Title
HOPT: A MYOPIC VERSION OF THE STOCHOPT AUTOMATIC FILE MIGRATION POLICY

Author
Olken, F.

Publication Date
1983-02-01

LBL-15554 c.2

Lawrence Berkeley Laboratory
UNIVERSITY OF CALIFORNIA

Physics, Computer Science & Mathematics Division

Submitted to the Association of Computing Machinery (ACM) Sigmetrics Conference on Measurement and Modeling of Computer Systems, Minneapolis, MN, August 29-31, 1983

HOPT: A MYOPIC VERSION OF THE STOCHOPT AUTOMATIC FILE MIGRATION POLICY

Frank Olken

February 1983

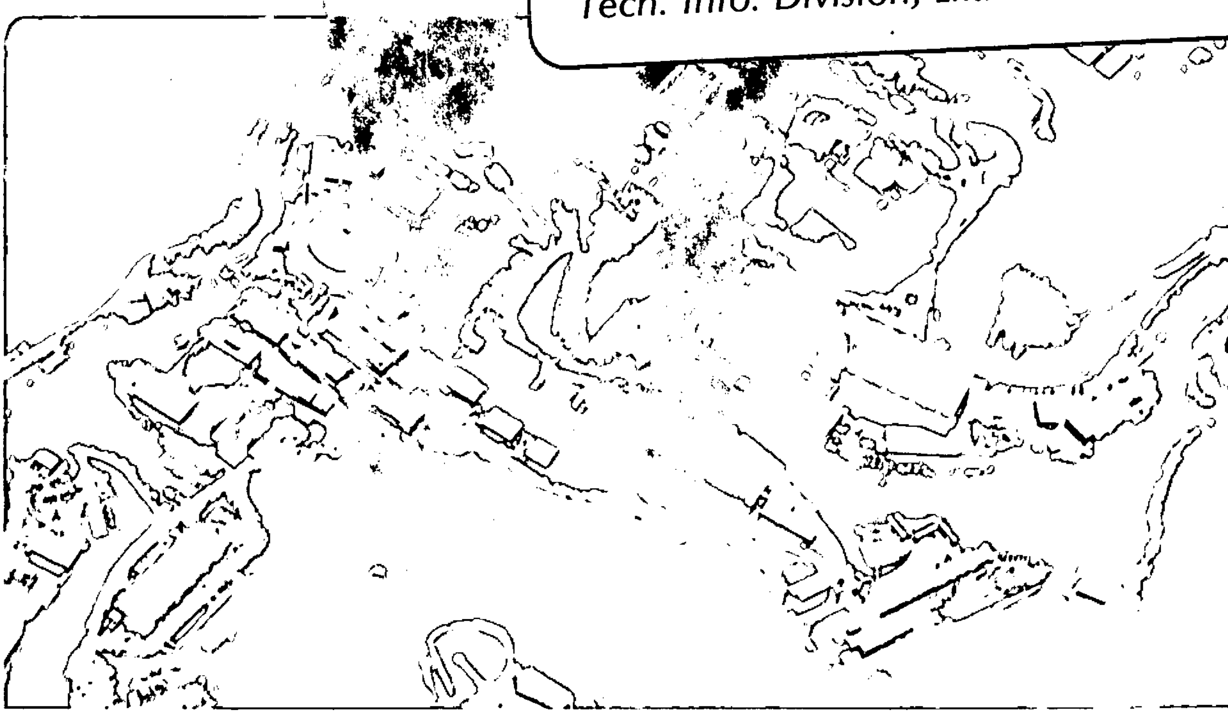

DISCLAIMER

LBL-15554

PROGRES-83.3

HOPT: A Myopic Version of the STOCHOPT Automatic File Migration Policy

Frank Olken

Computer Science and Mathematics Dept.
Lawrence Berkeley Laboratory
University of California
Berkeley, California 94720

February 7, 1983

This work was supported by the Director, Office of Energy Research, Office of Basic Research Sciences, Division of Engineering, Mathematical and Geosciences of the U.S. Department of Energy under Contract DE-AC03-76SF00098.

This manuscript was printed from originals provided by the author.

HOPT: A Myopic Version of the STOCHOPT Automatic File Migration Policy

ABSTRACT

We consider the application of the STOCHOPT automatic file migration policy (proposed by A.J. Smith) to a file system in which the file inter-reference time distributions are characterized by strictly monotonically decreasing hazard rates (also known as decreasing failure rates). We show that the STOCHOPT policy can be simply stated in terms of a scaled hazard rate (i.e., the hazard rate divided by the file size). The class of decreasing failure rate inter-reference time distributions includes mixtures of exponential distributions which are the continuous time analogues of the mixed geometric distributions proposed by Smith to model file inter-reference times in discrete time.

1. Introduction

We are concerned with replacement policies for deciding when to remove a file from a disk resident cache and send it to a teritary storage device (e.g. an automatic tape library).

In [Smith81b, pg 526] A.J. Smith proposed a stochastically optimal replacement policy (STOCHOPT) for automatic file migration. The model for which the policy is optimal is specified by known file inter-reference time probability distributions, $F_i(t)$, a unit cost for all file faults, and a positive constant rental price per unit of storage space, ϑ. The policy is specified in terms of the cache holding time, τ_i, for the file i which minimizes the expected cost of the next reference.[1] Smith considered the model in discrete time because of the discrete nature of the dataset he was studying [Smith81a]. We obtain our results by studying the problem in continuous time.

In reliability theory [Barlow75] it is commonplace to characterize probability distributions by their failure rate. The failure rate at time t (also known as the hazard rate) is the rate at which units which have survived until time t fail. We discuss this further in section 2. We shall show that if the file inter-reference time distribution has a strictly monotonically decreasing hazard rate the STOCHOPT policy can be simply stated in terms of an inequality on a scaled hazard rate (i.e., the hazard rate divided by the file size.) We call this policy HOPT (for Hazard Optimal). The policy is myopic (i.e., it does not explicitly look at the future behavior of the file, only its current hazard rate). Hence it lends itself to fixed space formulation, ranking files on the basis of their scaled hazard rates. We call the fixed space policy HMAX (for Hazard Maximal).

---

[1] Actually Smith uses a unit storage rental charge and a variable cost for tape accesses. The two formulations are equivalent, all that matters is the ratio between storage rental charge and tape access cost. The present formulation will simplify our discussion.

We begin by describing a continuous time version of Smith's model. We allow the file inter-reference time distribution to be an improper distribution (i.e., one which does not integrate to 1). This permits us to model the possibility that a file may never be referenced again (See Sect. 4). We also allow the cost of accessing a file from tape to depend on the file characteristics (e.g. size). These extensions to Smith's model make the model more realistic. We share Smith's assumptions that the file reference processes for distinct files are independent and that the cost of accessing a particular file from tape is invariant (e.g. with respect to time since last access).[2]

The expected cost until next reference for file i is comprised of three components:

(1) If the file is referenced before cache holding time τ there is only the cost of keeping it in the cache until it is referenced.

(2) If the file is not referenced before time τ there is the cost of keeping the file in the cache for τ seconds.

(3) If the the file is not referenced before time τ, but it is referenced again then there is the cost of accessing the file from tape.

Thus we define:

$C_i(\tau)$ = *expected cost until next reference for file i, holding file on disk* τ *seconds*

$$= \vartheta * s_i \left[\int_0^{\tau} u f_i(u)\, du + (1-F_i(\tau))\tau \right] + (F_i(\infty)-F_i(\tau)) * C_{TA}(i)$$

where

$C_{TA}(i)$ = *cost to access file i from tape*

$f_i(u)$ = *probability density function of inter-reference time interval for file i*

$F_i(u)$ = *cummulative distribution function of inter-reference time interval for file i*

$F_i(\infty)$ = *probability that file i is referenced again*

---

[2] Some file accesses may overwrite the entire file (and thus do not need to access the file from tape if it is not in the cache). We can model this by setting the tape access cost for such references to zero. Then the expected cost for a file access from tape depends on the ratio of reads to complete overwrites. We are implicitly assuming that this ratio is not dependent on the time since last reference. The dataset used by Smith [Smith81a] does not identify overwrite references.

s_i = *size of file i*

ϑ = *rental charge per unit space per unit time*

Note that we are measuring time here since the last reference to file i. This cost function corrresponds to that of Eqn. (4) of [Smith81b, pg. 526].

2. Monotone Decreasing Hazard Rate Distributions

The hazard rate (failure rate) is essentially a conditional probability density. It is the rate at time t since last reference at which files which have not yet been referenced are accessed. We define the hazard rate as $h_i(t) = \frac{f_i(t)}{1-F_i(t)}$.

The utility of the HOPT policy hinges upon the question of whether empirically observed file inter-reference time distributions have strictly monotonic decreasing hazard rates (SDHR). The evidence concerning this property is encouraging. Smith [Smith81a, pg 411] remarks that the hazard rate "declines sharply for a while and then becomes (after 20 days or so) relatively flat".

In reliability literature [Barlow75] hazard rates are called "failure rates". Distributions with monotonically decreasing hazard rates (DHR) are called "decreasing failure rate" (DFR) distributions by Barlow and Proschan [Barlow75]. In their usage DFR requires neither the existence of the probability density density nor strict monotonicity, i.e., they include the exponential distribution with a constant hazard in the DFR class. Thus all SDHR distributions are DHR, and all DHR distributions are DFR. The same authors state that if $F(t)$ is DHR for $t \geq 0$ then $f(t)>0$ for *all* $t>0$ [Barlow75, pg.79].

Barlow and Proschan show that all mixtures of exponential distributions are strictly DFR [Barlow75, pg 103]. Discrete mixtures of exponential distributions, are the continuous time analogues of the discrete time mixed geometric distributions used to model file inter-reference times in [Smith81a, pg 411]. Smith used method of moments estimation for fit his mixed geometric model in

[Smith81a]. Maximum likelihood estimation of parameters mixtures of exponentials is discussed in [Jewell82] and in the references cited therein.

3. HOPT Policy

We now show how STOCHOPT can be characterized in terms of a scaled hazard rate when $F_i(t)$ has a strictly decreasing hazard rate. The resulting policy we designate **HOPT.**

Define

$$\gamma_i(t) = \textit{scaled hazard rate for file } i, \textit{ as}$$

$$\gamma_i(t) = \frac{h_i(t) * C_{TA}(i)}{s_i}$$

Theorem (HOPT Policy)

If $F_i(t)$ is SDHR then $C_i(\tau)$ is minimized when $\gamma_i(\tau) = \vartheta$, if such τ exists, i.e., when the scaled hazard rate equals the storage rental rate. Otherwise, if $\gamma_i(t) > \vartheta$ *for all* $t \geq 0$, then the optimal $\tau = \infty$. If $\gamma_i(t) < \vartheta$ *for all* $t \geq 0$, then the optimal $\tau = 0$.

Proof

We proceed by calculating the derivative of $C_i(\tau)$ in terms of the hazard rate.

$$\frac{d}{d\tau} C_i(\tau) = \vartheta * s_i * \left[\tau f_i(\tau) + [(1-F_i(\tau)) - \tau * f_i(\tau)]\right] - f_i(\tau) * C_{TA}(i)$$

$$= \vartheta * s_i * (1-F_i(\tau)) - f_i(\tau) * C_{TA}(i)$$

$$= \vartheta * s_i * (1-F_i(\tau)) \left[1 - \frac{f_i(\tau) * C_{TA}}{\vartheta * s_i * (1-F_i(\tau))}\right]$$

We can write this in terms of the hazard rate as:

$$\frac{d}{d\tau} C_i(\tau) = \vartheta * s_i * (1-F_i(\tau)) \left[1 - \frac{h_i(\tau) * C_{TA}(i)}{\vartheta * s_i}\right]$$

Recall that $F_i(t)$ is SDHR implies that the density $f_i(t)$ is nonzero over the nonnegative real line, i.e., $(1-F_i(\tau)) > 0$ for all $\tau \geq 0$. Then $\frac{d}{d\tau}C_i(\tau) < 0$ whenever $\frac{h_i(\tau) * C_{TA}(i)}{s_i} > \vartheta$. We can restate this as $\frac{d}{d\tau}C_i(\tau) < 0$ whenever $\gamma_i(\tau) > \vartheta$. Now if the hazard rate $h_i(t)$ is a strictly monotonically decreasing function of t then we have

$$\gamma_i(T) \geq \vartheta \quad \Rightarrow \quad \frac{d}{d\tau}C_i(\tau) < 0 \; for\ all\ \tau < T$$

and

$$\gamma_i(T) \leq \vartheta \quad \Rightarrow \quad \frac{d}{d\tau}C_i(\tau) > 0 \; for\ all\ \tau > T$$

Hence conclude that the minimum cost is achieved by setting $\tau = T$ where $\gamma_i(T) = \vartheta$, if such T exists, otherwise $\tau = 0$ if $\gamma_i(t) < \vartheta$ *for all* $t \geq 0$, or $\tau = \infty$ if $\gamma_i(t) > \vartheta$ *for all* $t \geq 0$.

4. Improper Distributions

The reader will recall that our model permitted $F_i(t)$ to be an improper distribution. Smith only considered proper distributions for the file inter-reference times. A careful reading of [Smith81a] indicates that he discarded censored observations (presumably including infinite file inter-reference intervals) fitting his model only to the uncensored observations.[3] Anecdotal evidence from operators of large mass storage systems (e.g. at Lawrence Livermore National Lab) suggests that essentially infinite file inter-reference intervals are common. Users simply treat the filesystem as an archive, retaining dead files for backup. We propose to incorporate such infinite intervals into the model by scaling proper inter-reference time distributions down to an improper distribution. Thus we define $F_i(t) = F_i(\infty) * G_i(t)$ where $G_i(t)$ is a proper probability distribution. We shall show below that, if $G_i(t)$ is DHR, then $F_i(t)$ is strictly DHR. Thus HOPT will work for scaled (improper) versions of DHR distributions.

[3] See [Smith81a, pg. 405] for Smith's definitions of $I(i,j)$ and $C(i,j)$ and equations 10 through 13 on [Smith81a, pg. 410].

Theorem

If $G(t)$ is DHR then $F(t) = \alpha * G(t)$ is SDHR where $0 < \alpha < 1$.

Proof

$G(t)$ is DHR is equivalent to:

$$\frac{g(t)}{1-G(t)} \geq \frac{g(t+x)}{1-G(t+x)} \textit{ for all } t \geq 0 \text{ and } x > 0 \tag{4.1}$$

Hence

$$\frac{g(t)}{g(t+x)} > \frac{1-G(t)}{1-G(t+x)} \textit{ for all } t \geq 0, x > 0 \tag{4.2}$$

Since $G(t)$ is strictly monotonically increasing in t and always less than 1 we have $\frac{g(t)}{g(t+x)} > 1$ *for all* $t \geq 0, x > 0$. Thus $g(t)$ is strictly monotonically decreasing. Observe that

$$g(t)G(t+x) - g(t+x)G(t) > 0 \tag{4.3}$$

since $G(t)$ is strictly monotonically increasing and $g(t)$ is strictly monotoncally decreasing. Since $g(t) > 0$, $1 - G(t) > 0$ *for all* $t \geq 0$ inequality (4.1) is equivalent to

$$g(t)[1-G(t+x)] \geq g(t+x)[1-G(t)] \tag{4.4}$$

$\Longleftrightarrow$

$$g(t)-g(t+x) \geq g(t)G(t+x) - g(t+x)G(t) \tag{4.5}$$

Thus to show that $F(t)$ is SDHR we must show:

$$f(t)-f(t+x) > f(t)F(t+x) - f(t+x)F(t) \tag{4.6}$$

Substituting $F(t) = \alpha * G(t)$ and $f(t) = \alpha * g(t)$ we have

$$\alpha * \left[g(t)-g(t+x)\right] > \alpha^2 * \left[g(t)G(t+x) - g(t+x)G(t)\right] \tag{4.7}$$

$\Longleftrightarrow$

$$g(t)-g(t+x) > \alpha * \left[g(t)G(t+x) - g(t+x)G(t)\right] \tag{4.8}$$

but this follows from (4.3), (4.5), and $0 < \alpha < 1$.

5. Comments

The HOPT policy is variable space and myopic, i.e., the decision whether or not to keep a file in the cache is based solely on the current scaled hazard rate and does not (explicitly) consider future reference behavior. This suggests an obvious fixed space policy which ranks all files by their scaled hazard rates whenever a replacement decision must be made. The files with the largest scaled hazard rates are retained. We call this policy **HMAX**.

If all the files are completely homogeneous, i.e., have identical inter-reference time distributions, equal sizes, and equal tape access costs, then HMAX simply ranks files on their real hazard rates. Since the hazard rates are assumed here to be identical strictly monotonic decreasing functions of the time since last reference, HMAX reduces to LRU. Similarly HOPT reduces to the Working Set (WS) Policy when the files are completely homogeneous (i.e., the two policies differ only in their parameterization, WS(τ)=HOPT($h(\tau)$)). The successful experience with WS and LRU suggests that HOPT and HMAX may prove practical.

6. Acknowledgements

My thinking about automatic file migration policies has been greatly influenced by my former thesis advisor A.J. Smith. The idea for the paper was prompted by Smith's remarks [Smith81b, pg. 525] on the connection between the Working Set policy and decreasing hazard rates for page inter-reference times. The author is also grateful for the comments of Domenico Ferrari, Ron Wolff, Nick Jewell, Raphael Alonso, Arie Shoshani, John McCarthy, Paula Hawthorn, and Joe Sventek.

7. Conclusions

We have shown that, if the automatic file migration policy STOCHOPT is applied to a file system in which all files have inter-reference time distributions which have strictly monotonic decreasing hazard rates, then STOCHOPT reduces to keeping all files while their scaled hazard rate (hazard rate times tape access cost divided by file size) exceeds the storage rental rate. This policy we call HOPT. Scaled (improper) decreasing hazard rate distributions remain DHR; hence HOPT remains optimal. A fixed space analog of HOPT, HMAX, which ranks files by their scaled hazard rates, exists. HOPT and HMAX reduce to WS and LRU respectively for completely homogeneous file (i.e., paging) systems.

Bibliography

[Barlow75]

Barlow, Richard E. and Frank Proschan, *Statistical Theory Reliablity and Life Testing* , Holt, Rhinehart and Winston New York , 1975

[Jewell82]

Jewell, Nicholas P., Mixtures of Exponential Distributions , *Annals of Statistics* vol. 10, no. 2 , 1982 , pp. 479-484

[Smith81a]

Smith, A.J., "Analysis of Long Term File Reference Patterns for Application to File Migration Algorithms" , *IEEE Transactions on Software Engineering* vol. SE-7, no. 4 , July 1981 , pp. 403-417

[Smith81b]

Smith, A.J., "Long Term File Migration: Development and Evaluation of Algorithms" , *Communications of the ACM*, vol. 24, no. 8 , August 1981 , pp. 521-532

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