all files in existance on that day. This argument is closely related to that used to derive the distribution of time to the next event in a renewal process when the process is observed at a random point in time (Cox, 1962). The expected time to the next event (i.e. file death) is equal to the time since the last event (file birth). This approach is satisfactory if all files have finite lifetimes and if there is a large number of such files, so that a distribution may be collected. We have selected (a) as our means of estimating the file lifetime distribution; there are arguments to be made for (c) as well.

The measured file lifetimes are presented in figure 6 (by size) and figure 7 (by class). The file lifetime appears to be only slightly related to size, but is strongly influenced by class. Active files get scratched relatively quickly (see figure 7) and libraries tend to never be scratched. The mean lifetime is shown in table VIII for each size/class combination. Also shown in table VIII is the fraction of files of each size/class that died (were ever scratched). Among the libraries, a full 73% were still in existance at the

end of the measurement interval. For the reasons noted above, our measurements of file lifetimes may be biased, and considering the large number of files not scratched, that bias may be quite significant. The reader therefore should not attach too much significance to our file lifetime figures.

Further study of file lifetimes seems to have only slight payoff in terms of file migration, which is the ultimate aim of this study. Therefore, no attempt has yet been made to model file lifetimes. An effort in this direction is being considered for future research.

F. File Reference Patterns

1. Interreference Times

A very important aspect of file reference behavior from the point of view of this paper is the actual sequence of referenced - nonreferenced days, and in particular, the distribution of times between references. Let g(i,logsize,class) be the empirical probability mass function for the times between references to files belonging to size class "logsize" and type class "class". Let "*" denote the sum over all possible entries for that argument and let "-" denote an unspecified entry for that argument. Then

$$g(i,\text{logsize},\text{class}) = \sum_{j\in(\text{logsize},\text{class})} \sum_{k=1}^{\text{Nrefs}(j)-1} \begin{pmatrix} 1 \text{ if } C(j,k) = i; \\ 0 \text{ otherwise} \end{pmatrix} \Big/ \sum_{j\in(\text{logsize},\text{class})}^{\text{Nfiles}} \left(\text{Nrefs}(j) - 1 \text{ if } \text{Nrefs}(j) \geq 1\right) \quad (10)$$

$$g(i,*,\text{class}) = \sum_{j\in\text{class}} \sum_{k=1}^{\text{Nrefs}(j)-1} (1 \text{ if } C(j,k) = i;\ 0 \text{ otherwise}) \Big/ \sum_{j\in\text{class}} (\text{Nrefs}(j)-1 \text{ if } \text{Nrefs}(j) \geq 1) \quad (11)$$

$$g(i,\text{logsize},- \sum_{j\in\text{logsize}} \sum_{k=1}^{\text{Nrefs}(j)-1} \begin{pmatrix} 1 \text{ if } C(j,k) = i; \\ 0 \text{ otherwise} \end{pmatrix} \Big/ \sum_{j\in\text{logsize}} (\text{Nrefs}(j)-1 \text{ if } \text{Nrefs}(j) \geq 1) \quad (12)$$

$$g(i,*,*) = \sum_{j=1}^{Nfiles} \sum_{k=1}^{Nrefs(j)-1} (1 \text{ if } C(j,k) = i;\ 0 \text{ otherwise}) \Big/ \sum_{j=1}^{Nfiles} (Nrefs(j)-1 \text{ if } Nrefs(j) \geq 1) \tag{13}$$

We let G(i,logsize,class) be the cumulative distribution. In figure 8 we show as solid lines the distributions G(i,logsize,*) and in figure 9 the distributions G(i,*,class). (The lines showing fitted results are discussed below.) The number of interreference intervals and the mean interreference times are given in table VIII for each size and class combination.

Another column in table VIII shows the coefficient of variation of the interreference time distribution, as aggregated over all of the files in a size/class combination. Let CV(logsize,class) be the coefficient of variation for a size/class case. Then

$$CV(-,-) = \frac{\left(\sum_i g(i,-,-)\left(i - \sum_i ig(i,-,-)\right)^2\right)^{1/2}}{\sum_i ig(i,-,-)} \tag{14}$$

We observe that these coefficients of variation are all over 1.0 and the mean (i.e. for g(i,*,*)) is 3.08. The distribution is thus moderately skewed, but this implies nothing about the reference pattern to individual files. Such a skewed distribution could equally well be obtained from either of two circumstances (among many others): (a) each file is referenced with the same skewed interarrival time distribution peculiar to its size and class or (b) each file is referenced as a Bernoulli process but with the rate of reference varying between different files. (The Bernoulli process is the discrete time analogue of the Poisson process, by which a file i has a constant probability p(i) of being

Figure 9

INTERREFERENCE TIME DISTRIBUTIØN

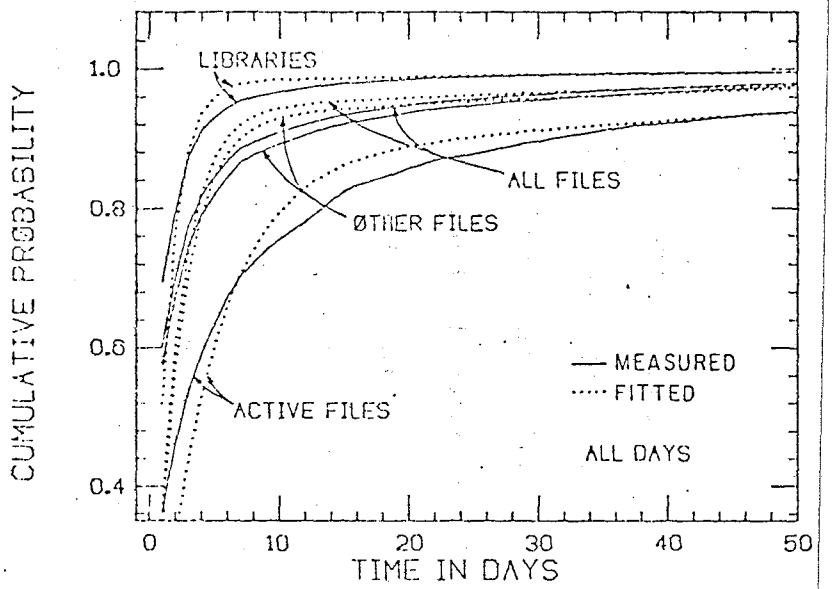

HAZARD RATE

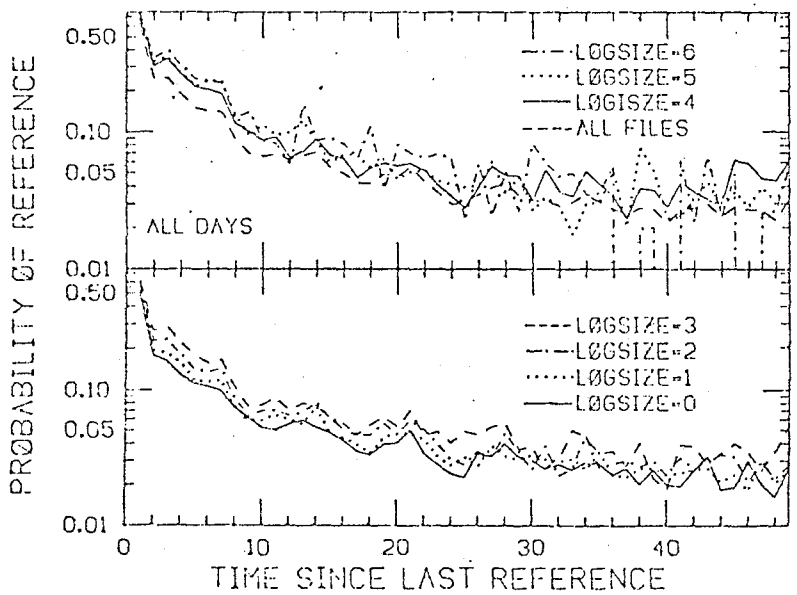

Figure 10

Figure 11

HAZARD RATE

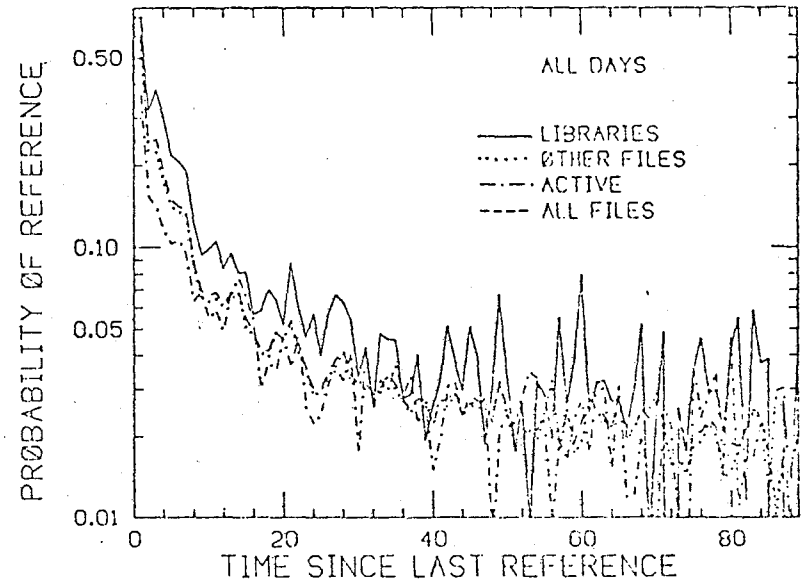

RATE ØF REFERENCE TØ FILES

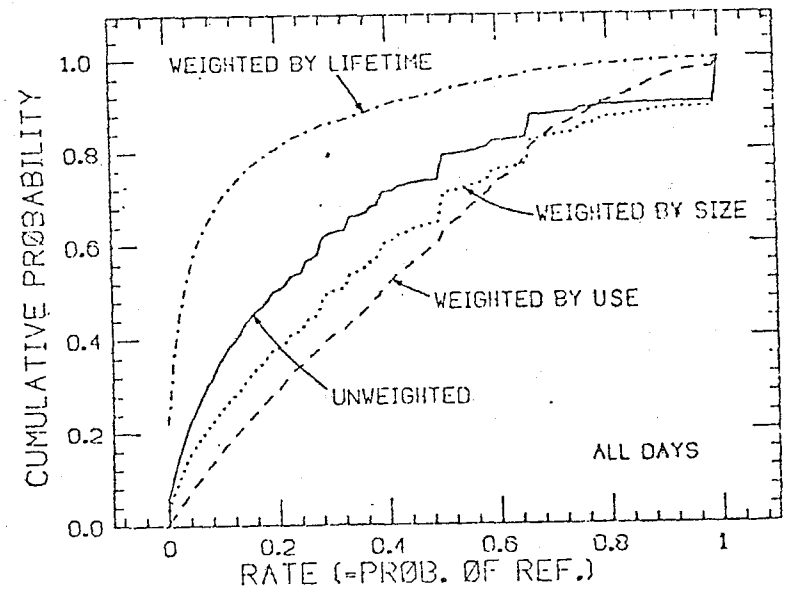

Figure 13

referenced each day. The interevent times are geometric.)

Another statistic which has been measured is the hazard rate of the empirical interreference time distribution. The hazard rate is defined as

$$h(i,-,-)=g(i,-,-)/(1-G(i-1,-,-)) \qquad (15)$$

Examination of the hazard rate function (see figures 10, 11) indicates that it declines sharply for a while and then becomes (after 20 days or so) relatively flat.

A distribution that is frequently used to model discrete empirical distributions of the form we have described (moderately skewed, declining but eventually flat hazard rate) is the hypergeometric distribution, which is the weighted sum of geometric distributions. We have selected a two part hypergeometric to fit g(i,-,-). Let gf(i,-,-) and Gf(i,-,-) be the appropriate fitted distributions. Then:

$$gf(i,-,-)= a\ b\ (1-b)**(i-1) + (1-a)\ c\ (1-c)**(i-1) \qquad (16)$$

for some a,b, and c to be selected in each case.

There are three methods in common use used to fit the parameters of a fitted distribution to empirical ones: least squares, maximum likelihood and method of moments. The method of moments is by far the simplest to use but is vulnerable to instability in the case that the number of points in the sample distribution is small. That does not appear to be a problem with our data, since there are a large number of interreference intervals in almost every size/class combination. We have therefore used the method of moments to estimate the parameters a, b and c in each case. The method

of moments involves calculating the moments (in this case, the first three) of the empirical distribution and selecting the set of values of the distribution parameters (a,b,c) that yield the same moments. This was done and the results appear in table VIII. We note that in two cases, no set of parameters a, b and c existed that yielded the observed moments.

In figure 8, the overall fitted distribution Gf(i,*,*) is shown as a dot-dash line and in figure 9, the fitted distributions Gf(i,*,-) are shown as dotted lines. As may be seen, the quality of the fit is at best fair. This observation is confirmed by the chi-square test for goodness of fit, which rejected the fit in almost all cases at the 99% confidence level. No attempt has yet been made to fit the observed distribution with other postulated distributions; we are considering doing so at a later time. Use of this fitted distribution is made in the companion paper to this, however, and despite the fair to poor quality of the fit, surprisingly good results are obtained for its use with file migration algorithms.

2. Testing File Reference Patterns

In the last section, we obtained the distribution of times between references to a file aggregated over all files in a size/class combination. As we noted, this says very little about the reference pattern to individual files. In particular, there are three questions that should be answered with regard to individual file reference patterns: (a) Is the rate of reference to a file stationary? (b) Are successive interreference intervals correlated? (c) Is the reference pattern to a file the Bernoulli process with a parameter

specific to that file? Testing a time series to obtain the answers to a, b and c would be straightforward were it not for the fact that most of the sequences of interreference intervals for the files are very short. As noted earlier, 50% of the files were referenced 2 or fewer times. It makes absolutely no sense to test a time series of 2 elements (or even 3 or 4). Further, results don't appear to be available in the published literature for the distributions to be expected when applying standard statistical tests to such small samples. Our basic methodology, which we summarize below, consists essentially of restricting our tests to those file reference patterns which can reasonably be tested, and then testing them against distributions obtained through pseudo-random number driven simulation. The details of our methodology and results will appear in a paper in preparation (Smith, 1978c).

The only files tested were those that (a) were referenced at least six times, (b) had a value of (Nref(i)-1)/(L(i)-F(i)) less than .95 and (c) for which L(i)-F(i) was greater than or equal to 10. The idea was to eliminate those files that either were referenced too few times, or were referenced over such a short period of time that the number of possible reference patterns was too small to allow for statistical testing. For example, a file referenced on six days in a seven day lifetime can only show five different possible reference patterns. We found that this elimination procedure left us with 5,334 (21.4% of all) files for all days and 4,818 (19.4% of all) files in the working days case. Approximately 80% of all of the file references remained, however, since those files eliminated were those which were referenced very

few times.

We tested for monotonic trend, serial correlation and skewness of the distribution. The null hypothesis in each case was that of a stationary, uncorrelated Bernoulli process. The statistics used are explained below. For short discrete time time series, there appear to be very few results known, and so the means and confidence intervals for the statistical tests were determined by generating a large number of random Bernoulli processes (of appropriate rates and lengths) and applying the statistical test to those time series. That is, a large number of Bernoulli time series were generated using a pseudo-random number generator. For each of these time series, each of the statistics given below in equations 17, 18 and 19 were calculated. This yielded the distribution of the statistic in the case of the null hypothesis - i.e. the Bernoulli process.

We do note one problem with our statistical procedure. Because of the short length of most of the time series we deal with, the values of the statistics (eq. 17-19) used will vary widely even in the case that the null hypothesis is valid. Thus it will be difficult to determine that our sample time series are not Bernoulli, and the test procedure may not actually be very powerful. An alternative approach is to find some (clever) way of aggregating the data from several time series, but we leave this for future research.

(a) Testing for Trend

Lewis and Shedler (1973) (see also Cox and Lewis, 1966) give a statistic for testing a continuous time time series for being a stationary uncorrelated Poisson process against the alternative of a Poisson process with monotonic trend. As

adapted for discrete time series, it is

$$Tr = \frac{z - T/2}{T/\sqrt{12Nref(i)}} \qquad \text{where } T = L(i) - F(i) + 1,$$

$$Z = \sum_{j=F(i)}^{L(i)} I(i,j)\,(j - F(i))/Nref(i) \qquad (17)$$

This statistic is simple to calculate and seems to be equally suitable for testing discrete time series.

As shown in table IX, approximately 35% of the file reference patterns displayed significant trend (at the 99% confidence level); the vast majority of such cases were of declining trend. This is as one might expect - a file is used less after it has been around for a while. We comment that a skewed but stationary distribution will display a larger variance in Tr than the Bernoulli process; thus the 99% confidence level was used rather than the more usual 95%, in order to minimize false rejection.

(b) Testing for Serial Correlation

The usual estimator (and usually the most powerful; see Lewis and Shedler, 1973) for serial correlation is given in equation (3); that estimator, however, is an approximation that is acceptable for long time series. The correct estimator is (from Kendall and Stuart, 1976):

$$S(k) = \frac{\frac{1}{n-k}\sum_{i=1}^{n-k}\left(x_i - \frac{1}{n-k}\sum_{i=1}^{n-k} x_i\right)\left(x_{i+k} - \frac{1}{n-k}\sum_{i=1}^{n-k} x_{i+k}\right)}{\left[\frac{1}{n-k}\sum_{i=1}^{n-k}\{x_i - \frac{1}{n-k}\sum_{i=1}^{n-k} x_i\}^2 \frac{1}{n-k}\sum_{i=1}^{n-k}\{x_{i+k} - \frac{1}{n-k}\sum_{i=1}^{n-k} x_{i+k}\}^2\right]^{1/2}} \qquad (18)$$

We warn the reader that this estimator is biased to the approximate extent of 1/n, where n is the number of elements

Table IX

Trend Test

Fraction That Reject (99% Confidence Level)
For Increasing/Decreasing Trend

	All Days	Working Days
Unweighted	.052/.282	.051/.274
Weighted by Lifetime	.087/.323	.077/.304

Table X

Serial Correlation Test

Fraction That Reject (99% Confidence Interval)
For Excessive Correlation

All Days	.042
Working Days	.066

Table XI

Distribution of Number of Times a File
is Used in 15 Working Days

Number	Probability	Cumulative
0	.600	.600
1	.112	.713
2	.060	.773
3	.037	.809
4	.025	.834
5	.020	.854
6	.017	.871
7	.015	.886
8	.013	.899
9	.012	.911
10	.011	.922
11	.012	.934
12	.012	.946
13	.013	.959
14	.014	.973
15	.027	1.00

Table XII

Coeficient of Variation Test

Fraction That Reject (99% Confidence Interval)
The Bernoulli Process

	All Days	Working Days
Unweighted	.354	.363
Weighted by Use	.494	.470

Table XIII

Mean/Median Rate of Reference

Weighted By	Mean	Median
Unweighted	.306	.19
Use	.410	.386
Lifetime	.121	.030
Size	.380	.296

in the time series. Thus knowlege of the actual mean and distribution of the estimator is essential.

The serial correlation was estimated using equation (18) for all files showing no significant trend. The fraction of files which displayed significant serial correlation is shown in table X, where we see that only about 5% of all files rejected the hypothesis of no correlation at the 99% confidence level. (Our earlier comment about skewed distributions holds here also.) In almost all cases, those significant correlations found were positive ones, not negative ones. Most correlation coefficient values observed were small (65% were less than .25) so that independent of their significance, their predictive power is low.

An additional test was made to look for possible serial correlations in interreference intervals. Figure 12 shows an x-y plot of successive interreference intervals taken over all interreference intervals for all files. Each entry is the (truncated) logarithm base 2 (in hexadecimal) of the number of points at that location. There is no pattern apparent to the author in this figure.

The cross correlation (between x and y) was computed over all of the points in this figure and a value of .295 (.307 for working days) was obtained. This a high and significant value, but it presumably reflects the fact that we have lumped all files together, rather than actual correlations within individual file reference processes. That is, files referenced at a given rate are likely to continue being referenced at a similar rate; this rate may be characteristic of the file or its size or its class.

We also collected data to look at the extent of

CONSECUTIVE INTERRESPONSE INTERVALS; 2*ND(Y) VS FIRST(X); ENTRY IS LOG2(HITS)

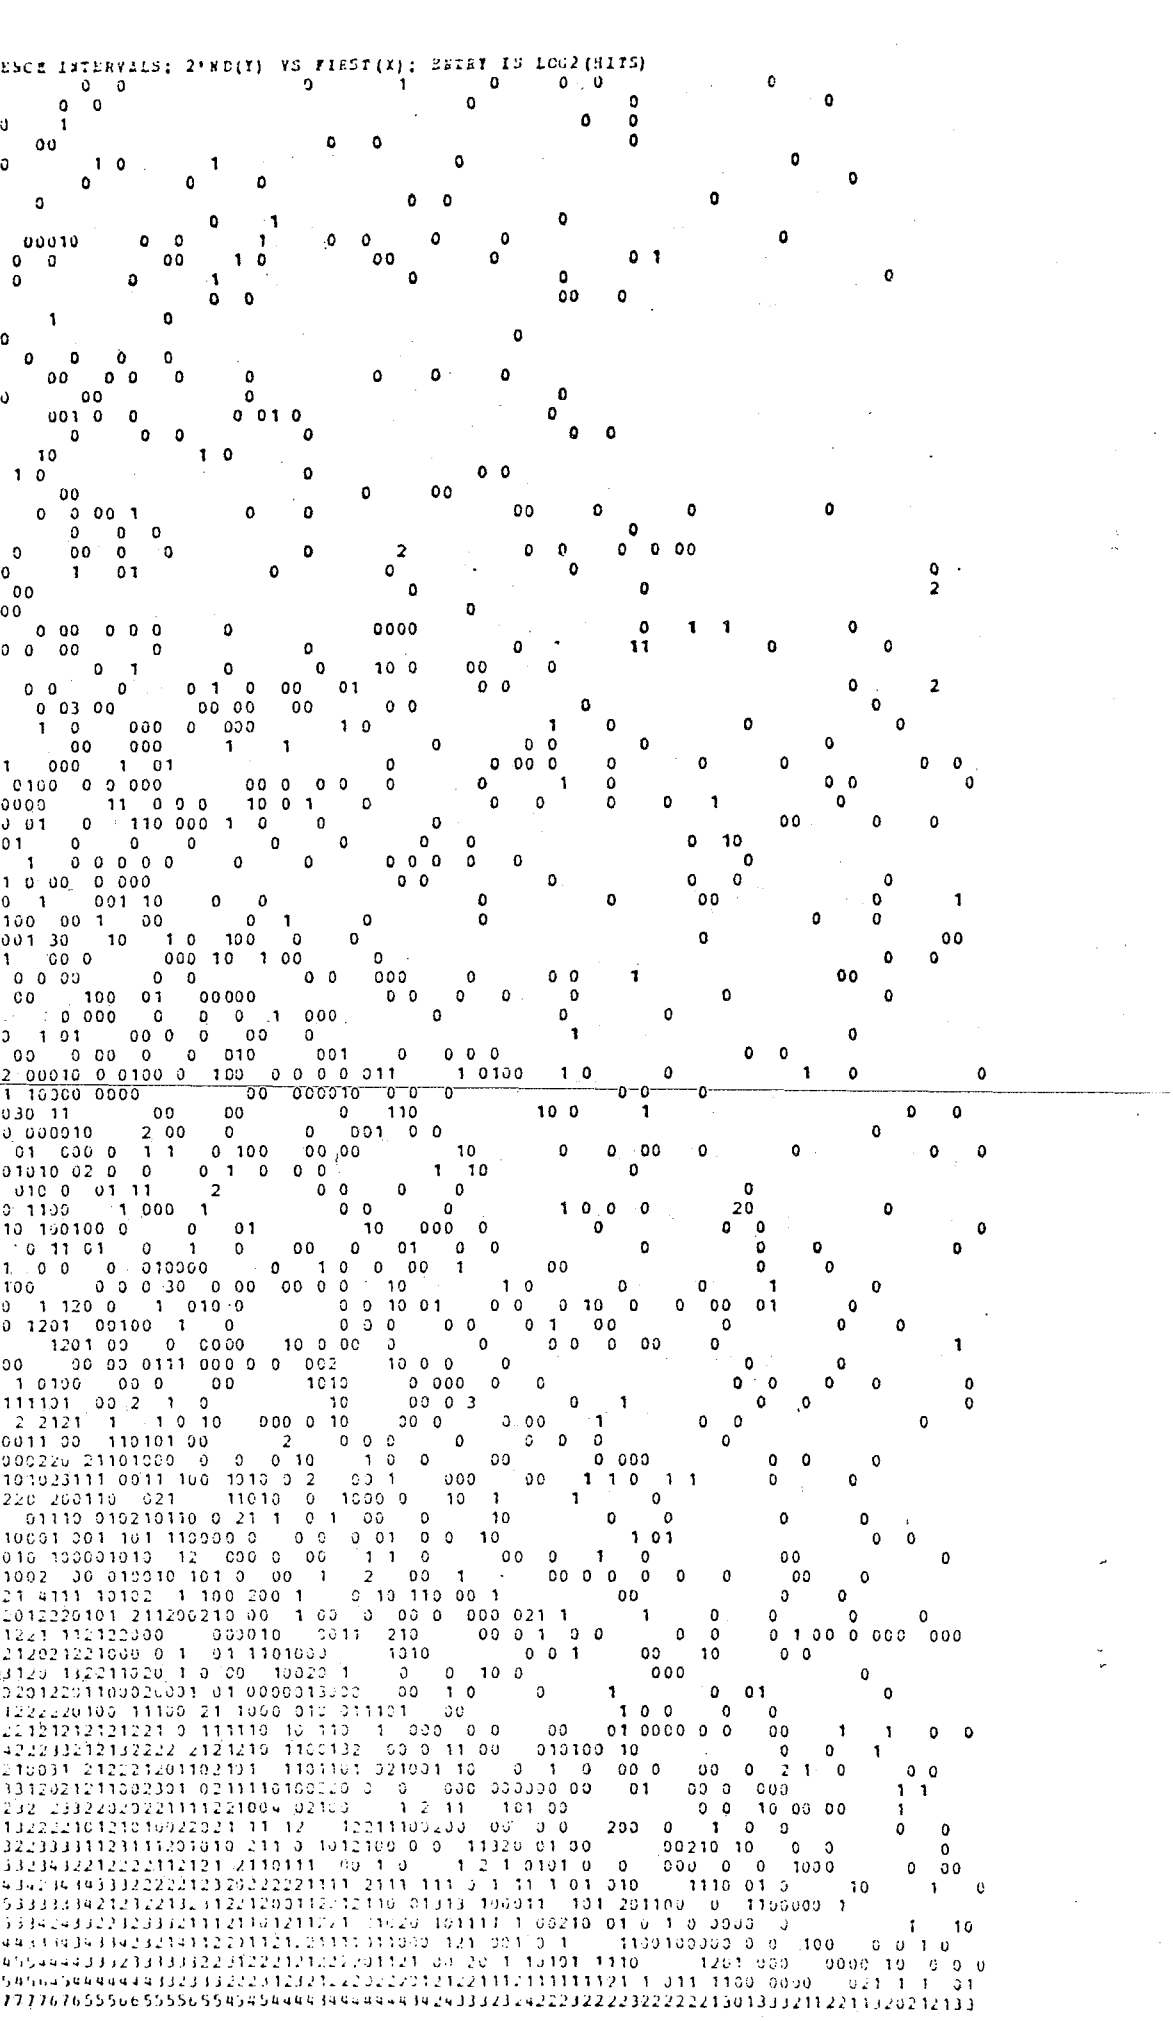

FIGURE 12

clustering in the reference pattern to a file. For every period of 15 days during which a file existed, the number of references to that file was counted. The distribution of the number of references to a file in 15 working days is given in table XI. As can be seen, there is substantial clustering; if a file is used a few times, it is just as likely to have been used on many more days. This confirms one's intuitive idea of file use - if a file is in active use (development or production), it is likely to be referenced with high probability each day over a several day period.

(c) Testing for the Bernoulli Process

Testing an interarrival process against a known statistical distribution is usually done with the chi-square goodness of fit test (see eq. 1) using the empirical interarrival distribution. This is only feasible, however, if the number of events in the arrival process is quite large (at least 25 or 50 in this case). We choose instead to measure the coefficient of variation of the interarrival intervals for each file and then determine whether that coefficient of variation is an acceptable one under the hypothesis of a Bernoulli process.

The estimator for the coefficient of variation used is:

$$\frac{\left[\frac{1}{Nref(i)-2}\sum_{j=1}^{Nref(i)-1}\left(C(i,j)-\left(\sum_{j=1}^{Nref(i)-1}C(i,j)/\left(Nref(i)-1\right)\right)\right)^2\right]^{1/2}}{\sum_{j=1}^{Nref(i)-1}C(i,j)/\left(Nref(i)-1\right)} \qquad (19)$$

We again warn the reader that this estimator does not approach its asymptotic distribution until the time series becomes quite long. We found (see table XII) that about 40% of all files tested (those with no trend) showed unacceptable

coefficients of variation at the 99% confidence level. In almost every case, this was because the coefficient of variation was too large. Thus our earlier caution in the trend and correlation tests regarding skewed distributions appears to have been appropriate.

It also happens that 96% of all coefficients of variation are below 2.0, whereas the mean coefficient of variation when computed over all interreference intervals for all files is 3.08. Similar high coefficients of variation are observed for individual size/class combinations (see table VIII). Thus the interreference distribution for a size/class combination may not be a very good estimate of the distribution for an individual file.

(d) Conclusions

From our tests in this section (III.F.2), we can see that files are frequently used less as they age, and that the interreference time distribution is more skewed than the Bernoulli. There is little if any significant serial correlation within the reference process to one file. A file migration algorithm would therefore find it useful to condition on the time since the last reference and perhaps the age of the file; conditioning on previous interreference intervals doesn't appear to be helpful.

It is worth mentioning that different results were reported by Stritter (1977) who analyzed the same data. His descriptions of the tests he ran are vague but he found that the number of files displaying either trend or large coefficients of variation to be small enough that he was willing to describe the file reference process as "Poisson". It appears that he tested only a subset of the file reference

processes (selection process unspecified); thus from his description of his results (which are not quantified) it isn't clear if there is actually a conflict with our results or not. We believe that our methodology has been more thorough and that the results presented here are correct.

3. Rates of Reference

Another interesting statistic is the measured rates of reference to files. The rate of reference to a file i, R(i), is just Nref(i)/(L(i)-F(i)+1). That is, it is just the number of times the file is referenced over the period of observation divided by the lifetime of the file during the period of observation.

We have computed the cumulative distribution of R(i) in four ways: RU(x) is the unweighted distribution:

$$RU(x) = \sum_{i=1}^{Nfiles} (1 \text{ if } R(i) <= x)/Nfiles \qquad (20)$$

This distribution can be weighted by file lifetime as

$$RL(x) = \sum_{i=1}^{Nfiles} (L(i)-F(i)+1 \text{ if } R(i) <= x;\ 0 \text{ otherwise})/ \sum_{i=1}^{Nfiles} (L(i)-F(i)+1) \qquad (21)$$

or by use as

$$RUS(x) = \sum_{i=1}^{Nfiles} (Nref(i) \text{ if } R(i) <= x;\ 0 \text{ otherwise})/Nref \qquad (22)$$

or by file size as

$$RS(x) = \sum_{i=1}^{Nfiles} (S(i) \text{ if } R(i) <= x;\ 0 \text{ otherwise})/ \sum_{i=1}^{Nfiles} S(i) \qquad (23)$$

The measured values for RU(x), RL(x), RUS(x), and RS(x) all appear in figure 13. The mean and median rates of reference are given in table XIII.

It is also interesting to compute the rate of reference to a file as a function of its age. The age of a file is defined to be the number of days since it was created. We compute the mean rate of reference RA(i,-,-) to a file as a function of its age i by considering those files which were created during the period of observation.

$$RA(i,-,-) = \sum_{j\in(-,-)} A(j,i) \Big/ \sum_{j\in(-,-)} (1 \text{ if } (L(j)-F(j)+1)>=i; \quad 0 \text{ otherwise}) \qquad (24)$$

Note that RA(1,-,-)=1. The function RA(i,-,-) is rather irregular for moderately large values of i, so we have smoothed it by using a five day rectangular window. The smoothed function RA(i,-,-) is shown in figures 14 and 15. The up-down regular oscillation of RA is due to the interaction of the smoothing window with weekends. Since most files are created during working days, a file of age 19 is less likely to be referenced than a file of age 21. (Since a file of age 21 is very likely to be in existance on a working day.) A smoothing window of a multiple of 7 days would have removed this oscillation, but was judged to be too wide.

Also shown in figures 14 and 15 is the value of RAf(i,-,-) obtained from our fitted two part hypergeometric distribution of equation (16). We assume that the file is referenced as a renewal process, with the gf(i,-,-) interarrival distribution, and that the file was referenced on

Figure 14

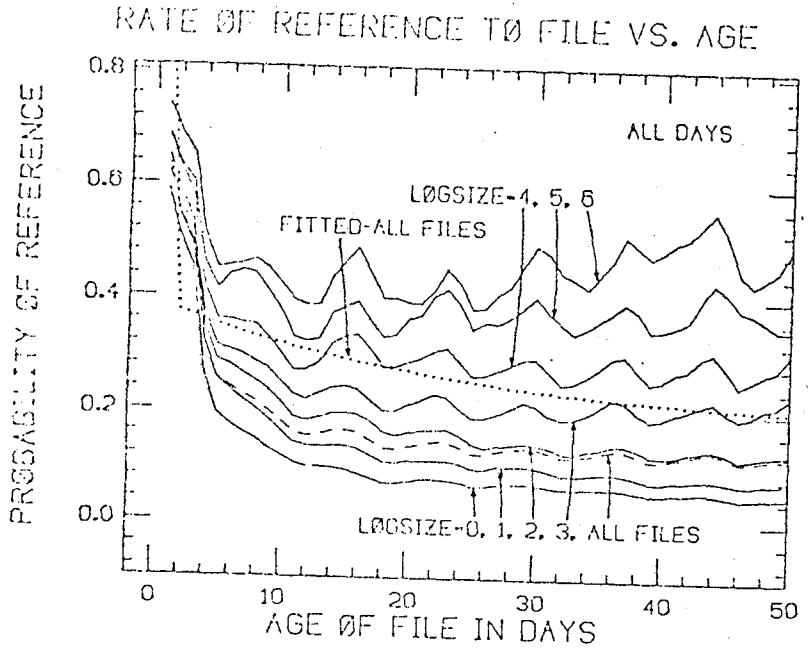

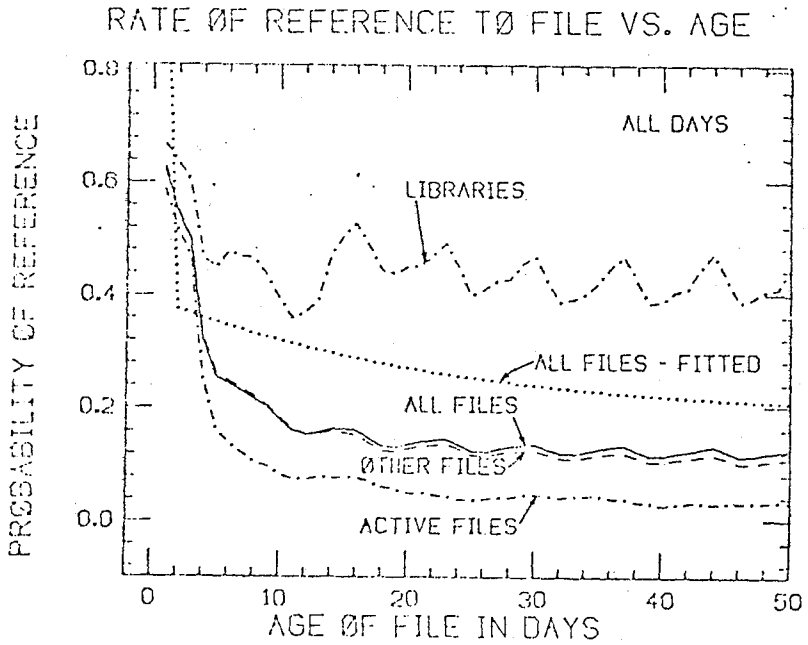

Figure 15

Figure 16

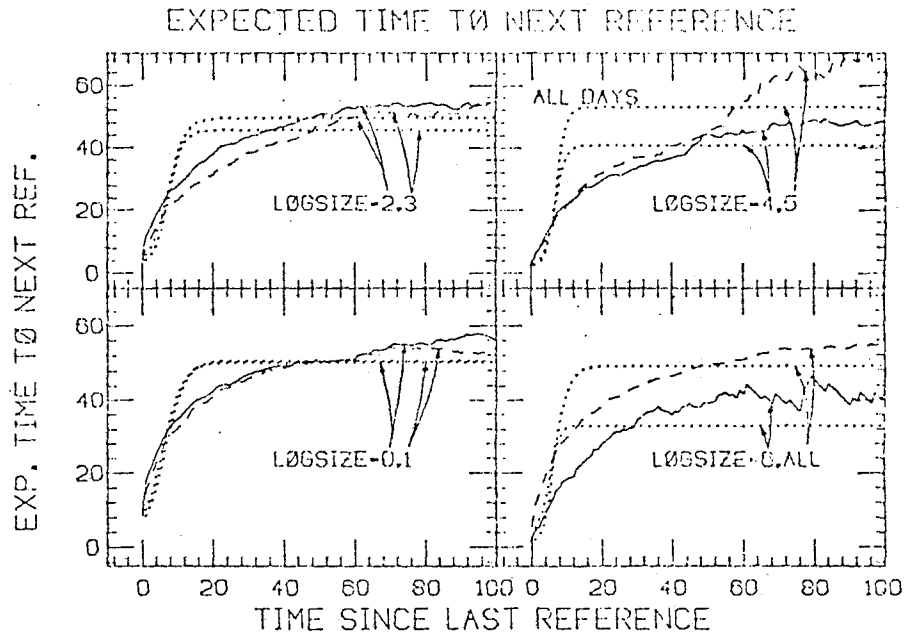

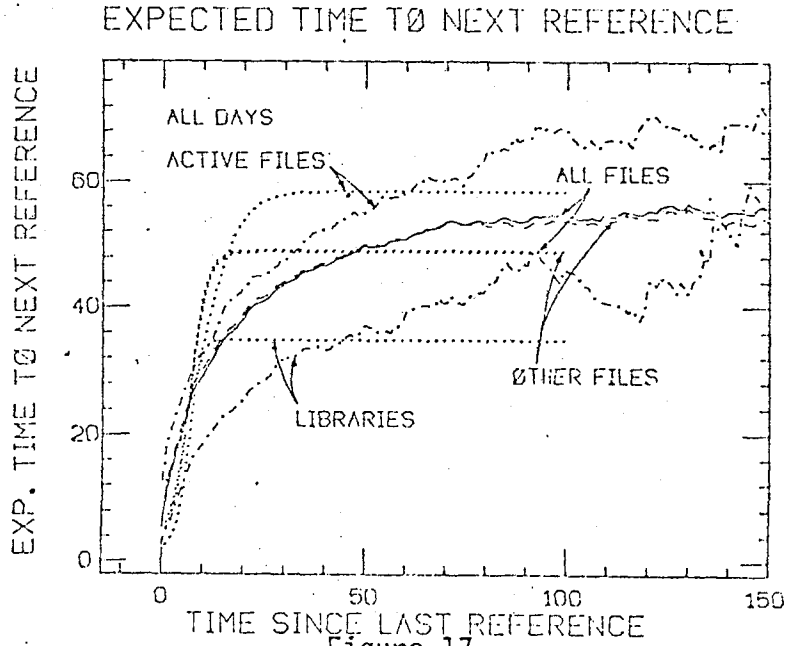

Figure 17

the day of its birth. Then from the renewal equation (Cox, 1962), one has:

$$RAf(i,-,-) = \sum_{k=1}^{i-1} RAf(k,-,-)\,g(i-k,-,-),\ i>1,$$
$$\text{where } RAf(1,-,-)=1 \qquad (25)$$

We note that because a file is referenced on the day it is created, the predicted rate of reference declines with age. The fact that the process is stationary (e.g. a renewal process) does not imply that the rate of reference need be constant as a function of age.

This declining rate of reference to a file with age is consistent with our earlier determination that about 35% of all tested files showed a declining rate of reference during their lifetime.

4. Expected Time to Next Reference

In performing file migration, one would generally like to remove that file with the largest space-time product to the next reference (size of file x time to next reference). (Criteria for file migration are discussed in detail in the companion paper to this.) One can estimate the time to the next reference by conditioning on the time since the last reference, among other things. Let E(i,-,-) be the expected time to next reference, given that the file hasn't been referenced for i days. If there have been i days elapsed since the last reference (i=0 means that the reference was that day), then by conditioning only on the time since last reference,

$$E(i,-,-) = \sum_{i-k+1}^{\infty} g(k,-,-)\,(k-i)/(1-G(i,-,-)) \qquad (26)$$

In figures 16 and 17 we show the expected time to next reference, both as calculated from the empirical distributions g(i,-,-) and from the fitted distributions gf(i,-,-). The dotted lines in each case are the fitted distributions.

The important observation to be made from figures 16 and 17 is that the expected time to next reference is a generally increasing function of the time since the last reference. Thus one is more likely to want to migrate a file which has not been used for a long period of time than a file which has been recently referenced. This issue is considered further in Smith (1978c).

IV. SUMMARY, CONCLUSIONS AND PLANS FOR FURTHER RESEARCH

We have observed that most files are referenced very few times. Of those files susceptable of statistical testing, we found that about one-third showed a declining rate of reference. Amost half of the files showing no trend displayed reference patterns which ruled out a Bernoulli process model for the reference process; almost no serial correlation was detected between successive interreference intervals. Thus for those files displaying no trend, a renewal process model with a moderately skewed interreference time distribution would appear to be appropriate. Agreeing with the observed frequent decline in the rate of reference to a file with its age is the overall decrease (when aggregated over all files) in the rate of reference to a file with age. We also note that we were not able to fit the interreference time distribution with a two part hypergeometric. Our interpretation of the results enumerated in this paragraph is

that a Markov chain model for the file reference process, if one exists, is likely to require a minimum of three or four states (including a state which represents file death). A very simple renewal process model or even a two state Semi-Markov Process model (e.g. Lewis and Shedler, 1973) seems insufficient to represent the observed properties of the data. Further research to develop a satisfactory model for the file reference process is planned.

The goal of our data analysis in this paper has been twofold: (a) describe and characterize to whatever extent reasonable the file behavior patterns observed and (b) to develop a basis for the specification of file migration algorithms. For a file migration algorithm, it appears to be useful to use: the time since last reference, the file size, the file class and the file age in order to predict the time to next reference for a file. The lack of serial correlation suggests that conditioning on previous file interreference intervals will not be useful. In the companion paper to this (Smith, 1978c), we do use the data analysis in this paper as a basis for the construction and evaluation of a number of file migration algorithms.

For our description and analysis of file reference data in this paper to be useful in the sense of being applicable to other installations, there must be some reason to think that file behavior will be similar accross users and computer systems. Until data is gathered and analyzed for other systems, it is of course impossible to say, but we believe the following: users use text editor files in much the same manner everywhere, subject to the distortions induced by the accounting or file migration algorithms. User access patterns

to non text editor files, such as data files or data bases, are still unknown and it is not reasonable to claim that our results here will apply.

ACKNOWLEDGEMENTS

Many thanks to Dr. Edward Stritter for collecting the data used in this paper and for making it available to the author for this research. This work was supported by the U. S. Department of Energy under Contract W-7405-ENG-48.

BIBLIOGRAPHY

Peter J. Bickel and Kjell A. Doksum, Mathematical Statistics: Basic Ideas and Selected Topics, Holden-Day, San Francisco, 1977.

George E. P. Box and Gwilym M. Jenkins, Time Series Analysis: Forecasting and Control, revised edition, Holden-Day, San Francisco, 1976.

C. J. Conti, "Concepts for Buffer Storage", IEEE Computer Group News, 2, August, 1969, pp. 9-13.

D. R. Cox, Renewal Theory, Methuen, London, 1962.

D. R. Cox and P. A. W. Lewis, The Statistical Analysis of Series of Events, Methuen, London, 1966.

Peter J. Denning, "Virtual Memory", Computer Surveys, 2, 3, September, 1970, pp. 153-189.

Peter J. Denning, "On Modeling Program Behavior", Proc. SJCC, 1972, pp. 937-944.

Roger Fajman and John Borgelt, "WYLBUR: An Interactive Text Editing and Remote Data Entry System", CACM, 16, 5, May, 1973, pp. 314-322.

K. R. Kaplan and R. O. Winder, "Cache Based Computer Systems", IEEE Computer, 6, 3, March, 1973, pp. 30-36.

Maurice Kendall and Alan Stuart, The Advanced Theory of Statistics, Volume 3, Design and Analysis and Time Series, third edition, Hafner Press, New York, 1976.

P. A. W. Lewis and G. S. Shedler, "Empirically Derived Micromodels for Sequences of Page Exceptions", IBM J. Res. Devel., 17, 2 March, 1973, pp. 86-100.

P. A. W. Lewis and G. S. Shedler, "Statistical Analysis of Non-stationary Series of Events in a Data Base System", IBM J. Res. and Devel., 20, 5, September, 1976, pp. 465-482.

Ron Revelle, "An Empirical Study of File Reference Patterns", IBM Research Report RJ 1557, April, 1975.

Alan Jay Smith, "Characterizing the Storage Process and Its Effect on the Update of Main Memory by Write-Through", 1976a, to appear, JACM.

Alan Jay Smith, "A Modified Working Set Paging Algorithm", IEEETC, C-25, 9, September, 1976b, pp. 907-914.

Alan Jay Smith, "Sequential Program Prefetching in Memory Hierarchies", 1977a, to appear, Computer.

Alan Jay Smith, "A Comparative Study of Set Associative Memory Mapping Algorithms and Their Use for Cache and Main Memory",

IEEETSE, SE-4, 2, March, 1978a, pp. 121-130.

Alan Jay Smith, "Bibliography on Paging and Related Topics", submitted for publication, August, 1978b.

Alan Jay Smith, "Long Term File Migration, Part II - File Replacement Algorithms", submitted for publication, August, 1978c.

Alan Jay Smith, "Testing for the Bernoulli Process in Long Term File Reference Patterns", 1978d, in preparation.

Edward P. Stritter, "File Migration", Stanford University Computer Science Report STAN-CS-77-594, (Ph.D. Dissertation), January, 1977.

This report was done with support from the Department of Energy. Any conclusions or opinions expressed in this report represent solely those of the author(s) and not necessarily those of The Regents of the University of California, the Lawrence Berkeley Laboratory or the Department of Energy.

Reference to a company or product name does not imply approval or recommendation of the product by the University of California or the U.S. Department of Energy to the exclusion of others that may be suitable.

TECHNICAL INFORMATION DEPARTMENT
LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA
BERKELEY, CALIFORNIA 94720