

Mathematical Models of File Growth

C. H. C. LEUNG AND K. WOLFENDEN

Department of Computer Science, University College London, Gower Street, London WC1E 6BT

The number of records in a file system is often recognised as a key determinant of efficiency. For example, the performance of sequential processing is $O(N(t))$ and that of tree search is $O(\log N(t))$, where $N(t)$ is the number of records in the file at time t . In this paper, the growth behaviour of files is studied in terms of quite general record insertion and deletion characteristics, and the performance evolution of some of the common systems is analysed. The growth data of an actual system are compared with the model results and reasonable agreement is observed.

1. INTRODUCTION

The efficiency of a file system is often related to its size. As the size of a file grows, its efficiency deteriorates, since the amount of time required to locate an item in a larger file tends to be longer. As one of the chief requirements of file systems is to enable record insertion and deletion to be carried out in a simple and routine manner, the variation in file efficiency over time is therefore to be expected and such variation is especially pronounced for volatile systems. For many of the common access techniques, their performance is directly related to the number of records present. For example, in sequential processing, the average number of accesses incurred in locating a given record varies approximately⁷ as $N/2$, where N is the number of records in the file. In the case of random processing, the average access distance separating two randomly located records varies approximately¹⁴ as $N/3$. In these cases, each record is responsible for a certain amount of contribution to the overall performance penalty – e.g. in sequential processing, this is one-half of an access. Thus the record insertion and deletion mechanism, which governs the file size, has a direct bearing on time-dependent file behaviour. Now insertions and deletions are random events occurring in time and often follow definite statistical patterns. They are dependent on factors such as the nature of the file system, its usage pattern, and the business activities to which the file relates. For instance, in order processing applications, the insertion and deletion of order records – which is related to order placement and delivery – depend on the seasonality of the products, the buoyancy of the market, the availability of stock, and the speed with which the orders are processed. These underlying mechanisms will consequently shape the overall statistical structure to which the insertion and deletion processes conform. The growth pattern of a file system is therefore not unlike that of a software system² in that it is related to the behaviour of the underlying application the characteristics of which are governed by the wider laws of, for example, business and economic systems.

Growth and deterioration of file systems have been previously considered in Refs. 1, 3–5, 8–11, 13, 15 and 17–19. In Refs. 15, 17 and 19 abstract deterioration patterns are adopted as basic assumptions without explicit reference to the underlying record insertion and deletion patterns for the purpose of formulating suitable reorganisation procedures. In Refs. 1, 3, 4, 8–11, 13 and 18 the impact of record insertions and deletions on performance deterioration is explicitly taken into

account, and in these studies the number of records present in the system is generally recognised to be a key determinant of efficiency. They are focused either on the analysis of particular file structures^{1, 3, 4, 8, 10, 13, 18} or the determination of optimal reorganisation strategy,^{1, 9, 11} and the patterns of record insertion and deletion are invoked solely for the purpose of supporting an overall evaluation. The characteristics of record insertions and deletions do not form the focus of these studies and they take on only a supporting role there; accordingly assumptions concerning them, although not always unreasonable, are sometimes restrictive and are mostly adopted for tractability or convenience. In Refs. 1 and 3 insertions are given in terms of the actual number of records, and the statistical pattern of insertion over time is not considered. In Refs. 13 and 18, the insertion pattern over time is assumed to follow a homogeneous Poisson process and the deletion process for individual records is likewise taken to be Poisson with the same rate. Ref. 8 adopts a slightly more general deletion pattern by allowing the deletion rate to be different from the insertion rate, both of which remain to be homogeneous Poisson processes. Ref. 4 relaxes the Poisson deletion assumption by allowing the lifetime of individual records to be generally distributed but retaining the homogeneous Poisson insertion assumption. In Refs 9 and 11, the insertion of records is assumed to follow a renewal process in which the inter-insertion interval of records is assumed to be independent, identically distributed; the statistical distribution of record lifetime is not considered and deletion is only permitted at file reorganisation. In Ref. 10, the insertion process is taken to be a non-homogeneous Poisson process, but again the deletion mechanism is not considered. A particularly noteworthy study which specifically aims to model the time-dependent performance of files is Ref. 5, where stochastic diffusion models are employed to approximate growth behaviour. However, the dynamic pattern of record insertion and deletion over time there is not explicitly represented and the study primarily focuses on file degradation after a given number of transactions have been entered.

The need for studying the impact of dynamic record insertions and deletions on file evolution has been pointed out in Ref. 1, but as yet there is no systematic study which specifically analyses the growth behaviour of general file systems over time. It is the aim of the present paper to provide such a study. The model presented here is quite general and includes all the above characteristics as

special cases. In addition, the present study allows us to provide meaningful justification of some of the abstract deterioration patterns used in Refs. 15, 17 and 19. The detailed model is developed in the next section, and in Section 3, a comparison is made between the growth pattern of a real-life system and that obtained from the model.

2. MODEL FORMULATION

We suppose that the time period of interest commences at time $t = 0$ and we shall mainly focus on development in the time interval $(0, t)$. For $b \geq a > 0$, we denote by $N(a, b)$ the random number of records inserted in the time interval (a, b) . Since for most storage devices and file systems multiple simultaneous insertions are not permitted, we suppose that $\Pr[N(t, t+\epsilon) > 1] = o(\epsilon)$. We let $N(t)$ denote the number of insertions in $(0, t)$, and the binary variable $dN(t)$ denote the number of insertions in the interval $(t, t+dt)$. The rate of insertion at time t is denoted by

$$\lambda(t) = E[dN(t)]/dt, \quad (2.1)$$

and the measure of insertion clustering is specified by the joint average

$$\beta(t_1, t_2) dt_1 dt_2 = E[dN(t_1) dN(t_2)], \quad t_1 \neq t_2.$$

For most practical systems, there are usually 'rush hours' and 'slack periods' alternating one another, so that $\lambda(t)$ is often a periodic function. The insertion process $N(t)$ here is quite general: *it need not be Poisson, stationary or uncorrelated*. During its lifetime in a file, a record generally experiences different phases of activity. For example, in certain applications a newly created record would typically experience an appreciable level of activity at the initial stage and then its activity would diminish before its eventual deletion from the file. In general, we suppose that the effect on file behaviour at time t produced by a record inserted at time τ is represented by a non-negative random effect $h(t-\tau)$ independent of the process $N(t)$. Since a record can only exert influence on a file after its insertion, $h(t)$ is assumed to be zero for $t < 0$; its correlation property is specified by

$$R(t_1, t_2) = E[h(t_1) h(t_2)].$$

The curve in Fig. 1 may be regarded as an example of a realisation of such an effect. The record is inserted at $t = 0$ and is removed from the file at $t = t_1$. The cumulative effect of these contributions due to insertions in the time interval $(0, t)$ on the file is denoted by a random quantity $f(t)$, which either measures the performance index of




Figure 1.

interest or from which the performance indices of interest can be derived, i.e.

$$f(t) = \int_0^t h(t-\tau) dN(\tau).$$

For example, if $f(t)$ measures the number of records in the system (which corresponds to $h(t) = 1$ for $t > 0$), then the performance¹² of sequential and random processing is $O(f(t))$, and that of tree search is $O(\log f(t))$.

We note that record deletion characteristics are already built into the function h . (Sometimes, in order to take into account the initial condition of the file system at $t = 0$, a further function $f_0(t)$, which is related to the file condition immediately after initial loading but before the operation of the process $N(t)$, would need to be added to $f(t)$ to represent the overall behaviour; however, in most of our discussions we shall omit $f_0(t)$ for convenience.) From the additivity of expectations, it follows that

$$\begin{aligned} E[f(t)] &= \int_0^t E[h(t-x) dN(x)] = \int_0^t E[h(t-x)] E[dN(x)] \\ &= \int_0^t \lambda(x) E[h(t-x)] dx = \lambda(t) * E[h(t)], \end{aligned} \quad (2.2)$$

where $*$ signifies convolution.¹⁶ For ease of reference $E[f(t)]$ will be called the deterioration curve, and $E[h(t)]$, the activity curve. For example, in the case of a file without deletions,[†] if we take

$$E[h(t)] = \begin{cases} 1/K & t \geq 0 \\ 0 & t < 0, \end{cases}$$

then for $K = 1$, $E[f(t)]$ gives the number of records in the file; for $K = 2$, $E[f(t)]$ represents the performance of sequential processing; and for $K = 3$, $E[f(t)]$ presents the performance of random processing. From (2.1) and (2.2) we have

$$E[f(t)] = \frac{1}{K} \int_0^t \lambda(x) dx = \frac{1}{K} \int_0^t E[dN(x)],$$

i.e.

$$E[f(t)] = E[N(t)]/K.$$

In the special case where $\lambda(t) = \lambda_0$ (constant), we have

$$E[f(t)] = \lambda_0 t / K. \quad (2.3)$$

It is worth pointing out that this special curve in fact corresponds to the deterioration models used in Refs. 15 and 17. More generally, if the insertion process is a cosine wave with period T , i.e.

$$\lambda(t) = \lambda_0 + \lambda_1 \cos \omega t, \quad \lambda_0 \geq \lambda_1, \quad (2.4)$$

where $\omega = 2\pi/T$, then we have

$$E[f(t)] = \lambda_0 t / K + \lambda_1 \sin \omega t / (\omega K).$$

The effect of deletion may be represented by the random function h shown in Fig. 2, where t_1 is a random variable signifying the lifetime of the record in question. This random function can be represented by

$$h(t) = [U(t) - U(t-t_1)]/K, \quad (2.5)$$

where

$$U(t) = \begin{cases} 1 & t > 0 \\ 0 & t \leq 0 \end{cases}$$

[†] Although files without deletions may seem artificial, this is implemented in some IBM access methods (e.g. ISAM[®]) in which no physical deletion of record is allowed.




Figure 2.

and t_1 is distributed according to an arbitrary density $g(t)$. By the law of total expectation, the activity curve is given by, for $t > 0$,

$$\begin{aligned} E[h(t)] &= \frac{1}{K} \left[1 - \int_0^\infty U(t-x) g(x) dx \right] \\ &= \frac{1}{K} \left[1 - \int_0^t g(x) dx \right] \\ &= \bar{G}(t)/K, \end{aligned}$$

where

$$\bar{G}(t) = \int_t^\infty g(x) dx$$

is the complementary distribution of t_1 . The deterioration curve is therefore

$$E[f(t)] = \lambda(t) * \bar{G}(t) / K. \quad (2.6)$$

In the case of random deletion with rate μ , $g(t)$ has an exponential distribution, i.e.

$$g(t) = \mu \exp(-\mu t)$$

so that $\bar{G}(t) = \exp(-\mu t)$ giving

$$E[f(t)] = K^{-1} \int_0^t \lambda(t-x) \exp(-\mu x) dx.$$

If the insertion process is periodic with rate given by (2.4) then the deterioration curve is

$$\begin{aligned} E[f(t)] &= \frac{\lambda_0}{K\mu} (1 - \exp[-\mu t]) + \frac{\lambda_1}{K\sqrt{\mu^2 + \omega^2}} \{ \cos \\ &\quad (\omega t - \phi) - \exp(-\mu t) \cos \phi \}, \end{aligned} \quad (2.7)$$



where $\phi = \tan^{-1}(\omega/\mu)$. We note that the deterioration curve contains a component which oscillates at the same frequency as the insertion process. The phase difference of ϕ indicates that there is a time lag before the effect of peak insertion is fully felt, which can be as much as a quarter of a cycle for rapid oscillations coupled with a low deletion rate (since $\phi \rightarrow \pi/2$ as $\omega/\mu \rightarrow \infty$). For constant insertion rate (i.e. setting $\lambda_1 = 0$ in (2.7)), we obtain an exponential deterioration curve. It is interesting to note that a similar exponential deterioration curve was, in fact, assumed in Ref. 10. Although random deletion is an important special case, it is not always an appropriate model. In a situation where no record in the file can have age exceeding a given value α (e.g. they may have to be transferred to an offline history file for archiving), then a plausible lifetime distribution would be

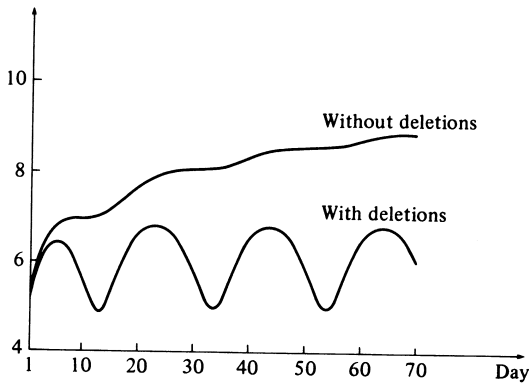
$$g(t) = \begin{cases} 1/\alpha & t \leq \alpha \\ 0 & t > \alpha, \end{cases}$$

which results in the following deterioration curve (with $\phi = \tan^{-1} \omega \alpha$)

$$E[f(t)] = \begin{cases} \frac{\lambda_0 t}{K} (1 - \frac{t}{2\alpha}) + \frac{\lambda_1}{K\omega^2 \alpha} [\sqrt{1 + \omega^2 \alpha^2} \cos(\omega t - \phi) - 1], & t \leq \alpha \\ \frac{\lambda_0 \alpha}{2K} + \frac{\lambda_1}{K\omega^2 \alpha} [\sqrt{1 + \omega^2 \alpha^2} \cos(\omega t - \phi) - \cos(\omega t - \omega \alpha)], & t > \alpha. \end{cases}$$

Figs 3 and 4 respectively show $E[f(t)]$ for this model and one in which there are no deletions. The parameters are $K = 1$, $\lambda_0 = \lambda_1 = 100$ (insertions per day) so that the average number of insertions ranges from 0 to 200 per day; the insertion period $2\pi/\omega$ is 20 days (i.e. $\omega = 0.1\pi$) and the maximum record lifetime for Fig. 3 is 10 days. Figs 3 and 4 show the development over a 70-day period. We see that there are a number of striking differences between the two. In Fig. 3 the deterioration curve, after the initial transient rise, settles to an equilibrium and oscillates in a regular pattern. It obviously varies with time, but the pattern of change remains the same as time goes on. Fig. 4, however, shows a steadily increasing pattern but the rate of growth is non-uniform. Unlike that in Fig. 3, the curve in Fig. 4 does not settle to an equilibrium but grows arbitrarily large as time goes on. The curve $\log_e E[f(t)]$ for these models are shown in Fig. 5. For the case with deletions, it behaves similarly to Fig. 3 and oscillates at the same frequency except that the peaks are slightly flattened and the troughs sharpened. For the other case, it rises swiftly at first, but unlike Fig. 4 the increase becomes rather heavily damped, which indicates relative insensitivity to fluctuation in growth rate.


 Figure 3. Deterioration curve $E[f(t)]$ – with deletions.

 Figure 4. Deterioration curve $E[f(t)]$ – without deletions.

Figure 5. Deterioration of $\log_e E[f(t)]$

More generally, if $\lambda(t)$ is an arbitrary non-negative periodic function with period T , then it can be represented by a Fourier series

$$\lambda(t) = \sum_{n=-\infty}^{\infty} c_n \exp(in\omega t)$$

with $\omega = 2\pi/T$ and $i = \sqrt{-1}$. Hence

$$\begin{aligned} E[f(t)] &= \sum_{n=-\infty}^{\infty} c_n \int_0^t \exp[in\omega(t-x)] E[h(x)] dx \\ &= \sum_{n=-\infty}^{\infty} c_n \exp(in\omega t) \int_0^t \exp(-in\omega x) E[h(x)] dx. \end{aligned}$$

If we denote the 'finite' Fourier transform

$$\int_0^t \exp(-i\omega x) E[h(x)] dx$$

by $A(\omega) \exp[i\phi(\omega)]$, then from the properties of the Fourier representation,¹⁶ we have

$$E[f(t)] = c_0 A(0) + 2 \sum_{n=1}^{\infty} \operatorname{Re}\{c_n A(n\omega) \exp(i[n\omega t + \phi(n\omega)])\},$$

where Re signifies the real part of a quantity.

For example, in the case of random deletion, we have

$$A(\omega) = \frac{1}{K(\mu^2 + \omega^2)} [(\mu - \mu \exp(-\mu t) \cos \omega t + \omega \exp(-\mu t) \sin \omega t)^2 + (\omega \cos \omega t + \mu \sin \omega t - \omega)^2]^{\frac{1}{2}},$$

and

$$\phi(\omega) = \tan^{-1} \left[\frac{\omega \cos \omega t + \mu \sin \omega t - \omega}{\mu - \mu \exp(-\mu t) \cos \omega t + \omega \exp(-\mu t) \sin \omega t} \right].$$

Thus far, we have been concerned only with the mean value function $E[f(t)]$. Of course, the autocovariance function $C(t_1, t_2) = \operatorname{Cov}[f(t_1), f(t_2)]$, which measures the association between the values of $f(t)$ at different times, is also of considerable interest. Since the random function $h(t)$ associated with different records is assumed independent, we have for $t_1 \leq t_2$

$$\begin{aligned} E[f(t_1)f(t_2)] &= \int_0^{t_2} \int_0^{t_1} E[h(t_1-x)] E[h(t_2-y)] \beta(x, y) dx dy \\ &\quad + \int_0^{t_1} R(t_1-x, t_2-x) \lambda(x) dx, \end{aligned}$$

where the first term is related to the contributions from different records and the second term is related to the effect produced by the same record at different times. In

the important special case where insertions follow a non-homogeneous Poisson process, then $\beta(x, y) = \lambda(x)\lambda(y)$ and so

$$C(t_1, t_2) = E[f(t_1)f(t_2)] - E[f(t_1)] E[f(t_2)] = \int_0^{t_1} R(t_1-x, t_2-x) \lambda(x) dx. \quad (2.8)$$

If $h(t)$ is given by (2.5), then for $x \leq y$

$$\begin{aligned} R(x, y) &= E[h(x)h(y)] = K^{-2} \operatorname{Pr}[h(x) = h(y) = K^{-1}] \\ &= K^{-2} \operatorname{Pr}[h(y) = K^{-1}] = K^{-2} \operatorname{Pr}[\text{record lifetime} \geq y] \\ &= K^{-2} \bar{G}(y), \end{aligned}$$

which gives

$$C(t_1, t_2) = K^{-2} \int_0^{t_1} \bar{G}(t_2-x) \lambda(x) dx,$$

and the variance is

$$\operatorname{Var}[f(t)] = C(t, t) = K^{-2} \int_0^t \bar{G}(t-x) \lambda(x) dx.$$

If $\lambda = \lambda_0$, then for random deletion these give

$$\begin{aligned} C(t_1, t_2) &= \frac{\lambda_0}{\mu K^2} \{\exp[-\mu(t_2-t_1)] - \exp(-\mu t_2)\}, \\ \operatorname{Var}[f(t)] &= \frac{\lambda_0}{\mu K^2} (1 - \exp[-\mu t]). \end{aligned} \quad (2.9)$$

If $h(t) = A\psi(t)$ where A is an independent random variable and $\psi(t)$, a fixed function – this corresponds to a situation where the records all have similar activity patterns but with different amplitude A – then

$$R(x, y) = \psi(x) \psi(y) E(A^2)$$

and so

$$C(t_1, t_2) = E(A^2) \int_0^{t_1} \psi(t_1-x) \psi(t_2-x) \lambda(x) dx$$

and

$$\operatorname{Var}[f(t)] = E(A^2) \psi^2(t) * \lambda(t). \quad (2.10)$$

We note that in (2.9), (2.10) and more generally in (2.8) $\operatorname{Var}[f(t)]$ increases with time, which indicates that the reliability of the mean function $E[f(t)]$ diminishes as the time horizon increases.

3. EMPIRICAL ASSESSMENT

In this section we shall assess the reasonableness of our model using the growth data of an actual system. Fig. 6 shows the growth data of the Reading University library system over a ten-year period from 1971/72 to 1980/81. The data are related to the number of printed volumes in the library and, since their withdrawal is very rare, the behaviour of the system can be approximated by means of a 'pure' growth mechanism without deletions. On taking λ_0 to be 20000/year and adding the constant $f_0 = 390000$ (to account for the initial system condition) to equation (2.3), we obtain the dotted line in Fig. 6. We observe that, although some discrepancy is evident, the overall agreement is quite good.

These empirical data represent growth on a relatively long-term (yearly) basis which probably mask most of the short-term fluctuations; in a university environment it is likely that some degree of seasonal variation exists, so that the growth pattern of Fig. 4 may possibly provide




Figure 6. Growth of an actual system.

a better representation of the system's evolution, but no monthly or quarterly data on the system are available to test this. Admittedly, the library file system is not very typical of most file systems in that there is very little deletion, and it is relatively non-volatile. Comparison between some of the more variable growth patterns of the previous section with shorter-term growth data collected on actual systems, especially volatile ones, will undoubtedly be valuable.

4. SUMMARY AND CONCLUSION

The growth and deterioration of file systems are analysed using a general record insertion model which incorporates

time variation and clustering effects. After insertion, the time-dependent behaviour of individual records is represented by an activity curve which allows the characteristics of deletions to be specified; such an activity curve not only permits the record lifetimes to be entirely generally distributed, but also enables the varying phases of record activity, which is typical of records in most real-life systems, to be quantitatively taken into consideration. The file evolution patterns for some of the common systems are studied and we find that, even for seemingly simple systems, such patterns can be quite complex. For example, the performance of certain systems could oscillate indefinitely over time while remaining in a state of equilibrium. The commonly adopted linear and exponential deterioration characteristics, while adequate for systems with fairly static insertion rate, are not representative of the behaviour of general systems. The growth data of an actual system have been compared with those obtained from the model, and reasonable agreement is observed.

Acknowledgements

The authors are grateful to Dr J. D. Roberts and H. Longe for valuable discussions and comments on an earlier version of this work, and to L. R. M. Bishop for providing the data for this study. This paper also benefited from the helpful comments and suggestions of the referee.

REFERENCES

1. D. S. Batory, Optimal file designs and reorganization points. *ACM Transactions on Database Systems* 7, 60-81 (1982).
2. L. A. Belady and M. M. Lehman, An introduction to growth dynamics. In *Statistical Computer Performance Evaluation*, edited by W. FREIBERGER. Academic Press, New York (1972).
3. Y. H. Chin, An analysis of distributed free space in an operating and data management systems environment. *IEEE Transactions on Software Engineering* SE-4, 436-440 (1978).
4. R. Fagin, J. Nievergelt, N. Pippenger and H. R. Strong, Extendible hashing - a fast access method for dynamic files. *ACM Transactions on Database Systems* 4, 315-344 (1979).
5. D. P. Heyman, Mathematical models of database degradation. *ACM Transactions on Database Systems* 7, 615-631 (1982).
6. IBM, *Introduction to IBM Direct-Access Storage Devices and Organization Methods*, order no. GY28-6607-9 (1976).
7. D. E. Knuth, *The Art of Computer Programming, Volume 3: Sorting and Searching*. Addison-Wesley, Reading, MA (1973).
8. P. A. Larson, Analysis of index-sequential files with overflow chaining. *ACM Transactions on Database Systems* 6, 671-680 (1981).
9. G. M. Lohman and J. A. Muckstadt, Optimal policy for batch operations: backup, checkpointing, reorganization, and updating. *ACM Transactions on Database Systems* 2, 209-222 (1977).
10. K. Maruyama and S. E. Smith, Optimal reorganisation of distributed free space files. *Communications of the ACM* 19, 634-642 (1976).
11. H. Mendelson and U. Yechiali, Optimal policies for data base reorganisation. *Operations Research* 29, 23-26 (1981).
12. J. Nievergelt, Binary search trees and file organisation. *Computing Surveys* 6, 196-207 (1974).
13. C. A. Olson, Random access file organisation for indirectly addressed records. *Proceedings of the 1969 ACM National Conference*, 539-549.
14. W. F. Sharpe, *The Economics of Computers*. Columbia University Press, New York (1969).
15. B. Shneiderman, Optimum data base reorganization points. *Communications of the ACM* 16, 362-365 (1973).
16. R. D. Stuart, *An Introduction to Fourier Analysis*. Chapman and Hall, London (1961).
17. W. G. Tuel, Optimum reorganisation points for linearly growing files. *ACM Transactions on Database Systems* 3, 32-40 (1978).
18. J. A. van der Pool, Optimum storage allocation for a file in steady state. *IBM Journal of Research and Development* 17, 27-38 (1973).
19. S. B. Yao, K. S. Das and T. J. Teorey, A dynamic database reorganisation algorithm. *ACM Transactions on Database Systems* 1, 159-174 (1976).