Performance of the Move-To-Front Algorithm with Markov-Modulated Request Sequences

E. G. Coffman, Jr. Predrag Jelenković
Bell Laboratories
Lucent Technologies
600 Mountain Avenue
Murray Hill, NJ 07974
{egc, predrag}@bell-labs.com

Aug 15, 1998

Abstract

We study the classical move-to-front (MTF) algorithm for self-organizing lists within the Markov-modulated request (MMR) model. Such models are useful when list accesses are generated within a relatively small set of modes, with the request sequences in each mode being i.i.d.. These modes are often called localities of reference and are known to exist in such applications as traffic streams of Ethernet or ATM networks and the locus of control or data accesses of executing computer programs. Our main results are explicit formulas for the mean and variance of the search-cost, the number of accesses required to find a given list element.

By adjusting the number of modes, one can use the MMR methodology to trade off the quality of an approximation with the computational effort it requires. Thus, our results provide a useful new tool for evaluating the MTF rule in linear-search applications with correlated request sequences. We illustrate the computations with several examples.

Keywords: Self-organizing lists, move-to-front algorithm, Markov-modulated process, hidden Markov chains, locality of reference, internet modeling.

1 Introduction

Performance analysis of self-organizing data structures, e.g., lists and trees, has a long history; references to the early work in this area can be found in [8, 9]. Quite recently, interest in these data structures has been rekindled by cache design problems in modern distributed networks. One of the most popular heuristics for self-organizing lists has been the *move-to-front* (MTF) rule. MTF is defined on sequences of requests for elements in a given list of N elements. In processing each new request, MTF moves the requested element to the first position (i.e., the left-most position or head) of the list, if it is not already there; the ordering of the remaining $N-1$ elements of the list remains unchanged.

In the analysis of self-organizing lists, there have been two approaches: probabilistic and combinatorial (amortized) analysis. We are interested in typical or average-case behavior, so we concentrate on the former approach. For the latter, the reader is referred to Bentley and McGeoch (1985)[1]. Our specific objective is a probabilistic analysis of the most common cost function: the position (search distance) of the currently requested element. The novelty of our contribution lies in our extension of MTF theory to Markov-modulated request (MMR) sequences.

There is a large literature on MMR models that spans a wide variety of applications in applied probability and engineering, including speech recognition, communications engineering, statistics, and risk theory. See [6] for many references. Equivalent names under which Markov-modulated models are known in the literature are: Markov-modulated random walks, random walks on Markov chains, functions (or random functions) of Markov chains, and hidden Markov chains. We define them formally as follows. Let $J = \{J_t, t = 0, 1, \dots\}$ be a discrete time, aperiodic, irreducible, finite Markov chain with state space $\{1, \dots, K\}$, and transition matrix $P = \{p_{ij}\}$, and let $R = \{R_t, t = 0, 1, \dots\}$ be a discrete-time process with state space $\{1, \dots, N\}$. We say that R is Markov-modulated if the pair process $\{(R_t, J_t), t \geq 0\}$ is Markovian in its second coordinate, i.e., if

$$\mathbf{P}[R_t = r, J_t = j | (R_{t-1}, J_{t-1} = i), \dots, (R_0, J_0)] = \\ \mathbf{P}[R_t = r, J_t = j | J_{t-1} = i] = p_{ij} q_{jr},$$

where $q_{jr} := \mathbf{P}[R_t = r | J_t = j]$. Define the stationary marginal distribution of R as $q_r := \mathbf{P}[R_t = r] = \sum_{i=1}^{K} \pi_i q_{ir}$, where $\{\pi_i, 1 \leq i \leq K\}$, is the stationary distribution of J. Since the process J, to be called the *modulating process*, is ergodic, its stationary distribution is unique and positive. The process R models sequences of requests for elements of the set $L = \{1, \dots, N\}$. A discrete-time process $\{\sigma_t, t = 0, 1, \dots\}$ is induced by R and the MTF rule, where $\sigma_t = (\sigma_t(1), \dots, \sigma_t(N))$ is a list (permutation) of the elements of L. We assume that, at $t = 0$, all permutations are equally likely, i.e., σ_0 is uniformly distributed over the set of all $N!$ permutations. According to MTF, σ_{t+1} is constructed from σ_t by bringing the element R_t to the first position of the list, if it is not already there, and keeping the ordering of the remaining elements unchanged; thus, the positions of the elements that were ahead of R_t are increased by one, while the positions of those behind R_t remain unchanged. Note that the joint process $\{(J_{t-1}, \sigma_t), t \geq 0\}$ is a Markov chain.

Early work on the probabilistic analysis of MTF dealt with i.i.d. requests; see [7] for key references. Our MTF model lies within the broader framework of Markovian request sequences. The MTF scheme with time-dependent Markovian requests was investigated in Lam at al. (1984) [10], where a formula for the expected search cost was derived. Special cases of the Markov model that are analytically more tractable were investigated by Rodrigues (1995) [13], who also examined convergence to stationarity [12]. Dobrow and Fill (1996) [5] derived transient and stationary probabilities for MTF in the Markov model. They also investigated spectral properties and the rate of convergence to stationarity.

Our interest in the MMR model stemmed from its flexibility as a tractable model of the high correlation (locality) structure encountered in the traffic streams (e.g., voice, video, and multimedia) of modern communication networks. (References to these models can be found in [7].) But more generally, our methods extend the computational tools for evaluating the performance of linear-search heuristics in an environment of Markovian request sources. They will be significantly more efficient than existing techniques when MMR models apply with the number K of modes relatively small.

The analysis of the MTF algorithm appears in Section 2, where we develop explicit formulas for the transient and steady state mean and variance of the cost function. In Section 3, we show that the sequence of list updates is Markovian on an appropriate extension of the set of all permutations of $\{1,\ldots,N\}$. In the same section, we investigate the spectral analysis and the stationary distribution of this Markov chain. Numerical examples are presented in Section 5 to illustrate the results.

2 Mean and Variance of the Search Cost

We derive explicit formulas for the mean and variance of the search cost C_t: the current position in the list state σ_t of the element R_t requested at time t. Thus, the search cost at time t is a measure of the time required by a linear search of the list to find the element requested at time t. The reversed Markov chain $\{\tilde{J}_t\}$ and the corresponding modulation process $\{\tilde{R}_t\}$ arise naturally in the following analysis. The transition probability matrix of $\tilde{J}$ is denoted by $\tilde{P}=\{\tilde{p}_{ij}\}$, where $\tilde{p}_{ij}=\pi_j p_{ji}/\pi_i$. For each $i\neq j$, let A_t^{ji} be the event that j is to the left of i in σ_t. Given that $R_t=r$, the cost C_t is the number of elements to the left of r in the list plus 1 to account for r itself. Thus, if we let $\sum^*{}_{r_1,\ldots,r_k}$ denote the sum over all sequences of *distinct* elements $r_1,\ldots,r_k$ with $r_i\in\{1,\ldots,N\}$, $1\leq i\leq k$, then we have the well-known indicator-function representation,

$$\begin{aligned} C_t &= 1+\sum_{k,r}{}^{*}\,\mathbf{1}(R_t=r,A_t^{kr}) \\ &= 1+\sum_{1\leq i\leq K}\,\sum_{k,r}{}^{*}\,\mathbf{1}(J_t=i,R_t=r,A_t^{kr}), \end{aligned} \tag{1}$$

and hence, after taking expectations,

$$\mathbf{E}C_t = 1+\sum_{1\leq i\leq K}\,\sum_{k,r}{}^{*}\,\mathbf{P}[J_t=i,R_t=r,A_t^{kr}]. \tag{2}$$

The ith summand $\mathbf{P}[J_t=i,R_t=r,A_t^{kr}]$ in the above expression can be written

$$\begin{aligned} \mathbf{P}[J_t=i,R_t=r,A_t^{kr}] &= \sum_{m=1}^{t}\mathbf{P}[J_t=i,R_t=r,R_{t-1}\notin\{r,k\},\ldots,R_{t-m+1}\notin\{r,k\},R_{t-m}=k] \\ &\quad +\mathbf{P}[J_t=i,R_t=r,R_{t-1}\notin\{r,k\},\ldots,R_0\notin\{r,k\},A_0^{kr}] \end{aligned}$$

Recall that σ_0 is a permutation chosen uniformly at random, so we have $\mathbf{P}[A_0^{kr}]=1/2$ for all k,r. In terms of the reversed processes, we can write

$$\begin{aligned} \mathbf{P}[J_t=i,R_t=r,A_t^{kr}] &= \sum_{m=1}^{t}\mathbf{P}[\tilde{J}_0=i,\tilde{R}_0=r,\tilde{R}_1\notin\{r,k\},\ldots,\tilde{R}_{m-1}\notin\{r,k\},\tilde{R}_m=k] \\ &\quad +\frac{1}{2}\mathbf{P}[\tilde{J}_0=i,\tilde{R}_0=r,\tilde{R}_1\notin\{r,k\},\ldots,\tilde{R}_t\notin\{r,k\}]. \end{aligned} \tag{3}$$

For a more compact notation, we introduce the matrices

$$Q_{r_1,\dots,r_k} = \{\tilde{p}_{ij}(q_{jr_1} + \cdots + q_{jr_k})\} \quad \hat{Q}_{r_1,\dots,r_k} = \{\tilde{p}_{ij}(1 - q_{jr_1} - \cdots - q_{jr_k})\}. \tag{4}$$

Then the mth summand in (3) can be expressed as

$$\begin{aligned}
&\mathbf{P}[\tilde{J}_0 = i, \tilde{R}_0 = r, \tilde{J}_1 = i_1, \tilde{R}_1 \notin \{r,k\}, \dots, \tilde{J}_{m-1} = i_{m-1}, \tilde{R}_{m-1} \notin \{r,k\}, \tilde{R}_m = k] \\
&= \sum_{i_1,\dots,i_m} \pi_i q_{ir} \tilde{p}_{ii_1}(1 - q_{i_1 r} - q_{i_1 k}) \cdots \tilde{p}_{i_{m-2} i_{m-1}}(1 - q_{i_{m-1} r} - q_{i_{m-1} k}) \tilde{p}_{i_{m-1} i_m} q_{i_m k} \\
&= \pi_i q_{ir} (\hat{Q}_{rk}^{m-1} Q_k \mathbf{e})(i),
\end{aligned} \tag{5}$$

where $\mathbf{e}$ is a column vector of ones, and $(\hat{Q}_{rk}^{m-1} Q_k \mathbf{e})(i)$ is the ith element of the vector $\hat{Q}_{rk}^{m-1} Q_k \mathbf{e}$. Similarly, we find that

$$\mathbf{P}[\tilde{J}_0 = i, \tilde{R}_0 = r, \tilde{R}_1 \notin \{r,k\}, \dots, \tilde{R}_t \notin \{r,k\}] \;=\; \pi_i q_{ir} (\hat{Q}_{rk}^t \mathbf{e})(i). \tag{6}$$

Finally, after substituting (3), (5), and (6) into (2), we arrive at the following result. Let $\rho(Q)$ denote the spectral radius of matrix Q.

Theorem 1 *If $\rho(\hat{Q}_{rk}) < 1$ for all r, k, then*

$$\mathbf{E}C_t = 1 + \sum_{k,r}^{*} \left[\mathbf{v}_r (I - \hat{Q}_{rk})^{-1} Q_k e + \mathbf{v}_r \hat{Q}_{rk}^t \left(-(I - \hat{Q}_{rk})^{-1} Q_k \mathbf{e} + \frac{1}{2}\mathbf{e} \right) \right],$$

where $\mathbf{v}_r = (\pi_1 q_{1r}, \dots \pi_K q_{Kr})$. Furthermore, the stationary expected cost is given by

$$\mu := \lim_{t\to\infty} \mathbf{E}C_t = 1 + \sum_{k,r}^{*} \mathbf{v}_r (I - \hat{Q}_{rk})^{-1} Q_k \mathbf{e}. \tag{7}$$

Computing the stationary search cost by (7) involves $N(N-1)/2$ inversions of $K \times K$ matrices $(I - \hat{Q}_{rk})$. When K is not very large this is fairly easy to do. By contrast, in the general Markov model, $N(N-1)/2$ inversions of $N \times N$ matrices are required, where N is normally large.

Theorem 1 gives classical results as special cases. In particular, Markov-modulated requests become i.i.d. if we reduce the matrix P to the scalar 1 and put $q_{1r} = q_r := \mathbf{P}[R_t = r]$. Then $\hat{Q}_{rk} = 1 - q_r - q_k$, $Q_k = q_k$, and $\mathbf{v}_r = q_r$ are also scalars which when substituted into Theorem 1 give [2, 11]

$$\mathbf{E}C_t = 1 + 2\sum_{r<k} \frac{q_r q_k}{q_r + q_k} + \sum_{r<k} \frac{(q_r - q_k)^2 (1 - q_r - q_k)^t}{2(q_r + q_k)}, \tag{8}$$

with the stationary mean

$$\mu = 1 + 2\sum_{r<k} \frac{q_r q_k}{q_r + q_k}. \tag{9}$$

If requests are generated by an aperiodic, irreducible, finite-state Markov chain, R, then the stationary expected search cost is given by [10]

$$\mu = 1 + \sum_{r,k}^{*} \frac{1}{m_{rk} + m_{kr}}, \tag{10}$$

where m_{ij} is the expected first passage time from state i to state j in R. To verify this, put $q_{ir} = \mathbf{1}(i = r)$ and $J = R$, and hence $K = N$. We have $\mathbf{v}_r = (0, \ldots, \pi_r, \ldots, 0)$ and $\hat{Q}_{rk} = \tilde{P}_{rk}$, where $\tilde{P}_{rk}$ is obtained from $\tilde{P}$ by deleting its rth and kth column. Similarly, $Q_k\mathbf{e}$ is a column vector with elements $\tilde{p}_{1k}, \ldots, \tilde{p}_{Nk}$, so for each $r \neq k$, we have $\mathbf{v}_r(I - \hat{Q}_{rk})^{-1}Q_k\mathbf{e}$ as the probability of starting in state r and reaching state k before again visiting state r in the Markov chain $\tilde{P}$. Then, by Lemmas 3.1.1. and 3.1.2. of [10], it follows that $\mathbf{v}_r(I - \hat{Q}_{rk})^{-1}Q_k\mathbf{e} = 1/(m_{rk} + m_{kr})$, which when substituted into Theorem 1 gives (10).

To compute the second moment of C_t, we square (1) and obtain

$$\begin{aligned} C_t^2 &= 1 + 3\sum_{i=1}^{K}\sum_{k,r}{}^{*}\, \mathbf{1}(J_t = i, R_t = r, A_t^{kr}) \\ &\quad + \sum_{i=1}^{K}\sum_{k_1,k_2,r}{}^{*}\, 1(J_t = i, R_t = r, A_t^{k_1 r} A_t^{k_2 r}). \end{aligned} \tag{11}$$

Let $A_t^{k_1 k_2 r}$ be the event that the relative (left-to-right) ordering of k_1, k_2, and r in σ_t is $k_1 k_2 r$. We observe that $A_t^{k_1 r} A_t^{k_2 r}$ is the union of the two disjoint events $A_t^{k_1 k_2 r}$ and $A_t^{k_2 k_1 r}$, and then take the expected value of (11) to obtain

$$\mathbf{E}C_t^2 = 3\mathbf{E}C_t - 2 + \sum_{i=1}^{K}\sum_{k_1,k_2,r}{}^{*}\, \mathbf{P}[J_t = i, R_t = r, A_t^{k_1 k_2 r}] + \mathbf{P}[J_t = i, R_t = r, A_t^{k_2 k_1 r}]. \tag{12}$$

The probabilities $\mathbf{P}[J_t = i, R_t = r, A_t^{k_1 k_2 r}]$ in the expression above are computed by

$$\begin{aligned} &\mathbf{P}[J_t = i, R_t = r, A_t^{k_1 k_2 r}] = \\ &\sum_{m_1=1}^{t-1}\sum_{m_2=m_1+1}^{t} \mathbf{P}[J_t = i, R_t = r, R_{t-1} \notin \{k_1, k_2, r\}, \ldots, R_{t-m_1+1} \notin \{k_1, k_2, r\}, R_{t-m_1} = k_1, \\ &\qquad R_{t-m_1-1} \notin \{k_2, r\}, \ldots, R_{t-m_1-m_2+1} \notin \{k_2, r\}, R_{t-m_1-m_2} = k_2] \\ &\quad + \sum_{m_1=1}^{t} \mathbf{P}[J_t = i, R_t = r, R_{t-1} \notin \{k_1, k_2, r\}, \ldots, R_{t-m_1+1} \notin \{k_1, k_2, r\}, R_{t-m_1} = k_1, \\ &\qquad R_{t-m_1-1} \notin \{k_2, r\}, \ldots, R_0 \notin \{k_2, r\}, A_0^{k_2,r}] \\ &\quad + \mathbf{P}[J_t = i, R_t = r, R_i \notin \{k_1, k_2, r\}, 0 \le i \le t-1, A_0^{k_1 k_2 r}], \end{aligned} \tag{13}$$

so we have

Theorem 2 *If* $\max_{k,r:k\neq r} \rho(\hat{Q}_{rk}) < 1$, *then*

$$\begin{aligned} \lim_{t\to\infty} \mathbf{Var}(C_t) &= \mu - 1 - (\mu - 1)^2 \\ &\quad + \sum_{k_1,k_2,r}{}^{*}\, \mathbf{v}_r(I - \hat{Q}_{rk_1k_2})^{-1}\left(Q_{k_1}(I - \hat{Q}_{rk_2})^{-1}Q_{k_2} + Q_{k_2}(I - \hat{Q}_{rk_1})^{-1}Q_{k_1}\right)\mathbf{e}. \end{aligned} \tag{14}$$

Proof: From (13), a time-reversal argument, and the fact that $\max_{\{r,k:r\neq k\}} \rho(\hat{Q}_{rk}) < 1$, it follows that

$$\begin{aligned} \lim_{t\to\infty} \mathbf{P}[J_t = i, R_t = r, A_t^{k_1 k_2 r}] &= \sum_{m_1=0}^{\infty}\sum_{m_2=0}^{\infty} \mathbf{v}_r \hat{Q}_{rk_1k_2}^{m_1} Q_{k_1} \hat{Q}_{rk_2}^{m_1} Q_{k_2}\mathbf{e} \\ &= \mathbf{v}_r(I - \hat{Q}_{rk_1k_2})^{-1}Q_{k_1}(I - \hat{Q}_{rk_2})^{-1}Q_{k_2}\mathbf{e}. \end{aligned}$$

Substituting the expression above into (12) and then subtracting the square of the mean (μ^2) gives the result of the theorem. ∎

If requests are i.i.d. with $\mathbf{P}[R_t = r] = q_r$, then by specializing quantities as before, we get (recall that $q_r = \mathbf{P}[R_t = r]$)

$$\lim_{t\to\infty} \mathbf{Var}(C_t) = \mu - 1 - (\mu-1)^2 + \sum_{k_1,k_2,r}^{*} \frac{q_r q_{k_1} q_{k_2} [(q_r + q_{k_1}) + (q_r + q_{k_2})]}{(q_r + q_{k_1})(q_r + q_{k_2})(q_r + q_{k_1} + q_{k_2})} \tag{15}$$

For the general Markov model, the expression for the variance of the search cost was first derived in Rodrigues (1995), Theorem 4.1, [13]. That result can also be derived directly from the Theorem 2, but we refrain from doing so, as the computation is quite awkward.

For a useful special case, assume that the list is partitioned into K disjoint subsets L_i, $1 \le i \le K$, $L_i = \{(i,1), \ldots, (i, N_i)\}$, $\sum_{i=1}^{K} N_i = N$, and that, when the underlying Markov chain is in mode i ($J_t = i$), the request sequence can only access items from the subset L_i, i.e., $q_{i,(i,r)} > 0$ for $1 \le r \le N_i$ and $q_{j,(i,r)} = 0$ for $j \ne i$. With a small abuse of notation we write $q_{(i,r)} = q_{i,(i,r)}$.

Theorem 3 *For the stationary expected cost, we have*

$$\begin{aligned} \mu \;=\; & 1 + \sum_i \pi_i \sum_{1\le k,l\le N_i}^{*} \frac{q_{(i,k)} q_{(i,l)}}{q_{(i,k)} + q_{(i,l)}} \\ & + \sum_{i,j}^{*} \sum_{1\le k\le N_i, 1\le l\le N_j} \frac{\pi_i \pi_j q_{(i,k)} q_{(j,l)}}{\pi_i q_{(i,k)} + \pi_j q_{(j,l)} - \pi_i \pi_j q_{(i,k)} q_{(j,l)} (1/\pi_i + 1/\pi_j - m_{ij} - m_{ji})}, \end{aligned} \tag{16}$$

where m_{ij} is the expected first passage time from state i to state j in J.

Remarks: Note that when the underlying Markov chain J represents an i.i.d. sequence (i.e., each row in P is equal to the distribution π), then $1/\pi_i + 1/\pi_j - m_{ij} - m_{ji} = 0$ and (16) reduces to (9). At the other extreme, if each subset reduces to $L_i = \{(i,1)\}$ and hence $J = R$, then (16) yields the result for Markovian requests stated in (10).

For the proof of Theorem 3, we first assemble a couple of well-known results for finite Markov chains. Let m_{ij} be the expected first passage time from state i to state j in J, and let ${}_j p_{ii}^{(n)}$ be the probability of going from state i to state i in n steps without visiting j (these are called taboo probabilities [4, p. 45]); let $\tilde{m}_{ij}$, ${}_j \tilde{p}_{ii}^{(n)}$ represent the analogous quantities for the reversed chain $\tilde{J}$. We denote by $\tilde{P}^{*i}$ the matrix obtained from $\tilde{P}$ by deleting its ith column; similarly, $\tilde{P}^{*ij}$ denotes the matrix resulting from the deletions of columns i and j.

Lemma 1 *Let $i \ne j$. Then*

(i) $\tilde{m}_{ii}/(\tilde{m}_{ij} + \tilde{m}_{ji}) = 1/\left(1 + \sum_{n=1}^{\infty} {}_j\tilde{p}_{ii}^{(n)}\right) = \det(I - \tilde{P}^{*i})/\det(I - \tilde{P}^{*ij})$

(ii) $m_{jj} = \tilde{m}_{jj} = 1/\pi_j$

(iii) $m_{ij} + m_{ji} = \tilde{m}_{ij} + \tilde{m}_{ji}$

Proof: The first equality in (i) follows from [4, eq. (14), pp. 49] and [4, Corollary 1, pp. 65]. The second equality is just an algebraic identity. Statements (ii) and (iii) represent Lemma 3.1.2 in [10]. ∎

Proof of Theorem 3: We will show that the expression in (16) follows directly from (7). First, note that $\mathbf{v}_{(i,k)} = (0,\dots,0,\pi_i q_{(i,k)},0,\dots,0)$, $(i,k) \in L_i$ and that $Q_{(i,k)}$ has only 1 nonzero column; it is the ith column of $\tilde{P}$ multiplied by $q_{(i,k)}$. Also, for $k \neq l$, we have $Q_{(i,k)(i,l)} = Q_{(i,k)} + Q_{(i,l)}$ and $\hat{Q}_{(i,k)(i,l)} = P - Q_{(i,k)(i,l)}$. These observations and simple algebra yield

$$\mathbf{v}_{(i,k)}(I - \hat{Q}_{(i,k)(i,l)})^{-1} Q_{(i,l)}\mathbf{e} = \pi_i \frac{q_{(i,k)}q_{(i,l)}}{q_{(i,k)} + q_{(i,l)}};$$

this justifies the first sum in (16).

When a pair of items (i,k) and (j,l) belong to different subsets ($i \neq j$), then $Q_{(i,k)(j,l)}$ has two nonzero columns which are the ith and jth columns in P multiplied by $q_{(i,k)}$ and $q_{(j,l)}$, respectively; also, we have $\hat{Q}_{(i,k)(j,l)} = P - Q_{(i,k)(j,l)}$. Then

$$\begin{aligned} \mathbf{v}_{(i,k)}(I - \hat{Q}_{(i,k)(j,l)})^{-1} Q_{(j,l)}\mathbf{e} &= \frac{\pi_i q_{(i,k)} q_{(j,l)}}{\det(I - \hat{Q}_{(i,k)(j,l)})} \sum_{j_1=1}^{K} (-1)^{i+j_1} \tilde{p}_{j_1 j} \det((I - \hat{Q}_{(i,k)(j,l)})^{*i}_{*j_1}) \\ &= \pi_i q_{(i,k)} q_{(j,l)} \frac{\det(I - \tilde{P}^{*j})}{\det(I - \hat{Q}_{(i,k)(j,l)})}, \end{aligned} \tag{17}$$

where $(I - \hat{Q}_{(i,k)(j,l)})^{*i}_{*j_1}$ is the matrix obtained from $(I - \hat{Q}_{(i,k)(j,l)})$ by deleting its j_1th row and ith column. The first equality in (17) just exploits the determinant representation of $(I - \hat{Q}_{(i,k)(j,l)})^{-1}$; the second equality follows from elementary properties of determinants, which also give us

$$\begin{aligned} \det(I - \hat{Q}_{(i,k)(j,l)}) &= q_{(i,k)} \det(I - \tilde{P}^{*i}) + q_{(j,l)} \det(I - \tilde{P}^{*j}) \\ &\quad - q_{(i,k)} q_{(j,l)} (\det(I - \tilde{P}^{*i}) + \det(I - \tilde{P}^{*j}) - \det(I - \tilde{P}^{*ij})) \end{aligned} \tag{18}$$

Next, let D_P be a determinant obtained by replacing one column in $(I - \tilde{P})$ with $\mathbf{e}$ (the value of D_P is independent of the column replaced). It is easy to show that

$$\pi_i = \frac{\det(I - \tilde{P}_i)}{D_P}. \tag{19}$$

Thus, after dividing (18) by D_P, and applying Lemma 1 and (19), we arrive at

$$\begin{aligned} \frac{\det(I - \hat{Q}_{(i,k)(j,l)})}{D_P} &= q_{(i,k)}\pi_i + q_{(j,l)}\pi_j \\ &\quad - q_{(i,k)} q_{(j,l)} (\pi_i + \pi_j - \pi_i \pi_j (m_{ij} + m_{ji})) \end{aligned} \tag{20}$$

Finally, divide the numerator and denominator in (17) by D_P and substitute (19) and (20) to obtain

$$\mathbf{v}_{(i,k)}(I - \hat{Q}_{(i,k)(j,l)})^{-1} Q_{(j,l)}\mathbf{e} = \frac{\pi_i \pi_j q_{(i,k)} q_{(j,l)}}{\pi_i q_{(i,k)} + \pi_j q_{(j,l)} - \pi_i \pi_j q_{(i,k)} q_{(j,l)} (1/\pi_i + 1/\pi_j - m_{ij} - m_{ji})}$$

which, when summed over i, j, k, l yields the second sum in (16). ∎

3 Examples and Discussion

As we have noted, in many applications, a state of the modulating process J determines that mode or context, of perhaps many, in which a list is being accessed. As a symmetric example for $K = 2$ possible accessing modes, consider a modulated Zipf's law. In mode 1, the request frequencies are $q_{1r} = 1/(jH_N),\ 1 \le j \le N$, with $H_N = \sum_{i=1}^{N} 1/i$. For maximum contrast in mode 2, we reverse the ordering of L by request frequency, i.e., we take the complementary probabilities $q_{2r} = 1/(N-j+1)\cdot 1/H_N,\ 1 \le j \le N$. The two-state Markov chain J is defined by the transition probabilities $p_{12} = p_{21} = (1-w)/2$, where w is a 'memory' parameter with $|w| \le 1$. As w decreases to -1, R increases its tendency to jump from one mode to the other, whereas it resides for long periods in a mode when w is close to 1. The stationary distribution of J is given by $\pi_1 = \pi_2 = 1/2$, and the unconditional request probabilities are $q_r = \pi_1 q_{1r} + \pi_2 q_{2r} = (q_{1r} + q_{2r})/2$.

Figure 1 plots the expected search cost (7) for $N = 50, 100$ as a function of w, and compares it to two i.i.d. models. The dashed line labeled 'MTF i.i.d.' refers to the MTF performance under independent requests each drawn from $\{q_r\}$. If the request frequencies are known in advance, and if requests are independent, then it is optimal to order L by decreasing request frequency and to keep this ordering fixed. This gives the 'optimal static i.i.d.' dashed line in the figure. As can be seen, for negative memory (w), MTF performs nearly as it would were requests i.i.d. (for $w = 0$ it would be exactly the same, as is easily verified). And for positive memory, the MTF expected search cost experiences a steep drop, especially as w nears 1. Indeed, for w sufficiently close to 1, MTF does even better than in the optimal static i.i.d. case. In this regime, MTF normally spends very long periods of time in making i.i.d. requests according to $\{q_{1r}\}$ or $\{q_{2r}\}$ before switching from one mode to the other; relatively little time is spent in making the major list restructuring that accompanies changes in mode. And as implied by the figure, when w is close to 1, MTF operating in either mode is better than the optimal static algorithm in the single 'combined' mode with independent requests drawn form $\{q_r\},\ q_r = (q_{1r} + q_{2r})/2$.

Next, consider a similar experiment with Markov modulated geometric request sequences. The probabilities are

$$q_{1r} = \lambda^{(r-1)}/c_N,\ q_{2r} = \lambda^{(N-r)}/c_N,\ 1 \le r \le N,$$

with the normalization constant $c_N = \sum_{r=0}^{N-1} \lambda^r$.

For $N = 100$, and $\lambda = 0.95$ the two distributions $Q1$ and $Q2$ are pictured on the left-hand side of Figure 2. On the right-hand side, we have displayed the dependence of the expected search cost on the memory parameter w. It is interesting to observe once again that for w close to one both 'MTF i.i.d.' and 'optimal static i.i.d.' are quite pessimistic in comparison with the values of the MTF algorithm in which locality of reference is modeled explicitly.

All of the observations made to this point apply also to our third and final example in which we assume that the list is partitioned into two equal-size blocks; Markov modulated Poisson request sequences are drawn alternately from the two blocks. Specifically, we choose $K = 2,\ N_1 = N_2 = N/2 = 20$, and take

$$q_{(1,k)} = q_{(2,k)} = \lambda^{(k-1)}/(c_N(k-1)!),\ 1 \le k \le N/2,$$

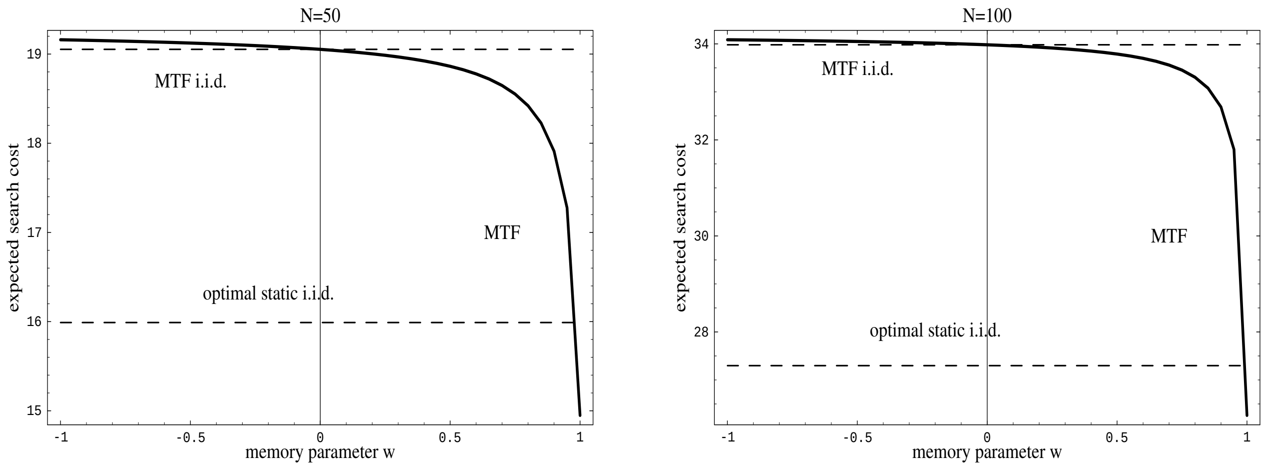

Figure 1: Illustration for the Markov modulated Zipf's law

with the normalization constant $c_N = \sum_{r=0}^{N/2-1} \lambda^k/k!$. Here, we use formula (16) with

$$m_{12} + m_{21} = \frac{1}{p_{12}} + \frac{1}{p_{21}} = \frac{4}{1-w}$$

to compute the expected search cost. For $\lambda = 10$, this computation is shown in Figure 3.

4 Final Remarks

The stationary distribution of search cost is of obvious interest, especially in studies of LRU caching where tail probabilities (fault rates) are needed. To obtain a formal solution to this problem, assume that both the MMR process and the search-cost process $\{C_t\}$ are stationary (initial states are samples from the stationary distributions). The derivation below follows closely that in [3], so we will be brief. For $k \geq 1$, let $X_k = \{\tilde{R}_1, \ldots, \tilde{R}_k\}$ be the set of distinct elements in the first k requests. Then

$$\begin{aligned} \mathbf{P}[R_0 = r, C_0 = n] &= \sum_{k=0}^{\infty} \mathbf{P}[\tilde{R}_0 = r, r \notin X_k, |X_k| = n-1, \tilde{R}_{k+1} = r] \\ &= \sum_{k=0}^{\infty} \sum_{\{B:|B|=n-1, r \notin B\}} \mathbf{P}[\tilde{R}_0 = r, X_k = B, \tilde{R}_{k+1} = r] \end{aligned} \tag{21}$$

By the inclusion-exclusion formula,

$$\begin{aligned} \mathbf{P}[\tilde{R}_0 = r, X_k = B, \tilde{R}_{k+1} = r] &= \sum_{\{A:A\subseteq B\}} (-1)^{|B-A|} \mathbf{P}[\tilde{R}_0 = r, X_k \subseteq A, \tilde{R}_{k+1} = r] \\ &= \sum_{\{A:A\subseteq B\}} (-1)^{|B-A|} \mathbf{v}_r Q^k(A) Q_r \mathbf{e}. \end{aligned} \tag{22}$$

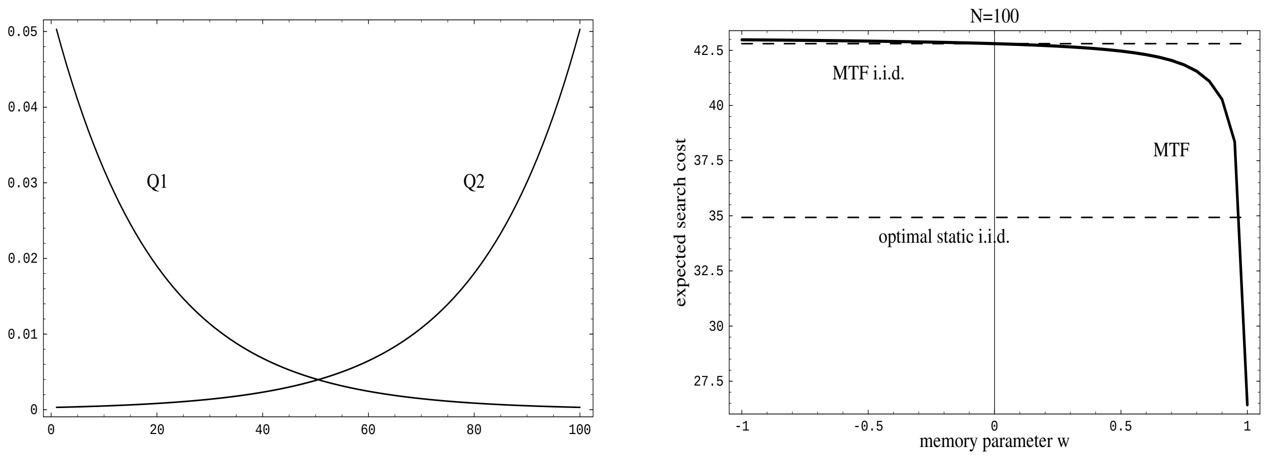

Figure 2: Illustration for Markov modulated geometric request sequences

Substituting (22) into (21) gives

$$\begin{aligned}\mathbf{P}[R_0 = r, C_0 = n] &= \sum_{\{B:|B|=n-1, r\notin B\}} \sum_{\{A:A\subseteq B\}} (-1)^{|B-A|}\mathbf{v}_r(I - Q(A))^{-1}Q_r\mathbf{e} \\ &= \sum_{\{A:|A|\le n-1, r\notin A\}} \sum_{\{B:B\supseteq A, |B|=n-1, r\notin B\}} (-1)^{|B-A|}\mathbf{v}_r(I - Q(A))^{-1}Q_r\mathbf{e}.\end{aligned}$$

The number of sets in $\{B : B \supseteq A, |B| = n-1, |A| = a, r \notin B\}$ is $\binom{N-1-a}{n-1-a}$, so after simplifying and summing over r, we obtain the following result, which reduces easily to the result in [3] when requests are i.i.d..

Theorem 4

$$\mathbf{P}[C_0 = n] = \sum_{r=1}^{N}\sum_{a=0}^{n-1}(-1)^{n-a-1}\binom{N-1-a}{n-1-a}\sum_{\{A:|A|=a, r\notin A\}} \mathbf{v}_r(I - Q(A))^{-1}Q_r\mathbf{e}.$$

Unfortunately, Theorem 4 does not give us a feasible computation for interesting values of N (e.g., for N in the thousands, at least). We leave as an interesting open question, the problem of estimating search-cost probabilities.

References

[1] J. L. Bentley and C. C. McGeoch. Amortized analyses of self-organizing sequential search heuristics. *Communications of the ACM*, 28(4):404–411, 1985.

[2] J. R. Bitner. Heuristics that monotonically organize data structures. *SIAM J. Comput.*, 8:82–110, 1979.

[3] P. J. Burville and J. F. C. Kingman. On a model for storage and search. *Journal of Applied Probability*, 10:697–701, 1973.

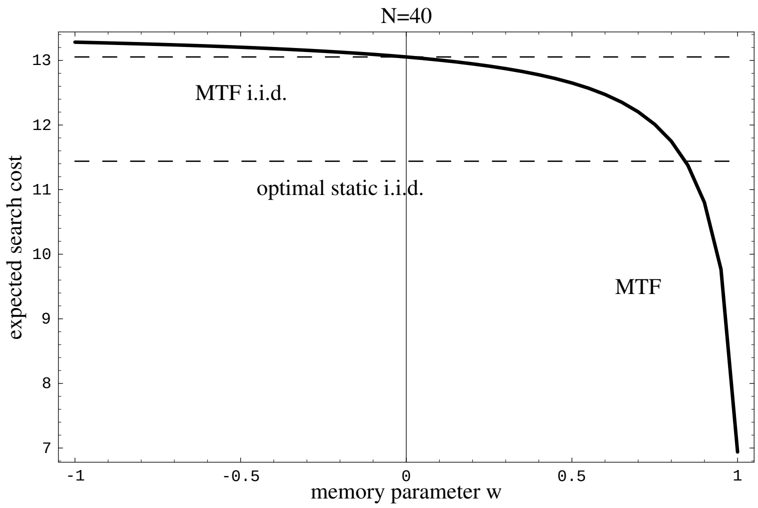

Figure 3: Illustration for Markov modulated Poisson request sequences

[4] Kai Lai Chung. *Markov Chains with Stationary Transition Probabilities.* Springer-Verlag, New York, 1960.

[5] R. P. Dobrow and J. A. Fill. The move-to-front rule for self-organizing lists with Markov dependent requests. In D. Aldous, P. Diaconis, J. Spencer, and J. M. Steele, editors, *Discrete Probability and Algorithms.* Springer-Verlag, 1996.

[6] B. M. Hochwald and P. R. Jelenković. State learning and mixing in hidden Markov models and the Gilbert-Elliot channel. *IEEE Transactions on Information Theory*, 1998, to appear.

[7] P. R. Jelenković. Asymptotic approximation of the move-to-front search cost distribution and least-recently-used caching fault probabilities. *Annals of Applied Probability*, 1998, to appear.

[8] E. G. Coffman Jr. and P. J. Denning. *Operating Systems Theory.* Prentice-Hall, 1973.

[9] D. E. Knuth. *The Art of Computer Programming, Vol. 3: Sorting and Searching.* Addison-Wesley, 1973.

[10] K. Lam, M. Y. Leung, and M. K. Siu. Self-organizing files with dependent accesses. *Journal of Applied Probability*, 21:343–359, 1984.

[11] J. McCabe. On serial files with relocatable records. *Operations Research*, 13:609–618, 1965.

[12] E. R. Rodrigues. Convergence to stationary state for a Markov move-to-front scheme. *Journal of Applied Probability*, 32(3):768–776, 1995.

[13] E. R. Rodrigues. The performance of the move-to-front scheme under some particular forms of Markov requests. *Journal of Applied Probability*, 32(4):1089–1102, 1995.