Analysis of a Mobile-Assisted Adaptive Location Management Strategy

R. Yates, C. Rose,
S. Rajagopalan & B. Badrinath

WINLAB, Rutgers University
PO Box 909 Piscataway NJ 08855

ryates@ece.rutgers.edu, crose@ece.rutgers.edu,
subha@nyquist.bellcore.com, badri@cs.rutgers.edu

April 20, 1995

ABSTRACT

Source messages intended for a mobile host can be routed in one of two ways. Either the source knows the direct route to the mobile host, and is informed of all location changes by the mobile host (informed routing), or the source directs messages to a home agent that forwards messages to the mobile host (triangle routing).

When the rate at which the mobile host changes location and the rate at which messages are directed to the mobile host are known and fixed, we show that the optimal routing policy is described by a threshold rule that depends on the normalized differential route cost between triangle and informed routing and the call to mobility ratio (source messaging rate divided by location update rate).

Since the call to mobility ratio may not be known *a priori* or may change slowly with time, we also derive an adaptive policy selection algorithm. The policy is derived from a maximum likelihood estimate of the call to mobility ratio based on observations of message arrivals and location changes. The algorithm is found to work well when there is a clear advantage to either triangle or informed routing. However, when the two routing schemes are relatively close in average cost, the algorithm performance is degraded by repeated policy reversals. For this reason, algorithms which use hysteresis and/or a preset preference (preference threshold) for one routing scheme or another were explored. It was found that neither hysteresis, nor preference threshold techniques alone performed well, but rather a combination of the two resulted in greatly improved performance for a wide range of values of the call to mobility ratio.

1 Introduction

Consider a mobile host whose point of network attachment is time varying.[1] As proposed in [1], a home agent residing on the nominal home network of the mobile host can be used to forward messages to the mobile host. That is, when another network host (a source), wishes to send a message to the mobile host, it can send the message to the home agent that forwards the message to the mobile host.

Unfortunately, routing all messages through the home agent creates unnecessary load on the network if the direct path between sender and the mobile host is less costly. For a source that communicates frequently with the mobile host, this inefficiency becomes significant and a reduction in network traffic may be possible if the source were to communicate directly with the mobile host without forwarding through the home agent. To examine these costs, we propose that for every source, the mobile host must choose one of the following two policies:

- *Informed Routing:* the mobile host always notifies the source of location changes. The source sends messages directly to the mobile host.
- *Triangle Routing:* the mobile host does not notify the source of location changes. The source routes its messages via the home agent.

The general problem of routing over the Internet subject to host mobility has only recently come under scrutiny [1–6]. A concise review of this work is available in [7]. In other work such as [8] which addresses location management in the context of call setup in a PCS network, the informed routing policy is called *per-user caching* since the source caches the location of the mobile host. Each source decides whether to cache the location of a mobile host. Cache misses occur whenever the source uses outdated location information to attempt to contact a mobile host that has changed its location. To minimize the occurrence of cache misses, [9], considers a threshold strategy to determine whether a cached location is obsolete.

Here we follow the *working set* approach of [7, 8] that is analogous to the working set model for efficient memory management [10]. However, unlike [8], we assume that the *mobile host* dictates the routing strategy used by each source. As we shall see, this mobile-assisted strategy has the advantage that the mobile host can easily measure both its own rate of location changes and the messaging rate of each source and thus make appropriate routing decisions.

Note that this mobile-assisted approach operates as an adjunct to the basic triangle routing scheme. That is, the mobile host always updates the home agent, whether or not a particular source is updated directly. This is done so that any source can locate the mobile host and establish communication. Consequently, neither the problem of paging mobile users [11–16] nor cache misses [8, 9] occurs.

We identify the optimal policy by which a mobile host can determine those sources with which it should communicate directly. This optimal policy depends on the differential cost of forwarding over direct communication and the mobile host *call to*

[1] Internet, local cellular network, in-building LAN etc.

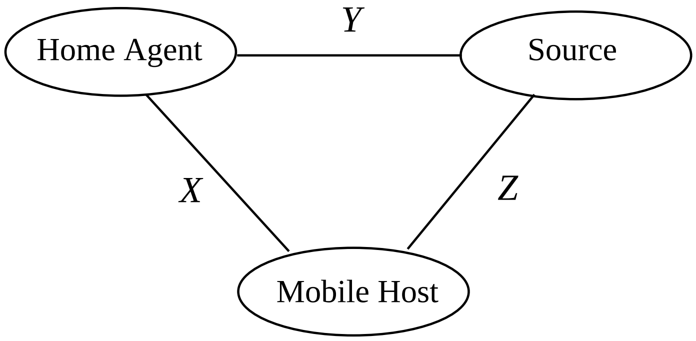

Figure 1: The model of network path costs.

mobility ratio (CMR). The CMR is defined as the ratio of source message rate to location update rate.[2]

In addition, we analyze an adaptive algorithm that adjusts to changes in the CMR. This adaptive algorithm is shown to perform near optimally over a broad range of system parameters. The lack of dependence of our approach on the network substrate should allow its application to many scenarios where information must be transferred to and from mobile hosts over some fixed network.

2 Non-Adaptive Optimal Location Management

We assume that the network has an underlying routing algorithm that is responsible for the establishment of communication paths between the source, the mobile host and the home agent. Given the complexity of routing algorithms that respond to changes in network load, we wish to adopt a simple model for the costs associated with these paths. In particular, we assume that messages between the mobile host and home agent, the source and home agent, and the source and mobile host incur costs described by the random variables X, Y, and Z, as depicted in Figure 1. We assume that each time a message is transmitted along a particular path, the cost of that message is independent of the cost of any other message, on that path or any other path. Furthermore, these communication costs are assumed to be independent of the movements of the mobile host.

We assume that source messages to the mobile host follow a Poisson process of rate λ. For the mobile host, location changes occur as a Poisson process of rate μ. Since the movement and message processes are assumed Poisson, location changes of the mobile host mark renewal periods. After each location change, the mobile host can choose one of the above policies by either notifying source or not. Since this constitutes the sole control action over the renewal period, either informed routing or triangle routing is an optimal policy.

[2] The CMR is the inverse of the *mobility index* [15, 17].

We will see that the performance of location management policies can be characterized in terms of the normalized expected differential route cost

$$\alpha = \frac{E[X] + E[Y]}{E[Z]} - 1 \tag{1}$$

and the CMR

$$\kappa = \frac{\lambda}{\mu} \tag{2}$$

To identify the optimal policy, consider a renewal process with renewals occurring whenever a source message or a location change takes place. Since the event process is Poisson with rate $\mu + \lambda$, the expected duration of the renewal period, i.e. the expected time between events, is

$$E[\tau] = 1/(\mu + \lambda) = \frac{1}{\mu}\frac{1}{1+\kappa} \tag{3}$$

Over the renewal period, the single event that occurs is either a location change with probability

$$\beta = \mu/(\lambda + \mu) = \frac{1}{1+\kappa} \tag{4}$$

or a source message with probability $1 - \beta$.

Let C_S and C_T denote the cost incurred under the informed routing and triangle routing policies during an arbitrary renewal period. Also, let $C_S(t)$ and $C_T(t)$ equal the corresponding total costs over the interval $(0, t)$. Under informed routing, each location update and each message from source to mobile host incurs an expected cost $E[Z]$. Over the renewal period, the expected cost under informed routing is

$$E[C_S] = \beta E[Z] + (1 - \beta)E[Z] = E[Z] \tag{5}$$

Under triangle routing, there is no cost associated with location changes since the source is never notified of location changes. Since each source message has an expected cost $E[X] + E[Y]$, the expected cost of triangle routing over the renewal period is

$$E[C_T] = (1 - \beta)\left(E[X] + E[Y]\right) \tag{6}$$

By renewal reward theory [18], the expected cost per unit time of each policy is

$$\gamma_S = \lim_{t\to\infty} \frac{C_S(t)}{t} = \frac{E[C_S]}{E[\tau]} \tag{7}$$

$$\gamma_T = \lim_{t\to\infty} \frac{C_T(t)}{t} = \frac{E[C_T]}{E[\tau]} \tag{8}$$

We see that informed routing is the optimal policy if $\gamma_S < \gamma_T$. That is, informed routing is optimal if

$$\mu E[Z] \le \lambda \left(E[X] + E[Y] - E[Z]\right) \tag{9}$$

Expressed in terms of the CMR κ and the differential cost α, the optimal policy is informed routing if

$$\kappa \geq 1/\alpha \tag{10}$$

and otherwise, triangle routing is optimal. Using $C^*(t)$ to denote the cost of the optimal policy over $(0, t)$, the optimal cost per unit time can be written

$$\gamma^* = \lim_{t\to\infty} \frac{C^*(t)}{t} = \mu E[Z]\,(\kappa + \min(1, \kappa\alpha)) \tag{11}$$

When the differential routing cost α and the CMR κ are both known, the optimal policy is easily specified. However, in practical systems, both the route costs and the message source rate may vary day-to-day or hour-to-hour. In this case, it will be necessary to estimate α and κ to determine an appropriate routing policy.

In this work, we assume that the variation of route costs with network load or time of day is modeled by the random variables X, Y and Z with *a priori* known mean values.[3] We further assume that these values change slowly with respect to the source message arrival rates and mobility rates. The problem then is to estimate the CMR κ using observations of location changes and source message arrivals. This estimate of κ will be the basis of an adaptive algorithm for location management.

3 Adaptive Location Management

In this section, we describe an adaptive location management strategy that responds to changes in the CMR κ. Let an epoch denote an interval of n events, where an event is either a location change or a source message transmission. The adaptive algorithm will use the location changes and source messages during an epoch to calculate $\hat{\kappa}$, a maximum likelihood estimate [19] of the CMR κ. During the next epoch, informed routing is used for the source if $\hat{\kappa} \geq 1/\alpha$.

Note that if during epoch k, the triangle routing was in effect, but during epoch $k+1$, the algorithm dictates that the informed routing be used, then at the start of epoch $k+1$, the mobile host must announce its location to the source. Similarly, if informed routing was in effect during epoch k, but triangle routing is to be used during epoch $k+1$, then the mobile host must send the source an invalidation message that instructs the source to route its packets via the home agent.

Let π_k denote the routing policy used in epoch k. That is, $\pi_k = S$ if informed routing is used; otherwise $\pi_k = T$. Suppose an arbitrary epoch k has just been completed. Since both location and message processes are Poisson with rates λ and μ, an event during the epoch is a location change with probability $\mu/(\mu+\lambda) = 1/(1+\kappa)$, independent of outcome of the other $n-1$ events in epoch k. Consequently, L_k, the number of location changes during epoch k, has the binomial distribution

$$P\{L_k = l\} = \binom{n}{l} \left(\frac{1}{1+\kappa}\right)^l \left(\frac{\kappa}{1+\kappa}\right)^{n-l} \qquad 0 \leq l \leq n \tag{12}$$

[3] The assumption of random route costs with fixed mean values is more general than but essentially similar to the standard assumptions of fixed routing costs; i.e. [8].

At the end of each epoch k in which $L_k = l$, we define the estimated CMR $\hat{\kappa}(l)$ by

$$\hat{\kappa}(l) = \arg\max_{\kappa} P\{L_k = l\} = \frac{n-l}{l} \tag{13}$$

Since the optimal policy is to use informed routing if $\kappa \geq 1/\alpha$, the policy for epoch $k+1$ is specified by the rule

$$\pi_{k+1} = \begin{cases} S & \hat{\kappa}(L_k) \geq 1/\alpha \\ T & \text{otherwise} \end{cases} \tag{14}$$

Hence, the probability that informed routing is used in epoch $k+1$ is

$$q = P\{\pi_{k+1} = S\} = P\{\hat{\kappa}(L_k) \geq 1/\alpha\} = P\left\{L_k \leq \frac{\alpha n}{1+\alpha}\right\} \tag{15}$$

Note that q is simply a function of n, α, and κ.

To deduce the cost per unit time of the adaptive algorithm, we must consider a renewal interval that accounts for the location announcement and invalidation costs. In particular, we can define a renewal process defined by those time instants at which an epoch is concluded and an announcement message is sent to the source. This announcement begins a sequence of N_S epochs during which informed routing is used. Since the decision to use informed routing in each epoch is independent of other epochs, N_S has a geometric distribution with mean

$$E[N_S] = \frac{1}{1-q} \tag{16}$$

Similarly, following an invalidation message, there will be N_T epochs in which triangle routing is used. Note that N_T has a geometric distribution with mean

$$E[N_T] = \frac{1}{q} \tag{17}$$

Over each epoch, the expected cost is $nE[C_S]$ when informed routing is in use or $nE[C_T]$ under triangle routing. In addition, over each renewal period, there is exactly one location announcement and one location invalidation, each at an expected cost $E[Z]$. Hence, the expected cost over the renewal period is

$$E[C_A] = nE[C_S]E[N_S] + nE[C_T]E[N_T] + 2E[Z] \tag{18}$$

The expected duration of the renewal period

$$E[\tau_A] = nE[\tau]E[N_S] + nE[\tau]E[N_T] \tag{19}$$

The cost per unit time of the adaptive algorithm is

$$\begin{aligned} \gamma_A &= \lim_{t\to\infty} \frac{C_A(t)}{t} && (20) \\ &= \frac{E[C_A]}{E[\tau_A]} && (21) \\ &= \mu E[Z]\left(\kappa + [q + \kappa\alpha(1-q)] + 2q(1-q)(1+\kappa)/n\right) && (22) \end{aligned}$$

Note that when $\kappa\alpha > 1$, then the adaptive algorithm chooses the correct policy with probability q. Similarly, when $\kappa\alpha < 1$, the correct policy is chosen with probability $1-q$. Note that in comparing γ_A to γ^* from (11), we identify components of γ_A. In particular, we can write

$$\gamma_A = \gamma^* + u_n(\kappa, \alpha) + f_n(\kappa, \alpha) \tag{23}$$

where

$$u_n(\kappa, \alpha) = \mu E[Z] \left(q + (1-q)\kappa\alpha - \min(1, \kappa\alpha)\right) \tag{24}$$

is the differential cost associated with uncertainty of which policy is correct. In addition, the cost of policy reversals that require location announcement and invalidation messages is

$$f_n(\kappa, \alpha) = 2q(1-q)(1+\kappa)/n \tag{25}$$

Choosing n large has two significant effects:

- The time between policy reversals is increased, decreasing $f_n(\kappa, \alpha)$.
- The accuracy of the CMR estimator is increased and either q or $1-q$ becomes closer to 1, decreasing both $u_n(\kappa, \alpha)$ and $f_n(\kappa, \alpha)$.

Hence, choosing larger n always reduces the long term expected cost when the message rate λ and movement rate μ remain constant. However, this estimation improvement is at the expense of reduced ability to adapt to changes in these rates.

4 Performance of Adaptive Location Management

For fixed α, κ and n, we see from equations (11) and (22) that γ^* and γ_A scale linearly in $E[Z]$ for fixed κ and α. Hence, we will assume $E[Z] = 1$. Moreover, in this section, we also assume a uniform cost model for the network in which the expected cost of each path is assumed to be the same. That is, we assume

$$E[X] = E[Y] = E[Z] = 1 \tag{26}$$

In this case, $\alpha = 1$.

We will compare γ^* and γ_A for a range of values of CMR κ. Note that the cost per unit time increases linearly with κ as the rate of source messages increases with κ. For purposes of comparison, it is much more instructive to compare the long term cost per source message. In particular if $M(t)$ is the number of source messages by time t, then $\lim_{t\to\infty} M(t)/t = \lambda$. The cost per message of the optimal policy is

$$\begin{aligned}
\hat{\gamma}^* &= \lim_{t\to\infty} \frac{C^*(t)}{M(t)} && (27)\\
&= \frac{\gamma^*}{\lambda} && (28)\\
&= E[Z]\left(1 + \frac{\min(1, \kappa\alpha)}{\kappa}\right) && (29)
\end{aligned}$$

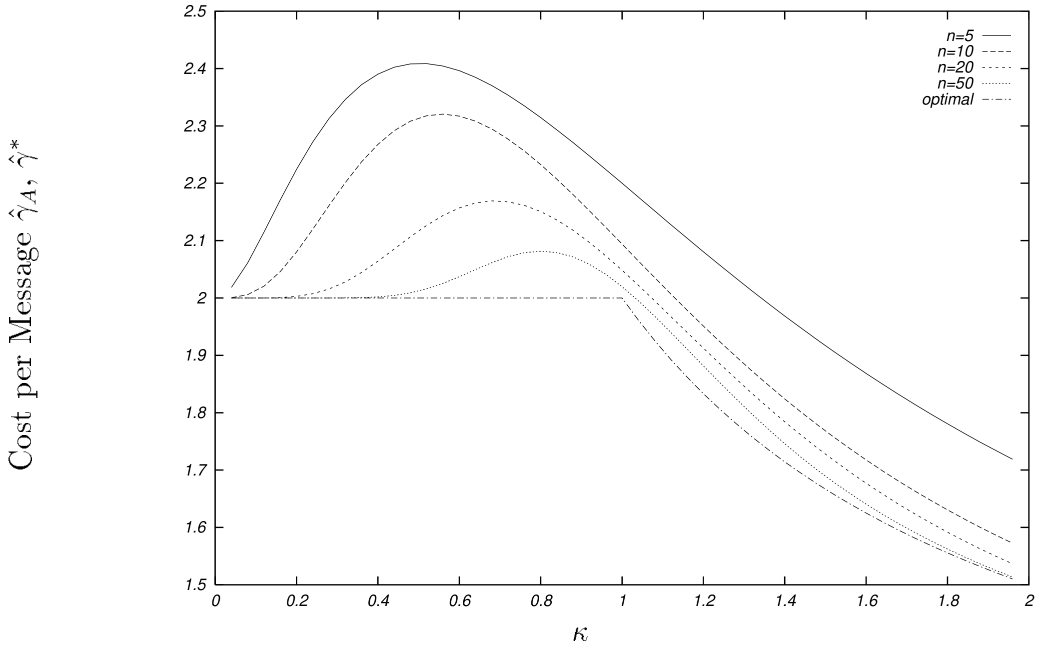

Figure 2: Cost per message for the adaptive algorithm ($\hat{\gamma}_A$) and for the optimal policy ($\hat{\gamma}^*$).

Similarly, the cost per message of the adaptive policy is

$$\begin{aligned} \hat{\gamma}_A &= \frac{\gamma_A}{\lambda} && (30) \\ &= E[Z]\left(1 + \frac{q + \kappa\alpha(1-q)}{\kappa} + \frac{2q(1-q)(1+\kappa)}{\kappa n}\right) && (31) \end{aligned}$$

In Figure 2, we compare $\hat{\gamma}_A$ to $\hat{\gamma}^*$ as a function of the CMR κ. We see that for $n \geq 20$, the adaptive algorithm performs only slightly worse than the optimal policy. In particular, we see that for very small κ or very large κ, the estimation of κ is sufficiently good that the optimal policy is almost always chosen. In these cases, the adaptive policy performs virtually as well as the optimal policy. However, when $\kappa \approx 1/\alpha$, either policy would be good, but the algorithm continually switches between policies. As n decreases, this reversal cost becomes increasingly significant.

5 Policy Selection Hysteresis and Offset

In the previous section, we observed that the adaptive policy performed nearly optimally when one policy was decidedly superior. However, from Figure 2, two deficiencies in the adaptive algorithm are apparent.

- When $\kappa \approx 1/\alpha$ and the two policies provide nearly equal long term costs, the reversals between policies penalized the adaptive policy.
- The threshold rule (14) using the ML estimator $\hat{\kappa}$ policy ignores the relative costs of the two policies. As a result, for low to moderate CMR, i.e. $\kappa < 1/\alpha$, the adaptive policy would do better if it were to choose the triangle routing more frequently.

In this section, we introduce adjustments to combat these two effects. In particular, we propose a hysteresis h that encourages the mobile host retain the existing policy and an offset factor d that simply increases the threshold to promote triangle routing. In combination, the hysteresis and the offset generate the two thresholds

$$\kappa_S = \frac{1}{\alpha}(1+d)(1-h) \tag{32}$$

$$\kappa_T = \frac{1}{\alpha}(1+d)(1+h) \tag{33}$$

The corresponding adaptation policy can be written

$$\pi_{k+1} = \begin{cases} S & \pi_k = S, \hat{\kappa}(L_k) \geq \kappa_S \\ T & \pi_k = S, \hat{\kappa}(L_k) < \kappa_S \\ S & \pi_k = T, \hat{\kappa}(L_k) \geq \kappa_T \\ T & \pi_k = T, \hat{\kappa}(L_k) < \kappa_T \end{cases} \tag{34}$$

Note that when the hysteresis h is zero, $\kappa_T = \kappa_S = (1/\alpha)(1+d)$. In this case, the adaptation policy is essentially the same as that used in the previous section, except the decision threshold has been offset from $1/\alpha$ to $(1/\alpha)(1+d)$. This encourages the use of triangle routing. For any offset $d \geq 0$, choosing $h > 0$ results in $\kappa_S < (1/\alpha)(1+d) < \kappa_T$. The effect of increasing h is that if a particular policy is used in epoch k, then the decision threshold requirement that must be met for that policy to be used in epoch $k+1$ becomes less stringent. Hence, the hysteresis h encourages the system to retain the same policy.

To analyze the effect of the offset and hysteresis, we say the mobile host is in state S when informed routing is the active policy or in state T otherwise. If in state S during epoch k, the mobile will transition to state T at the end of the epoch with probability

$$q_{ST} = P\{\hat{\kappa}(L_k) < \kappa_S\} = P\left\{L_k \geq \frac{\alpha n}{\alpha + (1+d)(1-h)}\right\} \tag{35}$$

Similarly, if in state T, the mobile will transition to state S at the end of epoch k with probability

$$q_{TS} = P\{\hat{\kappa}(L_k) \geq \kappa_T\} = P\left\{L_k \leq \frac{\alpha n}{\alpha + (1+d)(1+h)}\right\} \tag{36}$$

Once again, we consider a renewal to occur whenever the system reenters state S. In this case, $\tilde{N}_S$, the number of epochs spent in state S, has a geometric distribution with

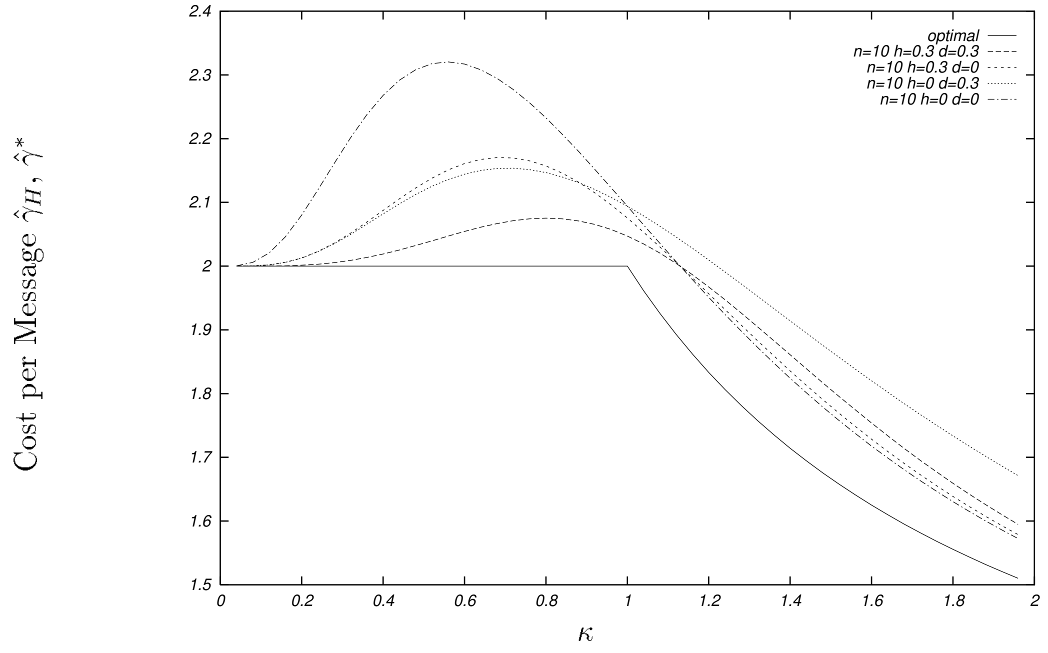

Figure 3: Cost per message $\hat{\gamma}^*$ for the optimal policy and $\hat{\gamma}_H$ for the adaptive algorithm with hysteresis h and offset d.

mean $1/q_{ST}$. Similarly, $\tilde{N}_T$, the number of epochs spent in state T is also geometric with mean $1/q_{TS}$. The expected cost of an epoch is $nE[C_S]$ under informed routing or $nE[C_T]$ under triangle routing. Since each renewal period also includes both a location announcement and a location invalidation, each at an expected cost of $E[Z]$, the expected cost of a renewal period is

$$E[C_H] = nE[C_S]E\left[\tilde{N}_S\right] + nE[C_T]E\left[\tilde{N}_T\right] + 2E[Z] \tag{37}$$

The expected duration of the renewal period is

$$E[\tau_H] = nE[\tau]E[N_S] + nE[\tau]E[N_T] \tag{38}$$

The expected cost per message becomes

$$\begin{aligned} \hat{\gamma}_H &= \frac{1}{\lambda}\lim_{t\to\infty}\frac{C_H(t)}{t} && (39)\\ &= \frac{1}{\lambda}\frac{E[C_H]}{E[\tau_H]} && (40)\\ &= E[Z]\left(1 + \frac{q_{TS} + q_{ST}\kappa\alpha}{\kappa(q_{TS}+q_{ST})} + \frac{2q_{TS}q_{ST}(1+\kappa)}{n\kappa(q_{TS}+q_{ST})}\right) && (41)\end{aligned}$$

From Figure 3, we see that for an appropriate choice of h and d, the adaptive algorithm can perform well over a wide range of CMRs κ. In addition, we see that either hysteresis or offset alone does not generally perform as well.

Since we have not assumed any particular *a priori* distribution on either the source message rate λ or the mobile host movement rate μ, it is not possible to use h and d to optimize the expected cost per message over the set of possible CMR values κ. However, Figure 3 shows that careful choice of h and d can result in an adaptive algorithm that provides significantly reduced costs for a wide range of possible values of CMR κ.

6 Conclusion

Source messages intended for a mobile host can be routed in one of two ways. Either the source knows the direct route to the mobile host, and is informed of all location changes by mobile host, or the source directs messages to a *home agent* which then forwards the messages to the mobile host. The former routing method is called *informed routing* and the latter is called *triangle routing*. The costs associated with each message sent via the either of the two methods are assumed random but with some known mean values.

If both the rate μ at which the mobile host changes location and the rate λ at which messages are directed to the mobile host are known and fixed, we have shown that the optimal routing policy can be determined using a threshold rule. This simple rule depends on the normalized differential route cost between triangle and informed routing (α), and the CMR ($\kappa = \lambda/\mu$).

However, since the mobility rate might not be known *a priori* or might change slowly with time, we also derived an adaptive policy selection algorithm. The policy selection is driven by maximum likelihood estimates of κ derived from observations of message arrivals and location changes. The algorithm works well for both larger and smaller κ where there is a clear advantage to either triangle or informed routing.

However, when the costs of either routing scheme are relatively close in mean value, the algorithm performance is degraded by repeated reversals of policy. Specifically, the source must be informed of any policy change. Since passing information incurs cost, if frequent policy changes occur owing to stochastic variation in the mobility index estimates, then information flow associated with unnecessary policy reversals becomes the dominant cost and performance suffers.

For this reason we explored algorithms which use hysteresis and/or a preset preference (preference threshold) for one routing scheme. It was found that neither hysteresis, nor preference threshold techniques alone performed well, but rather a combination of the two resulted in greatly improved performance for a wide range of κ. This robustness suggests that our adaptive routing policy decision algorithm might prove useful in a variety of practical situations.

References

[1] T. Blackwell *et al.* Secure short-cut routing for mobile IP. *USENIX Summer 1994 Technical Conference*, June. 1994.

[2] F. Teraoka and M. Tokoro. Host migration transparency in IP networks: The VIP approach. *ACM Computer Communications Rev.*, pages 45–65, Jan. 1993.

[3] J. Ioannidis and G.Q. Maguire Jr. The design and implementation of a mobile internetworking architecture. *USENIX Winter 1993 Technical Conference*, pages 489–500, Jan. 1993.

[4] W.A. Simpson. IP mobility support. *Mobile IP Working Group, IETF Internet Draft*, May 1994.

[5] C. Perkins and Y. Rekhter. Optimal routing for mobile hosts using IPs loose source route option. *Internet Draft*, Oct. 1992.

[6] T. Imielinski and B. Badrinath. Querying in highly mobile environments. *18th Proceedings VLDB*, Aug 1992.

[7] S. Rajagopalan. An adaptive location management scheme for mobile internetworking. Master's thesis, Rutgers University, January 1995.

[8] R. Jain, Y.-B. Lin, C. Lo, and S. Mohan. A caching strategy to reduce network impacts of PCS. *IEEE J. on Sel. Areas of Commun.*, 12(8):1434–1444, Oct. 1994.

[9] Y.-B. Lin. Determining the user locations for personal communications services networks. *IEEE Transactions on Vehicular Technology*, 43(3):466–473, Aug. 1994.

[10] P. J. Denning. The working set model for program behavior. *Comm. of the ACM*, II(5):323–333, May 1968.

[11] H. Xie, S. Tabbane, and D. J. Goodman. Dynamic location area management and performance analysis. In *Proc. of IEEE Vehicular Technology Conference*, 1993.

[12] A. Bar-Noy and I. Kessler. Tracking mobile users in wireless networks. In *INFOCOMM'93*, pages 1232–1239, March 1993. San Francisco, CA.

[13] J.Z. Wang. A Fully Distributed Location Registration Strategy for Universal Peronsal Communication Systems. *IEEE J. SAC*, 11(6):850–860, August 1993.

[14] C. Rose and R. Yates. Minimizing the average cost of paging under delay constraints. *ACM Wireless Networks*, 1995. (to appear).

[15] C. Rose. State-based paging/registration: A greedy technique. WINLAB-TR 92, WINLAB, Rutgers University, December 1994. (submitted for publication).

[16] C. Rose. Minimization of paging and registration costs through registration deadlines. *IEEE/ICC'95*, 1995. (to appear).

[17] C. Rose. Minimizing the average cost of paging and registration: A timer-based method. WINLAB-TR 83, WINLAB, Rutgers University, December 1994. (submitted for publication).

[18] S. M. Ross. *Stochastic Processes.* John Wiley & Sons, New York, 1983.

[19] A. Papoulis. *Probability, Random Variables, and Stochastic Processes, Third Edition.* McGraw-Hill, New York, 1991.