State-Based Paging/Registration: A Greedy Technique*

C. Rose

March 16, 1999

Department of Electrical and Computer Engineering
WINLAB
Rutgers University
PO Box 909
Piscataway NJ 08855

ABSTRACT

For a mobile unit moving according to some ergodic stochastic process, we show how to minimize the expected average cost of paging and registration based on system state information. Specifically we assume that both the system and mobile unit (user) have access to or can derive

1. the user location probability distribution as a function of time given the last location known to the system, and
2. the time elapsed since this last known location.

We then derive methods by which the user decides whether to register based on,

- its current location, x,
- the time elapsed since last contact with the system, t,
- the paging cost $F(x_t, x_0, t)$ to be incurred by the system at time t given the current location x_t and the last known location x_0; i.e., the user knows the paging strategy to be used by the system for each time t.

If x and t define the system state, the method can be called "state-based". Since the optimization is based only on the current expected cost rate and not that of all registration intervals, the method is "greedy".

*This work supported in part by NSF grant NCRI-9206148.

The greedy method was compared to a timer-based method using a simple diffusive motion process. Reductions in average paging registration cost of approximately 10% were observed. The more striking improvement was a reduction in the variability of paging/registration costs by a factor of three. Thus, taking both cost *and* variability reduction as a performance measure, even suboptimal inclusion of location information in the registration decision affords substantial improvement.

1 Introduction

User tracking is usually accomplished through some combination of paging and registration for mobile communications systems. In previous work we have derived optimal paging strategies given a probability distribution on user location $p_x(X)$ [1]. In subsequent work we used a timer-based method to minimize joint paging/registration costs given the time-varying conditional user location probability distribution [2, 3].

For the timer-based method, registration is mandated if no paging request is received before some carefully chosen deadline. This deadline is measured from the time of last user contact with the system. Here, however, we assume that the user knows not only the current time t, but also

- Its own location x_t
- The probably distribution on future location x_s $(s > t)$ given the location at any time t.
- The cost of paging $F(x_t, x_0, t)$

We then find simple decision rules for when and where the user should register.

The optimization criterion is minimized expected paging/registration cost per unit time as a function of mobility, paging cost and page arrival rate. The minimization is based on the expected cost rate of paging/registration for the current *roaming interval* where the roaming interval is defined as the the time between a call termination or registration and the next paging or registration event. The method can be considered greedy since it does not explicitly optimize the average cost rate over all roaming intervals.

We illustrate our method on a simple time-varying Gaussian user location distribution which often arises as a result of isotropic random user motion [4]. We then study the improvement obtainable over the simple timer-based method [2, 3].

2 Problem Formulation

2.1 Background

Let $p_{x_t|x_0}(X_t|X_0)$ be the probability distribution on a user residing in location x_t at time t given that its location was X_0 at $t = 0$. It is assumed that this distribution, which depends on the underlying user motion, is either known, can be calculated or measured. Implicit in the $x_t|x_0$ notation is the idea that future location x_t for $t > 0$

depends upon x_0 in the Markovian sense of

$$p_{x_t|x_0,x_j}(X_t|X_0, X_j) = p_{x_t|x_0}(X_t|X_0) \tag{1}$$

for any $j < 0$. However, this need not be the case. In a more general notation, x might be replaced by a location process state variable β which satisfies the Markov property. However, for the sake of clarity we retain the $x_t|x_0$ notation while recognizing that it may be extended whenever necessary.

We define a paging cost function, $F(x_t, x_0, t)$, as the cost of paging given that the last known location (or location state) was x_0 at $t = 0$ and the current location is x_t at time t. The cost function $F(x_t, x_0, t)$ can be defined arbitrarily, but we define it here as the total area searched when looking for a unit to be consistent with previous work [5, 6, 2, 3]. Thus, if $N(x_t, x_0, t)$ is the number of locations searched given a last sighting at $(x_0, t = 0)$ and δ is the size of a location (cell), then $F(x_t, x_0, t) = N(x_t, x_0, t)\delta$. This cost corresponds to (fixed network and wireless) signaling cost associated with issuing polling requests at each location searched.

We define the average number of locations searched at time t given x_0 as $\bar{N}(x_0, t)$ and likewise $\bar{F}(x_t, x_0, t) = \bar{N}(x_t, x_0, t)\delta$. $\bar{N}(x_0, t)$ is minimized by searching locations in decreasing order of probability [2].[1] Hence, the minimum $\bar{F}(x_t, x_0, t)$ is

$$\min_{Y_n^*} \bar{F}(x_t, x_0, t) = \sum_{n=1}^{\infty} n p_{x_t|x_0}(Y_n^*|X_0)\delta \tag{2}$$

where Y_n^* is an ordering function such that $p_{x_t|x_0}(Y_n^*|X_0) \geq p_{x_t|x_0}(Y_{n+1}^*|X_0)$. Therefore, under the assumption of optimal paging, $F(x_t, x_0, t)$ is

$$F(x_t, x_0, t) = \phi\delta \tag{3}$$

where $x_t = Y_\phi^*$ and Y_n^* is the ordering function from equation (2).

The average number of polling events, and hence the average polling delay, is identical to the average number of locations polled. Although procedures exist whereby the average polling delay can be greatly reduced at the expense of modestly increased average number of locations searched [1], we will concentrate on the simpler case without delay constraints since it provides a lower bound on the average paging cost. It is worth noting, however, that the methods developed in this paper can be applied to any paging method so long as the mobile user and the system agree upon the paging cost $F(x_t, x_0, t)$.

[1] Assuming the user does not move appreciably during the paging process.

2.2 Average Paging/Registration Cost

We assume a user's location is known to the system just after a conversation, a paging event or a registration. We will concern ourselves with the time between the last known location and the next paging/registration event: the *roaming interval*. We do not consider intervals terminated by a call initiation since no additional cost is incurred. We then assign costs to registration events and seek to minimize the time average cost of paging/registration. The cost of a registration event might be measured in power/bandwidth used by the mobile handset and any associated network signaling.

Let χ_i be the trajectory of the motion process during roaming interval i. χ_i is unknown to the system but may or may not be explicitly known by the user. Let Θ be a registration decision rule used by the mobile; i.e., given the time t, the current location x_t, and the last known location x_0, Θ is used to determine whether to register or not. Define $\mathcal{C}_{i|\Theta}$ and $d_{i|\Theta}$ respectively as the paging/registration cost and duration of roaming interval i under registration decision rule Θ. The average paging/registration cost per unit time is then,

$$\eta_\Theta = \lim_{n\to\infty} \frac{\sum_{i=1}^{n} \mathcal{C}_{i|\Theta}}{\sum_{j=1}^{n} d_{j|\Theta}} \tag{4}$$

The law of large numbers allows us to rewrite this expression as

$$\eta_\Theta = \frac{\sum_{X_0} \sum_{\boldsymbol{\chi}} \bar{\mathcal{C}}_\Theta(\boldsymbol{\chi}, X_0) p_{\chi|x_0}(\boldsymbol{\chi}|X_0) p_{x_0}(X_0)}{\sum_{X_0} \sum_{\boldsymbol{\chi}} \bar{d}_\Theta(\boldsymbol{\chi}, X_0) p_{\chi|x_0}(\boldsymbol{\chi}|X_0) p_{x_0}(X_0)} \tag{5}$$

where $\bar{\mathcal{C}}_\Theta(\boldsymbol{\chi}, X_0)$ and $\bar{d}_\Theta(\boldsymbol{\chi}, X_0)$ are respectively the expected cost and expected duration of the roaming interval given motion trajectory $\boldsymbol{\chi}$, last known location X_0 and registration decision rule Θ. The policy Θ is a collection of points in space time $\{x_t, t\}$ at which registrations are mandated.

3 Finding a Decision Rule

3.1 Assumptions

Assume the cost of registration is unity and the cost of paging is $\mathcal{P}$ per location. Since paging events are generated by incoming calls, also assume that paging events form a Poisson point process of intensity λ_p independent of user motion. Since Poisson processes are memoryless, the time until the next paging event is described by the probability distribution $s_t(t) = \lambda_p \mathbf{e}^{-\lambda_p t}$ where t is referenced to the start of the roaming interval.

3.2 Searching for Good Θ: Greedy Approaches

Finding the optimal decision rule Θ for the cost function of equation (5) is difficult since there is a complex interaction of the policy Θ and $p_{x_0}(X_0)$ when the motion process is space-variant. In particular, if the conditional location distribution $p_{\chi|x_0}(\boldsymbol{\chi}|X_0)$ varies as a function of location, then the registration policy can affect the relative frequency $p_{x_0}(X_0)$ of the last known location.[2] Of course, we may circumvent this particular difficulty by assuming that the motion process is spatially invariant; i.e. the motion characteristics are identical for any starting state.

Unfortunately even if spatial invariance is assumed, finding the optimal Θ for this simplified problem is either impossible or difficult using standard approaches. For example, the continuous time variable t frustrates attempts to directly apply dynamic programming/Markovian decision processes [7, 8, 9], and discretizing time can lead to a large state space making the method impractical.

Therefore, we propose a slightly different, if possibly suboptimal, approach. Consider that at some time s and location x_s, the mobile knows only that

- no paging event has occurred on $[0, s]$,
- the current location x_s,
- the cost of paging, $F(x_s, x_0, s)$,
- and the conditional distribution on location x_t for $t > s$, $p_{x_t|x_s}(X_t|X_s)$[3]

Using this information, the mobile can compute the expected cost of paging and registration to be incurred by the system at some time τ in the future given the current location process state x_s and the starting state x_0. If the mobile does not register at time s and plans to register at time $s + \tau$, then the expected cost is,

$$\bar{F}(x_s, x_0, \tau) = E_{x_{s+\tau}|x_s, x_0, \tau, s}\left[F(x_{s+\tau}, x_0, s + \tau)\right] \tag{6}$$

where $E_{x_{s+\tau}|x_s}$ is the expectation with respect to the random variable $x_{s+\tau}$ given x_s, x_0, τ and s.

[2] The interested reader is referred to [2] where this problem is discussed in more detail for a simpler registration model. Also, methods of deriving $p_{x_0}(X_0)$ as a function of the policy are suggested.

[3] i.e., we assume the mobile knows the statistics of its own motion process. In fact, the mobile unit might even know its itinerary, in which case the random variable x_t would be deterministic.

We define C_n and T_n as discounted cost functions [7] of the n most recent paging/registration costs $\{c_k\}$, and roaming interval durations,$\{t_k\}$, $k = 1, 2...n$. Specifically,

$$C_n = \sum_{k=1}^{n} \alpha^k c_k \tag{7}$$

and

$$T_n = \sum_{k=1}^{n} \alpha^k t_k \tag{8}$$

where $0 \leq \alpha \leq 1$.

We can now form an expression for the expected cost per unit time assuming a planned registration at time $s + \tau$. C_n is the total discounted cost incurred in the previous roaming intervals and T_n is the corresponding sum of discounted durations of all previous roaming intervals. Without considering the current interval the empirical average cost would be $\bar{\eta} = C_n / T_n$. Paging cost is incurred if a paging request arrives before time τ, otherwise a registration cost is incurred. Thus, the expected cost of the current roaming interval is the sum of the expected paging and registration costs (integral term and final $\mathbf{e}^{-\lambda_p \tau}$ term respectively). Given that s units of time have elapsed in the current roaming interval, the expected total length of the current roaming interval is $s + \frac{1-\mathbf{e}^{-\lambda_p \tau}}{\lambda_p}$. Combining these expressions yields

$$\bar{\eta}(x_s, x_0, s, \tau) = \frac{C_n + \mathcal{P} \int_0^\tau \bar{F}(x_s, x_0, t) \lambda_p \mathbf{e}^{-\lambda_p t} dt + \mathbf{e}^{-\lambda_p \tau}}{T_n + s + \frac{1-\mathbf{e}^{-\lambda_p \tau}}{\lambda_p}} \tag{9}$$

The numerator is therefore the sum of all past roaming interval costs plus the expected cost in the current roaming interval and the denominator is the duration of all roaming intervals plus the expected duration of the current interval. This construction is analogous to that in equation 5.

We then devise the following procedure:

Greedy Registration Algorithm

```
DO UNTIL registration or paging event
    Minimize η̄(x_s, x_0, s, τ) in τ.
    IF optimal τ > 0
        do not register
    ELSE
        register
```

Said another way, the mobile unit registers at the instant s^* the expected cumulative per-unit-time cost of paging and registration is minimized. Since the decision to register depends only on the current and last known locations, the time elapsed in the current roaming interval and functions of the past history (C_n and T_n), the algorithm provides a state-based policy Θ_g. For $\alpha = 0$ we have $C_n = T_n = 0$ thereby producing a completely *greedy* algorithm which only seeks to minimize the cost of the *current* roaming interval. For $n \to \infty$ and $\alpha = 1$ we seek to minimize the cumulative average cost rate at each step. The level of greediness is controlled by α.

4 RESULTS

We first develop a decision rule with certain general restrictions on $\bar{F}(x_s, \tau)$. Then, for illustration we apply the decision rules to a simple linear diffusion model with drift which allows concise analytic expressions to be obtained for $\bar{F}$. The details are relegated to Appendix A. We then compare the performance of the greedy policy to that of a simpler timer-based method [2, 3] in which registration is mandated if no paging event occurs before some optimally chosen deadline. The timer-based method was previously compared to a space-based scheme [5, 6] in which fixed-boundary location areas were derived for each user during each roaming interval. Since at zero velocity this fixed-boundary scheme is essentially equivalent to the distance-based method used elsewhere [10, 11], we are effectively comparing all three previous methods, timer-based, space-based and distance-based to the greedy state-based method.

4.1 Mapping the Greedy Policy

Using the Greedy Algorithm, we may obtain a map of registration points by first identifying those points (x_t, t) for which the optimum waiting time τ is *not* zero – all *other* points must registration points. Thus, for each time s we seek positions x_s for which the minimizing τ in equation (9) is nonzero. Assuming that $\bar{F}()$ is differentiable in τ, we may find the minimizing τ by using standard analytic methods. Differentiating $\bar{\eta}$ with respect to τ yields,

$$\frac{\partial \bar{\eta}}{\partial \tau} = \frac{\mathrm{e}^{-\lambda_p \tau}}{T_n + s + \frac{1-\mathrm{e}^{-\lambda_p \tau}}{\lambda_p}} \left(\lambda_p (\mathcal{P} \bar{F}(x_s, x_0, \tau) - 1) - \bar{\eta} \right) \tag{10}$$

Another differentiation with respect to τ yields

$$\frac{\partial^2 \bar{\eta}}{\partial \tau^2} = - \left(\lambda_p + 2 \frac{\mathrm{e}^{-\lambda_p \tau}}{T_n + s + \frac{1-\mathrm{e}^{-\lambda_p \tau}}{\lambda_p}} \right) \frac{\partial \bar{\eta}}{\partial \tau} + \frac{\mathcal{P} \lambda_p \mathrm{e}^{-\lambda_p \tau}}{T_n + s + \frac{1-\mathrm{e}^{-\lambda_p \tau}}{\lambda_p}} \frac{\partial \bar{F}(x_s, x_0, \tau)}{\partial \tau} \tag{11}$$

Extremal points occur when equation (10) is zero.

Now, looking to equation (11), we see that for any extremal $\tau^* < \infty$ we have

$$\frac{\partial^2 \bar{\eta}}{\partial \tau^2}|_{\tau=\tau^*} = \frac{\mathcal{P}\lambda_p \mathbf{e}^{-\lambda_p \tau^*}}{T_n + s + \frac{1-\mathbf{e}^{-\lambda_p \tau^*}}{\lambda_p}} \frac{\partial \bar{F}(x_s, x_0, \tau^*)}{\partial \tau} \tag{12}$$

Notice that if $\bar{F}(x_s, x_0, \tau)$ is strictly increasing in τ[4], then

$$\frac{\partial \bar{F}(x_s, x_0, \tau)}{\partial \tau} > 0 \tag{13}$$

and *every* extremal of $\bar{\eta}$ is a local minimum; i.e., $\frac{\partial^2 \bar{\eta}}{\partial \tau^2}|_{\tau=\tau^*} > 0$. For any continuous function, at least one local maximum must appear between each pair of local minima. However, since $\bar{\eta}$ has no extremals which are maxima, it must have exactly one minimum on $0 \leq \tau < \infty$.

Thus, the restriction of $\frac{\partial \bar{F}(x_s, x_0, \tau)}{\partial \tau} > 0 \; \forall \tau$ allows a simple test for whether a point (x_s, s) is a registration point. If $\frac{\partial \bar{\eta}}{\partial \tau}|_{\tau=0} < 0$, then the optimal τ is clearly greater than zero and registration should not occur. Likewise, if $\frac{\partial \bar{\eta}}{\partial \tau}|_{\tau=0} \geq 0$, then the minimum must occur at $\tau = 0$ since any minimum at some $\tau > 0$ would require an (impossible) interposed maximum.

Since

$$\frac{\partial \bar{\eta}}{\partial \tau}|_{\tau=0} = \frac{1}{T_n + s}[\lambda_p(\mathcal{P}\bar{F}(x_s, x_0, 0) - 1) - \frac{C_n + 1}{T_n + s}] \tag{14}$$

and $\bar{F}(x_s, 0) = F(x_s, x_0, s)$, the known paging cost at time s and user location x_s, we obtain the following decision rule:

$$\mathcal{P}F(x_s, x_0, s) - 1 \quad \underset{\text{do not register}}{\overset{\text{register}}{\gtrless}} \quad \frac{C_n + 1}{\lambda_p(T_n + s)} \tag{15}$$

4.2 Application to a Simple Motion Model

We use a linear diffusion process with diffusion constant $\mathcal{D}$ and drift velocity v. Since this motion process is spatially invariant we may set $x_0 = 0$ and drop x_0 from our expressions for clarity. We then have (see Appendix A)

$$\bar{F}(x_s, \tau) = 2\sqrt{\frac{\mathcal{D}\tau}{\pi}}\mathbf{e}^{-(x_s - vs)^2/\mathcal{D}\tau} + 2(x_s - vs)\mathrm{erf}(\frac{x_s - vs}{\sqrt{\mathcal{D}\tau}}) \tag{16}$$

[4] Since location uncertainty usually increases with time since last sighting, the assumption $\frac{\partial \bar{F}(x_s, x_0, \tau)}{\partial \tau} > 0$ is not unreasonable. This point is further considered in the Conclusion.

Differentiation with respect to τ yields,

$$\frac{\partial \bar{F}(x_s,\tau)}{\partial \tau} = \sqrt{\frac{\mathcal{D}}{\pi\tau}}\mathbf{e}^{-(x_s-vs)^2/\mathcal{D}\tau} > 0 \tag{17}$$

for any finite $\tau \geq 0$.

From equation (23) we have (for $x_0 = 0$)

$$F(x_s, 0) = 2|x_s - vs| \tag{18}$$

so that the decision rule of equation (15) becomes

$$2\mathcal{P}|x_s - vs| - 1 \quad \underset{\text{do not register}}{\overset{\text{register}}{\gtrless}} \quad \frac{C_n + 1}{\lambda_p(T_n + s)} \tag{19}$$

which translates to

$$\begin{gathered}\textbf{Do not register if:}\\ -\tfrac{\lambda_p(T_n+s)+C_n+1}{2\mathcal{P}\lambda_p(T_n+s)} + vs \leq x_s \leq \tfrac{\lambda_p(T_n+s)+C_n+1}{2\mathcal{P}\lambda_p(T_n+s)} + vs\end{gathered} \tag{20}$$

Typical registration regions for $n = 0$ are depicted in FIGURE 1. The horizontal axis represents time s since last contact with the system and the vertical axis represents position. The registration region is gray; if ever the mobile unit resides at a (time,place) point in the gray area, a registration is mandated. As can be seen, the optimal greedy registration procedure provides registration boundaries which vary as a function of time and which track the expected position of the mobile unit and whose "width" depends upon the call arrival rate λ_p.

When seeking to minimize the expected cost of only the current interval, the user should almost never register for s small since the interval cost rate varies inversely with interval duration given a registration. In addition, for smaller λ_p, the region where the user is not required to register becomes larger; i.e., larger paging cost is tolerable since paging events are less frequent.

4.3 Preliminaries: performance evaluation

Calculating the average cost of the state-based method is sufficiently difficult that we resorted to Monte-Carlo simulation to evaluate η. Details are contained in Appendix B. For all cases we assume $\lambda_p = 1$ and $\mathcal{P} = 0.1$ for ease of comparison with previous results [3, 2].

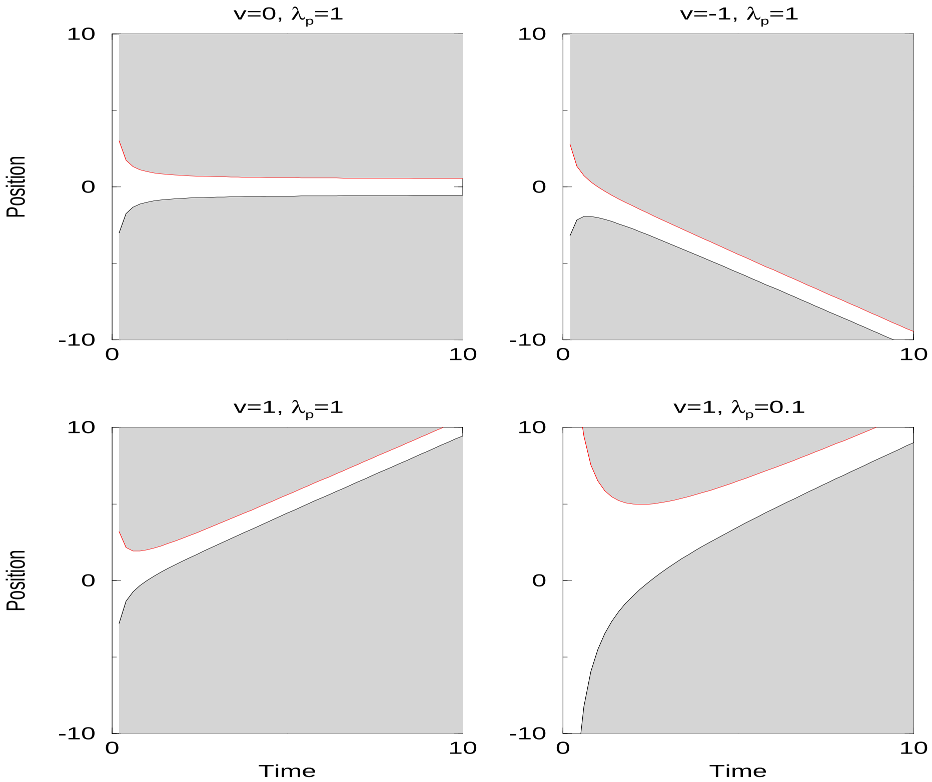

Figure 1: Registration regions using the greedy algorithm for various drift velocities v and call arrival rates λ_p. Paging cost factor $\mathcal{P} = 0.1$. If the mobile position and time correspond to a point in the gray region, registration takes place and the process is roaming interval timer is reset.

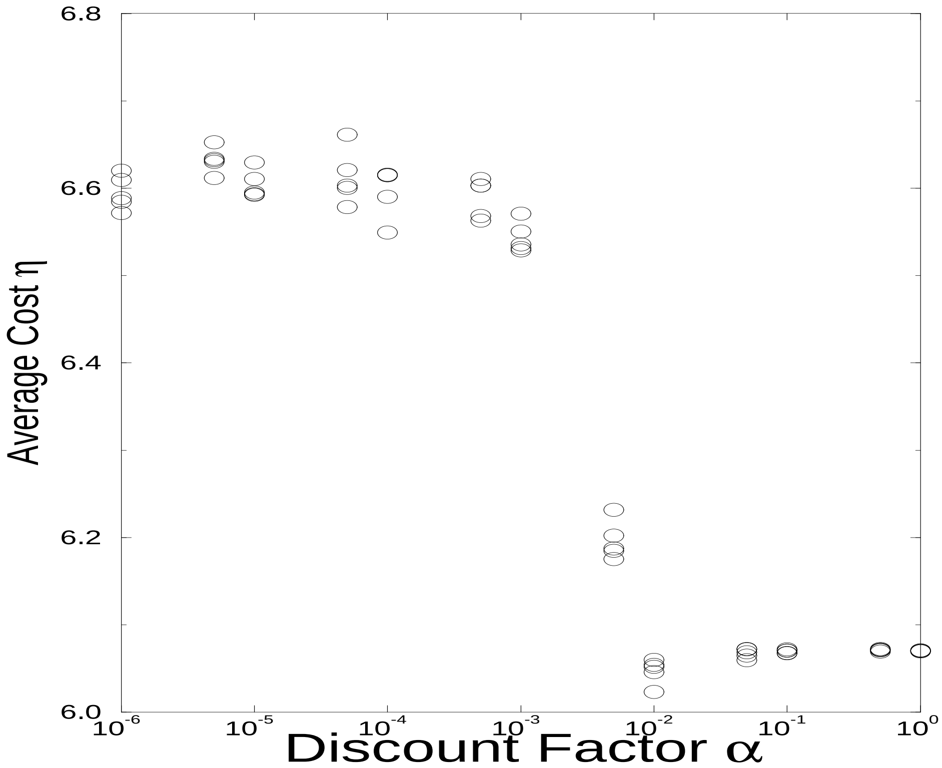

Figure 2: Variation of η with discount factor α for greedy algorithm. Five trials superimposed. $\mathcal{P} = 0.1$, $\lambda_p = 1$.

We first examined the variation of η with the discount factor α. The typical dependence is shown in FIGURE 2 where the average paging/registration cost is plotted as a function α for fixed $\mathcal{D} = 10^4$. Five trials are superimposed to show the range of experimental variability. As can be seen, η varies rapidly between its maximum and minimum values with α. The value of η achieved by the timer-based method for this case is $\eta = 6.57$[5]. Thus, the greedy algorithm, even with completely discounted history ($\alpha = 0$) performs on a par with the timer-based method. For undiscounted history, the performance improves by approximately 10%. This result was typical over a range of values for both v and $\mathcal{D}$.

We limited further tests to $\alpha = 1$ to achieve good greedy algorithm performance. We then examined the variation of η with velocity v. As shown in FIGURE 3, η did not vary systematically with v. Once again, five trials are superimposed to show experimental variability.

[5]Please see either Section 1 or the beginning of Section 4 for a description of the timer-based method

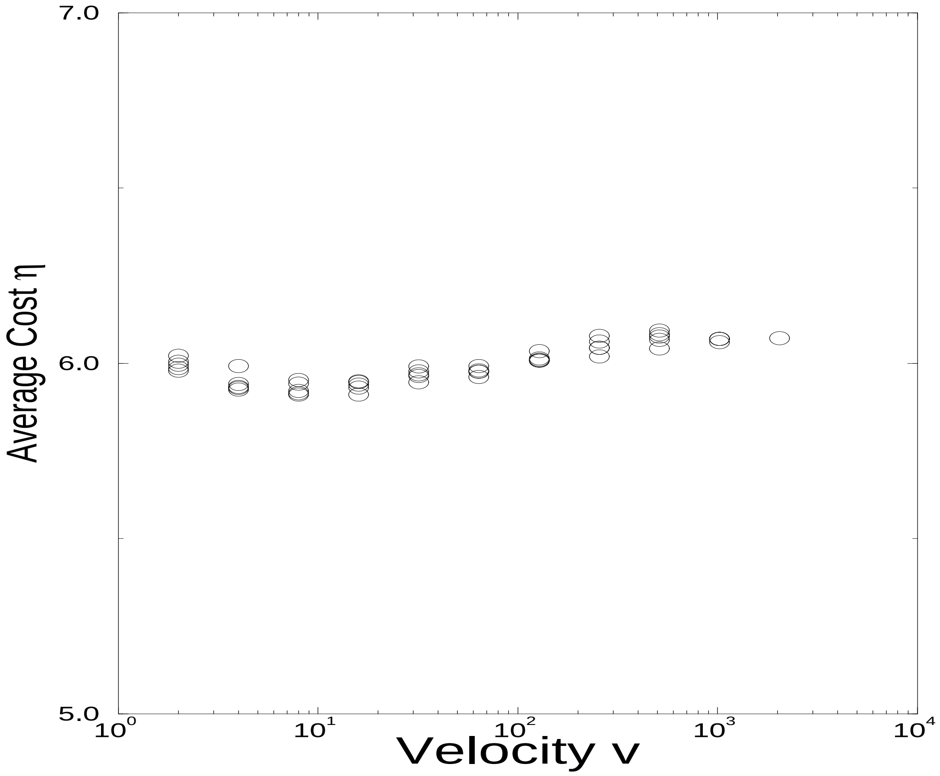

Figure 3: Variation of η for greedy algorithm with velocity v. $\mathcal{D} = 10^4$. Five trials superimposed. $\mathcal{P} = 0.1$, $\lambda_p = 1$.

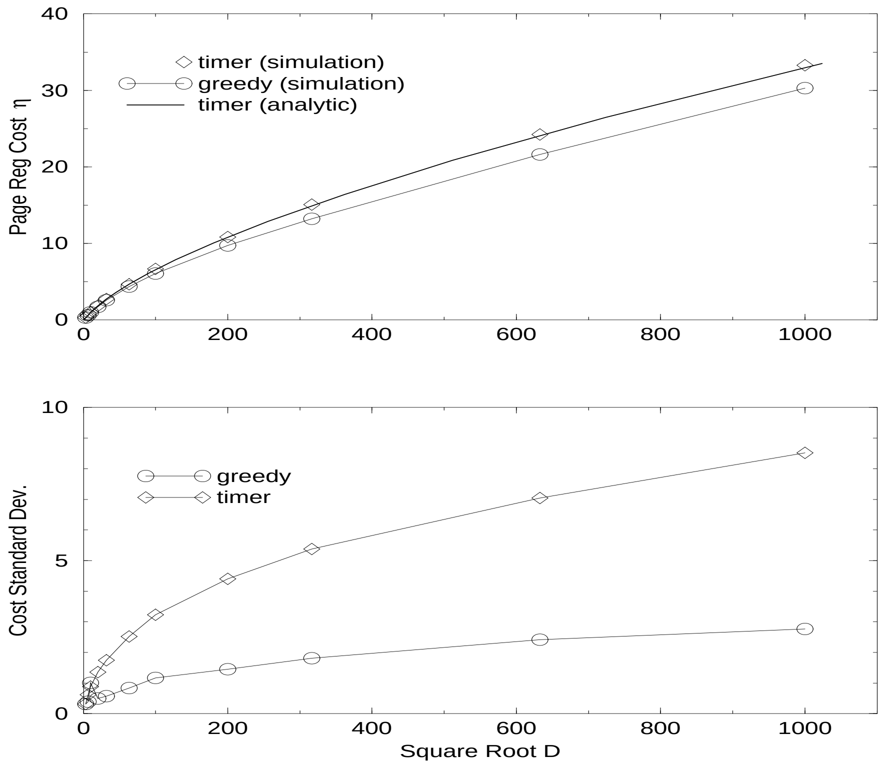

Figure 4: Comparison of timer-based and greedy registration method for various mobility indices. Upper plot: cost rate η vs. $\mathcal{D}$. Lower plot: standard deviation in roaming interval costs. $\mathcal{P} = 0.1$, $\lambda_p = 1$.

4.4 A Comparison to Timer-Based Registration

In FIGURE 4 we compare the average paging/registration cost rate η incurred by the greedy registration procedure to that of a timer-based method [2, 3]. As previously mentioned (Sections 1 and 4), the timer-based method mandates registration if no paging event occurs before an optimally chosen deadline. We also compared the standard deviation in cost per roaming interval achieved by the two methods. $v = 0$ was assumed for simplicity.

For the timer-based method, a theoretical curve and experimental values of η are shown. Experimental values are shown for the greedy algorithm with undiscounted past costs and interval durations; i.e., $\alpha = 1$ in equations (7) and (8). The greedy algorithm reduces average paging registration cost rates by approximately 10% as compared to the timer-based method for larger mobility indices.

However, the more pronounced difference is in the variability of the cost per roaming interval. The greedy algorithm shows much lower variability than the timer-based method. This result is not too surprising since the greedy algorithm has access to the current cost of paging and forces registration before the paging cost becomes too large. In comparison, the timer-based method has no access to the current paging cost and therefore might more frequently sustain large paging costs.

5 Summary and Conclusion

We formulated the general state-based paging/registration problem and found it to be difficult. We then simplified the problem by considering a "greedy" method which optimizes the average cost per unit time of the current roaming interval and incorporates past information regarding paging/registration cost rates. A timer-based method was previously shown to outperform other strictly distance or place-based methods under a variety of conditions [2, 3]. Thus, we restricted comparison of the greedy method to the timer-based method and found that the greedy technique achieved a 10% improvement in average cost while reducing the average cost variability by a factor of three.

The greedy method uses a simple decision rule (equation (15)) based on the point (time,place) of last contact with the system, the current position and the time since last contact. This result is similar to the optimal thresholding rule based on position derived in [11]. In this previous work paging cost was assumed only to be a function of location and not time. This assumption allowed a renewal theory approach for spatially invariant processes which resulted in optimal registration policies which were threshold rules. In contrast we consider time as a decision variable as well. Thus, though generally suboptimal, the greedy method presented here is applicable to a broader class of motion processes. For example, even the simple diffusion with constant drift model considered here does not satisfy the conditions for the place-based thresholding policy to be optimal since paging cost is explicitly a function of place *and time*.

The only requirement for application of equation (15) is that the expected paging cost increase monotonically with time from any known (time, place) point. Furthermore, even if this condition is relaxed, the general Greedy Registration Algorithm is still applicable, although explicit computation of η would probably be required.

As an example for which paging cost might not increase monotonically from time of last sighting, consider a commuter with the following daily cycle: home $\rightarrow$ commute $\rightarrow$ work $\rightarrow$ commute $\rightarrow$ home. Suppose the time of last sighting is just before the morning commute. If this commuter only receives calls once or twice a day, then the location uncertainty, and thus paging cost, varies cyclically from time of last sighting. However, if this commuter receives calls frequently relative the epochs associated with each mobility segment (home/commute/work), then the model of increasing variability in location with time is reasonable.

Although not explicitly pursued here, the same basic idea might be used with spatially varying motion processes as can be done for a timer-based method [2]. For the completely greedy algorithm where no past information is used, the application would be straightforward. Using past information, sets of $\{T_n\}_i$ and $\{C_n\}_i$ would have to be compiled corresponding to each last known location x_i; i.e., different thresholds in equation (15) for each location.

6 Suggestions for Future Work

We did not consider the joint optimization of paging strategy and registration strategy in that the paging cost incurred by the system might be significantly reduced if the state-based registration procedure were known to the system (as is the case for fixed location area methods [6]). For example, knowing the registration locations, the system need only search over non-registration locations at any given point in time.

Unfortunately, we believe such a game-theoretic problem [12] to be difficult in general. Nonetheless, a clearer theoretical understanding of the strategies most mutually beneficial for both user and system would greatly improve our ability to design robust and simple procedures for real systems.

A Paging Cost and Linear Diffusion with Drift

A number of motion models which are specified in terms of independent increments result in Gaussian distributions on location probability [4],[13]. The simplest is diffusive motion on a line with drift has probability density,

$$p_g(x_t|x_0,t) = \frac{1}{\sqrt{\pi \mathcal{D} t}} \mathrm{e}^{\frac{-(x_t - vt - x_0)^2}{\mathcal{D} t}} \tag{21}$$

where x_t is the location variable, $\mathcal{D}$ is the diffusion coefficient (units of length^2/time), v is the mean velocity and x_0 is the position at $t = 0$.

To obtain minimum average paging we must search each location in decreasing order of probability [1]. The most likely location is the mean of the distribution with symmetric equally likely locations to either side of the mean. Assuming that locations are quantized in steps of δ symmetrically about the mean vt, a mobile at position x_t will require

$$N(x_t, x_0, t, \delta) = \lfloor 2\frac{|x_t - vt - x_0|}{\delta} \rfloor + \mathrm{sgn}(x_t - vt - x_0) \tag{22}$$

polling operations. Note that the system assumes a last sighting of the mobile unit at position x_0 and time $t = 0$.

If we define the paging cost as the number of locations polled times δ we obtain $F(x_t, x_0, t) = N(x_t, x_0, t, \delta)\delta$. That is, we use *length* as a surrogate for number of locations searched analogous to [5, 6]. Assuming δ small we have

$$F(x_t, x_0, t) \approx 2|x_t - vt - x_0| \tag{23}$$

which has units of length for our one-dimensional motion process. We will now assume that δ is very close to zero and take equation (23) as an equality.

We may now calculate the expected cost at time $s + \tau$. For a spatially invariant motion process such as diffusion on a line we may assume $x_0 = 0$ with no loss of generality. Given that the user resides at x_s at time s, the distribution on location at time $s + \tau$ is

$$p(x_{s+\tau}|x_s, \tau) = \frac{1}{\sqrt{\pi \mathcal{D} \tau}} \mathrm{e}^{\frac{-(x_{s+\tau} - v\tau - x_s)^2}{\mathcal{D}\tau}} \tag{24}$$

The expected paging cost at time $s + \tau$ is then,

$$\bar{F}(x_s, \tau) = E[\mathcal{F}(x_{s+\tau})|x_s, \tau] = 2\int_{-\infty}^{\infty} |x_{s+\tau} - v(s+\tau)| \frac{1}{\sqrt{\pi \mathcal{D}\tau}} \mathrm{e}^{\frac{-(x_{s+\tau} - x_s - v\tau)^2}{\mathcal{D}\tau}} dx_{s+\tau} \tag{25}$$

Rewriting we have

$$\bar{F}(x_s,\tau) = 2\int_{-\infty}^{\infty} |x| \frac{1}{\sqrt{\pi\mathcal{D}\tau}} \mathrm{e}^{\frac{-(x-x_s+vs)^2}{\mathcal{D}\tau}} dx \tag{26}$$

which simplifies to

$$\bar{F}(x_s,\tau) = 2\sqrt{\frac{\mathcal{D}\tau}{\pi}}\mathrm{e}^{-(x_s-vs)^2/\mathcal{D}\tau} + 2(x_s - vs)\mathrm{erf}(\frac{x_s - vs}{\sqrt{\mathcal{D}\tau}}) \tag{27}$$

B Derivation of Simulation Parameters

A Monte-Carlo simulation of diffusive motion with drift parameter v and diffusion coefficient $\mathcal{D}$ requires proper choice of step increment Δx, time increment Δt and increment probability p. We provide simple derivations for these parameters in this appendix.

Let the process steps be of size Δx and let the steps occur at times which are integer multiples of Δt. Define $t = \Delta t T$ where T is the number of steps. We can define the particle position $x(t)$ as

$$x(t) = x(0) + \sum_{i=1}^{T} \Delta x_i \tag{28}$$

Assuming that each step is independent and $x(0) = 0$ we have approximately

$$p_{x(t)}(x(t)) \approx N[t\frac{\Delta x}{\Delta t}(2p-1), t\frac{\Delta x^2}{\Delta t}4p(1-p)] \tag{29}$$

I.e., the distribution is approaches normal for a large number of independent steps, T. Following convention [4], we define diffusion constant $\mathcal{D}$ and mean velocity v as

$$\frac{\mathcal{D}}{2} = \frac{\Delta x^2}{\Delta t}4p(1-p) \tag{30}$$

and

$$v = \frac{\Delta x}{\Delta t}(2p-1) \tag{31}$$

respectively.

The definition for v yields

$$p = (v\Delta t/\Delta x + 1)/2 \tag{32}$$

and combined with the definition of $\mathcal{D}$ yields

$$\Delta x = \sqrt{v^2 \Delta t^2 + \mathcal{D} \Delta t / 2} \tag{33}$$

which yields p as a function of Δt,

$$p = \frac{1}{2} \left(1 + \sqrt{\frac{v^2 \Delta t}{v^2 \Delta t + \mathcal{D}/2}} \right) \tag{34}$$

Thus by choosing Δt, Δx and thence p can be found.

Care was taken in the simulation to make the time and space steps sufficiently small so that quantization effects in the results were avoided. In addition, average results were obtained using at least $100,000$ roaming intervals.

References

[1] C. Rose and R. Yates. Minimizing the average cost of paging under delay constraints. *ACM Wireless Networks*, 1(2):211–219, 1995.

[2] C. Rose. Minimization of paging and registration costs through registration deadlines. In *IEEE International Conference on Communications ICC'95*, pages 735–739, 1995. Seattle.

[3] C. Rose. Minimizing the average cost of paging and registration: A timer-based method. *ACM Wireless Networks*, 2(2):109–116, June 1996.

[4] C. H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: algorithms & complexity.* Prentice Hall, 1982.

[5] S. Tabbane. An Alternative Strategy for Location Tracking. *IEEE Journal on Selected Areas in Communications*, 13(5):880–892, June 1995.

[6] H. Xie, S. Tabbane, and D.J. Goodman. Dynamic location area management and performance analysis. In *Proc. IEEE Vehicular Technology Conference, VTC'93*, May 1993. Secaucus, NJ.

[7] D.P. Bertsekas. *Dynamic Programming.* Prentice Hall, Englewood Cliffs, NJ, 1987.

[8] K.W. Ross and D. Tsang. Optimal circuit access policies in an isdn environment: A markov decision approach. *IEEE Transactions. Commun.*, 37(9):934–939, 1989.

[9] R.A. Howard. *Dynamic Programming and Markov Processes.* M.I.T. Press, 1960.

[10] U. Madhow. Power and bandwidth optimization for personal communications mobility tracking. In *Proceedings, IEEE Int. Symp on IT*, page 231, June 1993. Trondheim, Norway.

[11] U. Madhow, M.L. Honig, and K. Steiglitz. Optimization of Wireless Resources for Personal Communications Mobility Tracking. *IEEE Transactions on Networking*, 3(6):698–707, December 1995.

[12] E.M. Valsbord and V.I. Zhukovski. *Introduction to Multi-Player Differential Games and Their Applications.* Gordon and Breach Science Publishers, 1988.

[13] W. Feller. *An Introduction to Probability Theory and Its Applications, Volume I, Chapters IV & XIV.* Wiley, third edition, 1968.