Estimating Disc Access Patterns using Diffusion Models

C. H. C. LEUNG AND K. WOLFENDEN
Department of Computer Science, University College London, Gower Street, London WC1E 6BT

Computer disc accesses necessitate the mechanical movement of a read/write head over a sequence of storage locations. The pattern of such random movement has a significant bearing on access times and its estimation is of primary importance to the meaningful prediction of performance. The procedure presented, which is based on a random walk description, allows such estimates to be derived quickly, incorporating any prior knowledge and partial information on data characteristics. Although only an approximation, it is able to produce good agreement with published measurements.

1. INTRODUCTION

One of the chief reasons for installing computer systems is the amount of data the computer can handle, and the speed and flexibility with which the data can be retrieved. In present-day systems the commonest form of mass storage medium is the magnetic disc. Although disc time is of the order of 10^{-2} seconds, it is considered slow compared with other components of the computer system whose speed is typically measured in units of 10^{-6} seconds. The reason is that discs depend on *mechanical positioning* rather than *electronic switching*. The bulk of the disc access delay is due to the mechanical movement of the read/write head over the data recording surface. A model of the disc is shown diagrammatically in Fig. 1;

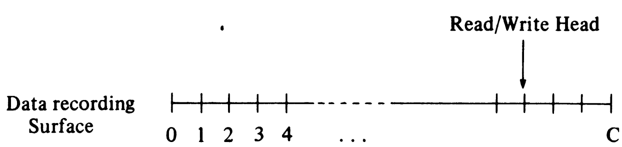

Fig. 1.

here the read/write head has to move physically among the positions $\{0, 1, \ldots, C\}$ for data processing. Such pattern of movement can be highly diverse and is dependent on the random requirements of data access. Owing to the slow speed of the read/write head, it is a frequent source of performance bottleneck, and the estimation of its motion pattern is of vital importance to performance prediction. From the representation in Fig. 1 it does seem plausible, as a first approximation at any rate, to describe the movement of the read/write head as a one-dimensional random walk on the coordinates $\{0, 1, \ldots, C\}$; in fact, such an analogy is first noticed by Sharpe (1969).[1] However, this simple analogy has not yet been explored further and there appears to be no development to date which exploits the relatively well-developed theory of random walk and related stochastic processes to study disc behaviour. It is the purpose of the present note to report such a development and, in particular, to demonstrate how such methodology could be practically applied to the estimation of disc access patterns.

The present approach enables the analyst to express his knowledge of data characteristics in quantitative terms and allows approximate performance estimates to be derived quickly and with little dependence on information which is mostly unavailable at the time of estimation. Such knowledge may be subjective or incomplete, but it must be incorporated in the resultant estimates without which meaningful prediction of performance cannot proceed. Compared with empirical measurements, the present approach is able to produce good agreement. It offers a practical and efficient alternative to the computation intensive method proposed in Leung and Wolfenden;[2] these two approaches incorporate different degrees of detail and are complementary to each other.

2. DIFFUSION REPRESENTATION

Here we shall replace the discrete storage positions by the continuous interval $I = [0, C]$. Such approximation is normally acceptable as the number of positions in most devices is large (≈ 1000), and is a commonly adopted approximation (see e.g. Ref. 1). For the moment let us suppose that the motion of the read/write head is deterministic. Then it may be possible to describe its movement by means of a deterministic equation of motion. For example, if $x(t)$ denotes the position of the read/write head at time t, then ignoring end effects the motion may be described by the following differential equation

$$\dot{x}(t) - \mu(x, t) = 0,$$

where a dot signifies differentiation with respect to time. Thus the behaviour of the system can be determined on knowledge of $\mu(x, t)$, which may be interpreted as the speed of the read/write head. Analogously, in the case of stochastic movement, one might introduce randomness by means of a random noise term $n(t)$ so that the equation of motion becomes

$$\dot{x}(t) - \mu(x, t) = n(t),$$

which is now a stochastic differential equation. The noise $n(t)$, for example, could be a white Gaussian noise with zero mean, and if, in addition, $\mu(x, t)$ is constant, then it can be shown that in this case $x(t)$ corresponds to a Brownian motion process with drift μ, which is a limiting version of a simple random walk.[3] More generally, diffusion processes are describable by means of stochastic differential equations of the form

$$\dot{x} - \mu(x\ t) = \dot{y}(x, t),$$

where $y(x, t)$ is an additive process.[3] Such stochastic *equation of motion* can be shown (Papoulis (1965)) to correspond to the following *distribution equation*

$$\frac{1}{2}\frac{\partial^2}{\partial x^2}[\sigma^2(x, t)f(x, t)] - \frac{\partial}{\partial x}[\mu(x, t)f(x, t)] = \frac{\partial f(x, t)}{\partial t}, \qquad (2.1)$$

where $f(x, t)$ is the density function of the position of the read/write head at time t under appropriate initial conditions. The coefficients $\mu(x, t)$ and $\sigma^2(x, t)$ admit the following intuitive interpretation

$$\mu(x, t)\,dt = E[dx|x]$$
$$\sigma^2(x, t)\,dt = E[dy^2|x] = E[(dx-\mu dt)^2|x],$$

i.e. μ may be regarded as the local average rate of displacement and σ the corresponding standard deviation. As we shall see below, it is through such intuitive characterisation that meaningful estimates of storage performance are constructed. Thus far, we have considered only the unrestricted movement of the read/write head. However, since its movement is always restricted to the interval I, it is reasonable to suppose that the boundary conditions of reflecting barriers apply at the points $x = 0$ and $x = C$; these conditions are expressible mathematically as[4]

$$\frac{1}{2}\frac{\partial}{\partial x}[\sigma^2(x, t)f(x, t)] = \mu(x, t)f(x, t) \quad x = 0, C. \quad (2.2)$$

The time-dependent solution $f(x, t)$ of (2.1) and (2.2) is in general very difficult to obtain. The equilibrium solution is much easier to obtain, however, and is also of considerable practical interest. Removing all time dependencies in (2.1) and (2.2), they become quite manageable ordinary differential equations. It can be verified that the resultant (equilibrium) solution is

$$f(x) \propto \frac{1}{\sigma^2(x)}\exp\left[2\int\frac{\mu(x)}{\sigma^2(x)}dx\right]. \quad (2.3)$$

Hence, once $\mu(x)$ and $\sigma(x)$ are specified the resultant behaviour of the read/write head can be determined in a straightforward manner. For instance, if there is no preferred direction of motion, then we can take $\mu(x) = 0$, which is an appropriate structure for situations where virtually no information concerning data characteristics is available; in this case, we have from (2.3)

$$f(x) \propto 1/\sigma^2(x).$$

If, in addition, we regard the disc surface as homogeneous so that $\sigma^2(x) = \sigma^2$ (a constant), then we have $f(x)$ = constant. It is worth pointing out that this special estimate represents a widely adopted pattern for *random access*; for example, it is used in Coffman and Denning,[5] Martin[6] and Sharpe.[1] Thus, by adopting suitable forms for $\mu(x)$ and $\sigma^2(x)$, a wide variety of storage processing characteristics could be built into the resultant estimate. The next section provides an empirical assessment of this approach.

3. EMPIRICAL ASSESSMENT

Here we consider a disc storage with an activity centre located at position $b \in (0, C)$ towards which the read/write head tends to drift. A possible way of representing this knowledge is to let $\mu(x) = -\beta(x-b)$, where β is a constant. If we again take σ^2 to be constant, then (2.3) gives, for $x \in$ I,

$$f(x) \propto \exp\left\{-\frac{1}{2}\left(\frac{x-b}{k}\right)^2\right\},$$

where $k^2 = \sigma^2/(2\beta)$. The separation between two independent samples X_1, X_2 from $f(x)$ gives the seek distance

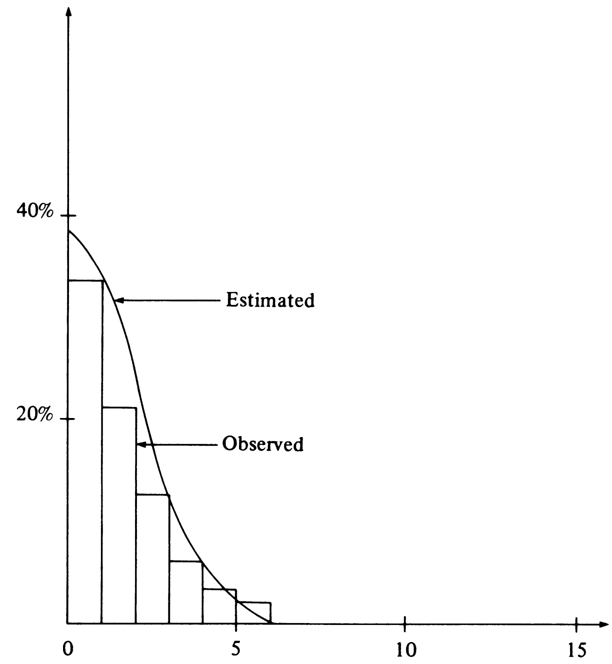

Fig. 2.

$S = |X_1 - X_2|$, and it is shown in Appendix A that the density of S is approximately equal to

$$g(s) = (k\sqrt{\pi})^{-1}\exp\{-s^2/(4k^2)\}, \quad s \geq 0. \quad (3.1)$$

In Ref. 7 measurements are taken on a disc storage having 405 positions. There is an activity centre consisting of five highly localised positions which absorb nearly 75% of the total activity. Fig. 2 compares the measurements published there with the estimate (3.1) for $k = 1.45$. (In Fig. 2 it is not possible to represent seek distances greater than 15 because of the scale of the diagram, but these are insignificant as none of them absorbs more than 0.8% of the activity.) We see that, although (3.1) does not yield perfect agreement with the empirical data, it does not seem to be unreasonable as an approximation. Of course, one should not expect to obtain a very close fit with this method since, at the outset, it is not intended to be a faithful reflection of the full situation: it aims to provide a sufficiently meaningful estimate using only a small fraction of the computational effort. (The method proposed in Ref. 2 would undoubtedly give a closer fit but it requires fine details and the manipulation of a 405×405 matrix.)

4. SUMMARY AND CONCLUSION

A novel procedure based on a diffusion approximation for the estimation of disc access patterns is presented. The procedure is efficient in that it allows estimates to be very simply derived from the salient features and prior knowledge relating to data characteristics. Although this approach involves a certain degree of approximation and does not purport to be a totally faithful and accurate reflection of reality, it does seem to agree reasonably well with published empirical data. It furnishes a pragmatic and credible tool for the meaningful prediction of disc performance, and offers a practical alternative to rival computation-intensive methods.

Acknowledgements

The authors are grateful to Professor B. W. Conolly and D. Girmes for their valuable comments and suggestions on a previous version of this paper. This paper also benefited from the helpful comments of the referee.

REFERENCES

1. W. F. Sharpe, *The Economics of Computers*. Columbia University Press, New York (1969).
2. C. Leung and K. Wolfenden, Disc database efficiency: a scheme for detailed assessment based on semi-Markov models. *The Computer Journal* **26**, 10–14 (1983).
3. A. Papoulis, *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, New York (1965).
4. D. R. Cox and H. D. Miller, *The Theory of Stochastic Processes*. Chapman and Hall, London (1965).
5. E. G. Coffman and P. J. Denning, *Operating Systems Theory*. Prentice-Hall, New Jersey (1973).
6. J. Martin, *Design of Real-Time Computer Systems*. Prentice-Hall, New Jersey (1967).
7. L. Boi *et al.*, A performance evaluation of the CII SIRIS 8 operating system – methodology, tools and first results. In *Modelling and Performance Evaluation of Computer Systems*, edited H. Beilner and E. Gelenbe, pp. 67–80. North-Holland, Amsterdam (1976).
8. M. G. Kendall and A. Stuart, *The Advanced Theory of Statistics*, 3rd edn, vol. I. Griffin, London (1969).

APPENDIX A. DERIVATION OF EQUATION (3.1)

It is evident from the form of the density $f(x)$ that it is a truncated normal density. When the disc surface is large, then it may be approximated by the ordinary normal density with mean b and standard deviation k. The seek distance S is simply the range of two independent samples from such a density. From Ref. 8, the density of S is given by

$$g(s) = \int_{-\infty}^{\infty} 2f(x)\,f(x+y)\,dx.$$

Substituting the normal density for $f(x)$ in the above, and applying the following formula (for constants a, b and c)

$$\int_{-\infty}^{\infty} \exp\{-(ax^2+bx+c)\}\,dx = \sqrt{(\pi/a)}\exp\{(b^2-4ac)/(4a)\},$$

we obtain (3.1) after simplification.