On the Performance of Sequential Paging for Mobile User Location

Bhaskar Krishnamachari, Rung-Hung Gau, Stephen B. Wicker, Zygmunt J. Haas
School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853
{bhaskar,runghung,wicker,haas}@ece.cornell.edu

Abstract

We present some results regarding the paging cost gains and the cost-delay tradeoff that can be achieved by using sequential paging to locate mobile users in cellular networks. We describe tight bounds on the average paging cost and the paging delay and quantify the intuition that greater gains are achieved when the mobile user's location probabilities are concentrated in a small portion of the location area. We also examine the impact of errors in the location estimates on the paging cost.

1. Introduction

When a call arrives for a mobile user in a cellular radio network, it is necessary to locate the current cell of that user in order to set up the call. Typically, the mobile informs the network of its location using *location updates*, and a set of cells where the user is known to be is then *paged* to determine the exact location when a call arrives.

A *sequential paging* scheme is one where the n cells in a service region are partitioned into w indexed groups referred to as *paging zones*. When a call arrives for a user, the cells in first paging zone are paged simultaneously in the first round and then, if the user is not found in the first round of paging, all the cells in the second paging zone are paged, and so on. Rose and Yates, in their seminal paper on the subject [1], show how to use the cell-wise location probabilities of a user to partition a location area of n cells into w paging zones in such a way that the average average cost of paging is minimized. Sequential paging permits a trade-off between the cost of paging and the paging delay. This can be seen by considering the two extreme cases: when $w = 1$, all n cells are paged simultaneously for every call, and the worst case delay is only 1 paging round; when $w = n$, each paging zone has only one cell, so that although the average paging cost will be reduced, the worst case delay is now increased to n paging rounds.

This paper studies a number of performance issues concerning optimal sequential paging. The rest of it is organized as follows. In section 2, we introduce some useful notation and definitions. A dynamic programming algorithm for minimizing average paging cost under a worst case delay constraint is specified in section 3 along with a description of its time and space complexity. In section 4, we present bounds on the delay incurred and paging cost gains achieved by sequential paging. We consider selective paging, a version of sequential paging particularly useful in conjunction with dynamic location update algorithms, in section 5. We show conditions under which this technique is optimal, and use it to quantify some of the paging cost - delay tradeoffs for a geometric zone location distribution. In section 6 we compare the paging cost and delay of optimal sequential paging scheme with simple, suboptimal, sequential paging schemes in which the k-most probable cells are paged first. We discuss the effect of estimation error in the user's location on the performance of sequential paging in section 7. We present our concluding comments in section 8.

2. Preliminary Notation and Definitions

Let $\Pi_n = \{\pi_1, \pi_2, \ldots, \pi_n\}$, where π_i is the estimated probability that the user is located in the i^{th} cell. We adopt the convention in this paper that the cells are numbered in decreasing order of location probabilities as paging in this order minimizes the paging costs [1]. Let the number of cells in the i^{th} paging zone be denoted by n_i, and let p_i be the corresponding *zone location probability* of the user. Let the paging zones be partitioned into w zones in such a way that the average paging cost is minimized. The corresponding *average paging cost*, i.e. the average number of cells paged per call arrival, is $\overline{L}_w^{\Pi_n} = \sum_{i=1}^{w}((\sum_{j=1}^{i} n_j)p_i)$; and the corresponding *average paging delay*, i.e. the average number of zones paged per call arrival, is $\overline{D}_w^{\Pi_n} = \sum_{i=1}^{w}(ip_i)$.

We can normalize the average paging cost with respect to the average paging cost when using only one partition. Let the *normalized reduction in average paging cost* be defined as follows:

$$\Lambda_{\Pi_n}(w) = \frac{\overline{L}_1^{\Pi_n} - \overline{L}_w^{\Pi_n}}{\overline{L}_1^{\Pi_n}} \tag{1}$$

We would like $\Lambda_{\Pi_n}(w)$, which represents the reduction in paging cost that is gained by using multiple paging zones, to be as close to 1 as possible. If $\Lambda_{\Pi_n}(w)$ is 0.5 for example, then we can get a reduction of 50% in the paging cost by using w paging zones instead of 1.

3. Algorithm for Optimal Sequential Paging

If we are given a location distribution Π_n, a dynamic programming algorithm based on the following recursion provides the partition into w zones that minimizes the average paging cost. Let $h[n', w']$ be the minimum average paging cost achievable by partitioning the first n' cells into w' paging zones. Then:

$$h[n', w'] = min_{j=w'-1}^{n'} \left(h[j, w'-1] + n' \sum_{i=j+1}^{n'} \pi_i \right) \tag{2}$$

We can think of $h[.,.]$ as an $n \times w$ array which is filled up using the above recursion, and the initial conditions $h[n', 1] = n', \forall n'$. The solution we desire is $h[n, w]$. The optimal partition can be determined by keeping track of the decisions at each step and tracing back through this array as is done in dynamic programming algorithms [2]. Due to the size of the array, the space complexity of the algorithm is $O(wn)$, and the time complexity is $O(wn^2)$ as at most n previous array entries must be compared for each new entry.

4. Bounds for Sequential Paging

It is shown in [1] that sequential paging schemes have the worst case average paging cost and paging delay when the mobile users are equally likely to be in any cell, i.e. a uniform distribution U_n. Under this distribution, the average is paging cost is minimized when each paging zone has the same number of cells (plus/minus one). The following inequalities hold in this case:

$$p_i \le \frac{1}{n}\lceil \frac{n}{w} \rceil \tag{3}$$

$$\sum_{j=1}^{i} n_j \le i \lceil \frac{n}{w} \rceil \tag{4}$$

It can be verified that the expressions (3) and (4), in conjunction with the definitions provided in section 2, yield the following tight bounds on the average paging delay, the average paging cost and the normalized gain in average paging cost:

$$\overline{D}_w^{\Pi_n} \le \overline{D}_w^{U_n} = \frac{w(w+1)}{2n} \lceil \frac{n}{w} \rceil \approx \frac{w+1}{2} \tag{5}$$

$$\overline{L}_w^{\Pi_n} \le \overline{L}_w^{U_n} = \frac{w(w+1)}{2n} \lceil \frac{n}{w} \rceil^2 \approx \frac{n}{2}(1 + \frac{1}{w}) \tag{6}$$

$$\Lambda_{\Pi_n}(w) \ge \Lambda_{U_n}(w) = 1 - \frac{w(w+1)}{2n^2} \lceil \frac{n}{w} \rceil^2 \approx \frac{1}{2}(1 - \frac{1}{w}) \tag{7}$$

The approximate expressions hold exactly when the number of cells is an integer multiple of the number of paging zones, and in the limit as $n \to \infty$. The result (7) is particularly appealing. It tells us, for example, that we can obtain at least a 25% reduction in paging costs by using 2 paging zones, and at least a 40% reduction in paging costs by using up to 5 paging zones *irrespective of what the user location probabilities are*.

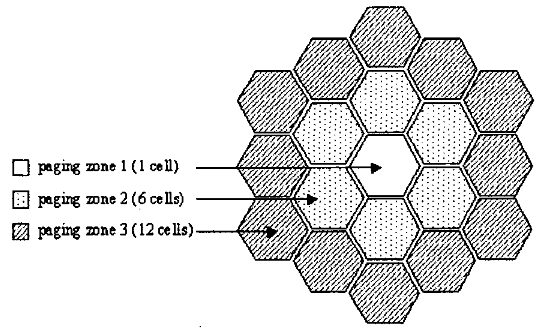

Figure 1. An example of selective paging with hexagonal cells, $w = 3$ paging zones

5. Selective Paging and Geometric zone-location probabilities

The cluster paging scheme introduced in [3], more commonly referred to as selective paging in the literature (e.g. [5]), is a special case of sequential paging. In selective paging cells are paged from the last known location outwards in successively concentric rings. An example with a maximum of 3 rounds of paging can be seen in figure 1. Paging

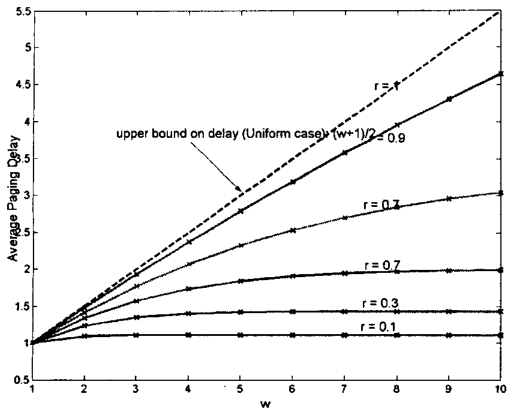

Figure 2. Average paging delay for selective paging with geometric location distribution

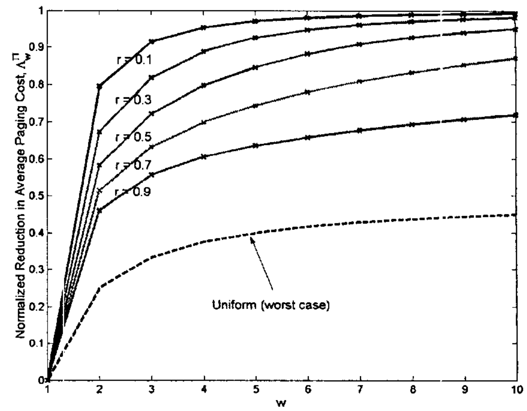

Figure 3. Normalized reduction in average paging cost for selective paging with geometric location distribution

zone 1 consists of the center cell, paging zone 2 consists of the 6 cells in the ring surrounding it, and paging zone 3 is the ring of 12 cells on the outside. Selective paging finds applications particularly in the context of dynamic location update schemes. The following theorem tells us when selective paging is optimal:

Theorem: *If the user location probability is equal for all cells in each ring, and each successive outer ring has a lower location probability, then selective paging minimizes the average cost of paging under the constraint that there be no more paging zones than the number of rings.*

Due to space limitations, we only provide a sketch of the proof here. Further details can be found in [4]. The proof proceeds by showing that every possible partition of the n cells into w paging zones can be constructed from the partition used in selective paging (placing all cells of the i^{th} ring in the i^{th} paging zone), by making a series of simple "moves" of cells from one partition to another. It is then shown that, under the assumption about the location probabilities, these moves can never result in an overall decrease in the paging cost. This in turn implies that the partition used in selective paging is indeed optimal under these assumptions.

Empirical data described in [3] suggests that in some real-world situations, the zone location probabilities are distributed in a geometric manner, i.e. for some $0 < r \leq 1$, $p_i = r^i/(\sum_{i=1}^{w} r^i)$. The number of cells in the i^{th} paging zone can be described by the relation $n_i = M \cdot i$ for $i \geq 2$ and $n_1 = 1$, where M is a positive integer constant that depends on the shape of the cells (for example, M = 6 for hexagonal cells, and 8 for rectangular cells). Under these assumptions, the following expressions for the average paging delay and average paging cost can be easily derived:

$$\overline{D}_w = \frac{1-r}{1-r^w} \cdot \sum_{i=1}^{w} i r^{i-1} \qquad (8)$$

$$\overline{L}_w = \frac{1-r}{1-r^w} \cdot \sum_{i=1}^{w} (1 + M\frac{i(i-1)}{2}) r^{i-1} \qquad (9)$$

The average paging delay $\overline{D}_w$ and the normalized reduction in average paging cost $\overline{\Lambda}(w)$ versus w are plotted in figures 2 and 3, based on the above equations[1]. The dashed lines represent the bounds provided by the uniform distribution discussed in section 4. These curves confirm the intuition that the greatest gains are obtained when the user location probabilities are concentrated in a relatively small portion of the location area, which is the case when r has a low value.

6. Suboptimal Sequential Paging

The algorithm present in section 3 provides the optimum partition of the location area into paging zones based on the particular user location distribution. One can envision simpler, suboptimal, sequential paging schemes which are independent of the location distribution. One example, when

[1] It should be kept in mind that the number of cells is not kept constant in figures 2 and 3, but rather grows linearly with the number of paging zones.

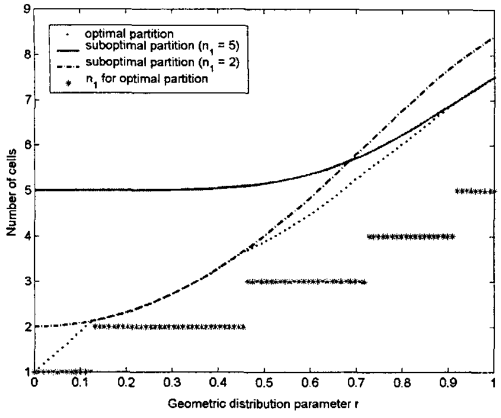

Figure 4. Average paging cost for optimal and suboptimal partitions (n=10, w=2)

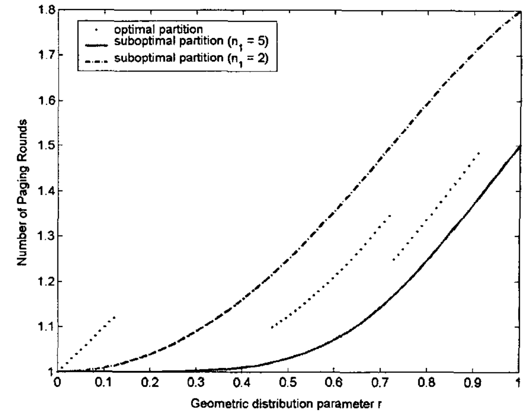

Figure 5. Average paging delay for optimal and suboptimal partitions (n=10, w=2)

$w = 2$, is to put the k most probable cells in the first paging zone (i.e. $n_1 = k$), and the rest in the second paging zone. Such a scheme may be appealing for practical applications because of ease of implementation. It is reasonable to inquire how such a suboptimal but distribution-independent scheme would compare with the optimal solution.

To examine this, we consider a sample scenario where the *cell* location probabilities have a geometric distribution[2], i.e. for some $0 < r \leq 1$, $\pi_i = r^i/(\sum_{i=1}^{n} r^i)$. The average paging cost and paging delay incurred when $n_1 = k$ is then:

$$\overline{L} = n - (n-k)\left(\frac{1-r^k}{1-r^n}\right) \tag{10}$$

$$\overline{D} = 2 - \left(\frac{r-r^k}{1-r^n}\right) \tag{11}$$

Figures 4 and 5 show the average paging cost and average paging delay with respect to the parameter r for two such suboptimal schemes when $n = 10$. In one case $n_1 = k = 2$, and in the other $n_1 = k = 5$. These are shown along with the corresponding plots for the optimum obtained using the dynamic programming algorithm. A number of observations can be made about these results. The first is that in the optimum solution, the number of cells in the first paging zone depends upon the parameter r, increasing monotonically as the distribution becomes less concentrated. Hence the performance of the k-cells-first schemes, though suboptimal in general, coincides with that of the optimal over certain ranges of r. In this particular example, the 2-cells-first scheme performs generally better than the 5-cells-first scheme in terms of average paging cost, being nearly optimal for the entire range of r values. In general a suitable value of k that is independent of the probability distribution would depend upon the total number of cells.

The optimal sequential paging scheme minimizes the average paging cost, but clearly at the expense of greater average paging delay. For example, from figure 5, it is clear that paging 5 cells first always has delay less than or equal to that of the optimal scheme. Another interesting observation that can be made regarding the average delay of the optimal sequential paging scheme is that it is discontinuous with respect to the parameter r. This is understandable, since the discontinuities occur when the number of cells in the first paging zone (and hence the probability of locating the user in one round) also undergo discontinuous jumps.

[2]Note the difference from section 5 where we assumed that the *zone* location probabilities were geometric. Geometric location distributions are generally useful for studying the performance of paging algorithms because they allow us to tune the "concentration" of the user's locations. When the parameter r is near 0, the user's movements are restricted to a small portion of the location area, when it is near 1, the user is equally likely to be anywhere in the location area.

7. Sensitivity to Estimation Error

The use of sequential paging schemes is predicated upon the ability to obtain good estimates of the cell-wise user location probabilities. We now briefly consider what happens when there are errors in the location estimates for the user. We perform some experiments again with the same parameters as in section 6: $n = 10, w = 2$, geometric cell location probabilities with parameter r. We introduce noise to the probability distribution using a noise level parameter η as follows. We start with the correct distribution Π_n, and to each element of this distribution, we add a random number uniformly from $[-\eta, +\eta]$ clamping the values at 0 to prevent negative numbers. The new distribution is then normalized so that it sums to 1. The noise level η is thus a direct indicator of the estimation error.

Figure 6 shows the average paging cost obtained when noisy estimates of the location probability are used, with respect to the geometric parameter r. The figure was obtained by running Monte Carlo simulations 30 times for each noise level and r value, with 1000 sample locations generated in each simulation according to the correct location distribution. When the error is very low ($\eta = 0.05$), there is no significant difference between the sequential paging using the correct distribution (optimum) and paging using the noisy distribution[3]. As is to be expected, the average paging costs increase with increasing η, but sequential paging appears to be fairly robust to errors in the location distribution, as for all noise levels studied the average paging cost is significantly lower than 10, which would be the case if all n cells were paged at once. It can also be seen from the figure that the noise affects the paging costs the least when the user is equally likely to be anywhere ($r = 1$). Another way to put this is that the performance of sequential paging is most sensitive to estimation errors when the correct distribution is concentrated in a small portion of the location area.

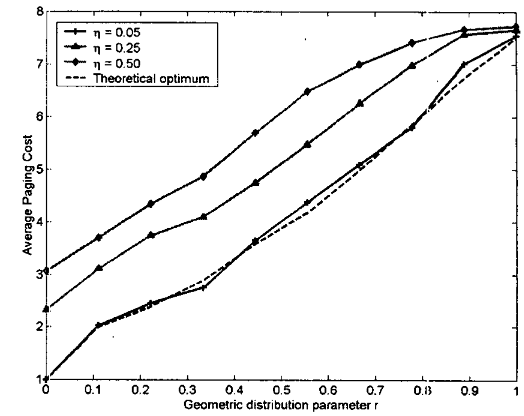

Figure 6. Effect of estimation error on paging cost (n=10, w=2)

8. Conclusions

In this paper we have presented some results concerning the performance of sequential paging for mobile location in terms of the average paging cost and paging delay. We gave tight bounds on the average paging delay, average paging cost and the normalized reduction in average paging costs that can be obtained using sequential paging. We described conditions under which selective paging is the optimal sequential paging scheme, and studied the gains obtained in average paging cost and the average paging delay as the user location distribution is varied as per a geometric zone location probability. The plots we obtained showed the greatest gains are obtained when the user location is concentrated in a small portion of the location area.

We then compared the optimal sequential paging scheme to simpler suboptimal schemes which may be easier to implement, showing that such schemes can give near-optimal results over a wide range of location distributions. Finally, we considered the effect of errors in the estimates of the location distribution on the performance of sequential paging. It appears that sequential paging is fairly robust to small estimation errors, being most sensitive to these errors precisely where the greatest gains are obtained with perfect estimates - when the user location is concentrated in a few cells.

[3] Note that one of the data points of the $\eta = 0.05$ curve appears to be below optimum - this is due to experimental error; the data values for non-zero η are averages based on simulation whereas the zero-noise level curve represents the theoretical expectation.

References

[1] Christopher Rose and Roy Yates, "Minimizing the average cost of paging under delay constraints", *Wireless Networks*, 1995.

[2] T.H. Cormen, C.E. Leiserson, R.L. Rivest, *Introduction to Algorithms*, MIT Press, 1999.

[3] D. Munoz-Rodriguez, "Cluster paging for traveling subscribers," *Proc. IEEE Vehicular Technology Conference*, 1990.

[4] Bhaskar Krishnamachari, Rung-Hung Gau, Stephen B. Wicker, and Zygmunt J. Haas, "Optimal Sequential Paging in Cellular Networks," *unpublished* (available upon request), 2001.

[5] I. F. Akyildiz and J. S. M. Ho, "Movement-based location update and selective paging for PCS Networks," *IEEE/ACM Transactions on Networking*, vol. 4, no. 4, p. 629-638, December 1995.