

An Auxiliary User Location Strategy Employing Forwarding Pointers to Reduce Network Impacts of PCS

Ravi Jain* and Yi-Bing Lin

Applied Research, Bellcore, Morristown, NJ 07960

Abstract. We propose a per-user forwarding strategy for locating users who move from place to place while using Personal Communications Services (PCS). The forwarding strategy augments the basic location strategy proposed in existing standards such as GSM and IS-41, with the objective of reducing network signalling and database loads in exchange for increased CPU processing and memory costs. With the forwarding strategy, calls to a given user will first query the user's Home Location Register (HLR) to determine the first Visitor Location Register (VLR) which the user was registered at, and then follow a chain of forwarding pointers to the user's current VLR. This strategy is useful for those users who receive calls infrequently relative to the rate at which they change registration areas. It can be shown, under certain assumptions, that forwarding can reduce total network costs by 40-60% for users with call-to-mobility ratio below 0.5.

1 Introduction

We consider the problem of locating users who move from place to place while using Personal Communications Services (PCS). Previous studies [9, 12, 8, 10] have shown that, with predicted levels of PCS users, there will be significant loads upon the signalling network and network databases, and that these loads are dependent upon the data management strategies adopted. We present a user location strategy which has the potential to reduce these loads significantly. The strategy we discuss here is an *auxiliary* strategy, in that it augments the *basic* user location strategies proposed in standards such as the North American IS-41 cellular standard [2] and the European GSM standard for mobile communications [11, 15]. (For surveys see [5, 14].)

The *basic* user location strategies proposed in the IS-41 [2] and GSM [11, 15] standards use a two-tier system of Home Location Register (HLR) and Visitor Location Register (VLR) databases (see also [14]). The user performs a registration at the HLR every

time the user changes registration areas, and deregisters at the previous VLR. The key observation we make is that, in many cases, it should be possible to avoid these registrations at the HLR, by simply setting up a forwarding pointer from the previous VLR. Calls to a given user will first query the user's HLR to determine the first VLR which the user was registered at, and then follow a chain of forwarding pointers to the user's current VLR; several techniques can be used to limit the maximum length of the chain. This observation results in a strategy which will be useful for those users who receive calls infrequently relative to the rate at which they change registration areas. This idea attempts to exploit patterns in the call reception and mobility patterns of PCS users, and is called *per-user forwarding*.¹

The per-user forwarding strategy attempts to reduce the network signalling and database loads of the basic strategies in exchange for increased CPU processing and memory costs. Since technology trends are driving the latter costs down, deploying the forwarding strategy on a system-wide basis will become increasingly attractive. Once deployed, whether the forwarding strategy should be invoked for a particular user is a function of the user's mobility and communications patterns.

The idea of forwarding itself is not novel, as it has been used in various forms in computer systems software designs (e.g. [3]). Forwarding strategies have also been proposed for locating mobile users in [1, 4]. However, in this paper we propose a different forwarding strategy, concentrate on the issues which arise in its application to a PCS telecommunication signalling network, and attempt to understand the classes of users and the system parameters for which the strategy is likely to yield benefits.

The outline of this paper is as follows. In sec. 2, we describe the PCS network architecture that we assume. (Some of this background material in this paper is repeated from our earlier paper [7] in order to keep

*Address correspondence to: Ravi Jain, Bellcore, 445 South St, Morristown, NJ 07960. Phone: 201-829-4756, Fax: 201-829-5888, E-mail: rjain@thumper.bellcore.com

¹We have previously designed an auxiliary strategy, called *per-user caching*, which attempts to exploit the calling and mobility patterns of a different population of users, namely those who receive calls frequently relative to the rate at which they change registration areas [7].

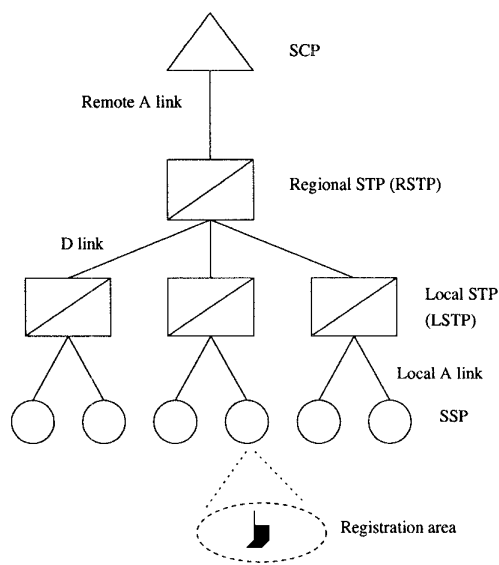


Figure 1: Reference CCS network architecture

this paper self-contained). In sec. 3, we describe the forwarding strategy. We use the definition of the user's call-to-mobility ratio (CMR) in sec. 4 to quantify the costs and benefits of forwarding for different classes of users.² Due to space limitations we only provide a brief sketch of the performance model; see [6] for details. Section 5 provides a summary.

2 PCS network architecture

We define the location of a PCS user, as known by the wire-line network, as the *registration area* (RA) in which the user is located. The geographical region covered by a PCS network is divided into radio port coverage areas, or *cells*. Each cell is primarily served by one radio *base station*, although a base station may serve one or more cells. The base station locates a user, and delivers calls to and from the user, by means of paging within the cell(s) it serves. Base stations are connected to the rest of the wire-line network by wire-line links. An RA consists of an aggregation of cells, forming a contiguous geographical region.

We assume that a Common Channel Signalling (CCS) network, with a Signaling System No. 7 (SS7) protocol [13], as commonly used for telecommunications, is used to set up calls. Fig. 1 illustrates the reference signalling network assumed in this study.³ The cells of the geographical region are served by base

²Note that the network, and not the user, bears the burden of maintaining information about the user's CMR [6].

³This architecture is meant as a reference architecture for a hypothetical geographical region, and is not necessarily the architecture corresponding to any particular implementation.

stations and are aggregated into RAs. The base stations of an RA are connected via a wire-line network to an end-office switch, or Service Switching Point (SSP). Each SSP serves a single RA. SSPs of different RAs are in turn connected to a two-level hierarchy of Signalling Transfer Points (STPs), comprised of a Regional STP (RSTP) connected to all Local STPs (LSTPs) in the region, which perform message routing and other SS7 functions.⁴ The RSTP is also connected to a Service Control Point (SCP), which is assumed to contain the functionality associated with a Home Location Register (HLR) database. For simplicity, we assume that in functional terms the MSC is collocated with the SSP, and a distinct Visitor Location Register (VLR) database is associated with each MSC [7]. (In the rest of this paper, the terms switch or SSP are used interchangeably). Each switch is assumed to serve exactly one RA, which, in turn, may be comprised of one or more cells.

3 Per-user Forwarding

Two operations which incur signalling are defined (we use the terminology of [7, 1]): *MOVE*, where the PCS user moves from one RA to another, and *FIND*, where the RA in which the PCS user is currently located is to be determined.

3.1 Basic user location procedures

We first present the *MOVE* and *FIND* procedures used in current PCS standards proposals such as IS-41 [2]; we call these the *BasicMOVE* and *BasicFIND* procedures.⁵ We use the following pseudo-code to describe the *BasicFIND*() and *BasicMOVE*() procedures.

```

BasicMOVE( ) {
  The mobile terminal detects that it is in a new
  registration area;
  The mobile terminal sends a registration message
  to the new VLR;
  The new VLR sends a registration message to the
  user's HLR;
  The HLR sends a registration cancellation message
  to the old VLR;
  The old VLR sends a cancellation confirm message
  to the HLR;
}

```

⁴In practice each STP actually consists of two STPs in a mated-pair configuration for redundancy [13]; for simplicity Fig. 1 only shows one of the two STPs of each mated pair.

⁵Note that the *BasicMOVE* and *BasicFIND* procedures we present only attempt to capture the major interactions between the HLR and VLR databases relevant to us. See [2, 15] for their detailed specification.

The HLR sends a cancellation confirm message to the new VLR;

Notice that the *BasicMOVE()* procedure involves notifying the HLR every time a user enters a new registration area. The registration information is used to find the user when a call is to be delivered to that user.

```

BasicFIND() {
  Call to PCS user is detected at local switch;
  if called party is in same RA then return;
  Switch queries called party's HLR;
  HLR queries called party's current VLR, V;
  VLR V returns called party's location to HLR;
  HLR returns location to calling switch;
}

```

3.2 Forwarding user location procedures

In the forwarding procedures, when a user moves from one RA to another it informs the switch (and VLR) at the new RA of the old RA from which it arrived. The switch at the new RA determines whether to invoke the basic *MOVE* or the forwarding *MOVE* strategy.

In *FwdMOVE()*, the new VLR exchanges messages with the old VLR to set up a forwarding pointer from the old to the new VLR, but does not involve the user's HLR. (See Fig. 2). A subsequent call to the user from some other switch will invoke the *FwdFIND()* procedure, which queries the user's HLR as in the basic strategy, and obtains an outdated pointer to the old VLR. The pointer from the old VLR is then followed to the new VLR to determine the user's correct current location. (See Fig. 3). To bound the time taken by *FwdFIND()*, and hence the call setup time, the length of the chain of forwarding pointers must be limited. The chain is allowed to grow to at most *K* pointers during the *FwdMOVE()* process. (We use the shared global variable *i* in the following pseudocode).

```

FwdMOVE() {
  /* Initially, i is 0 */
  if (i < K - 1) {
    User registers at new RA/VLR, passing id of former RA/VLR;
    New VLR deregisters user at old VLR;
    Old VLR sends ACK and user service profile to new VLR;
    i := i + 1;
  }
  else {
    BasicMOVE();
    i := 0;
  }
}

```

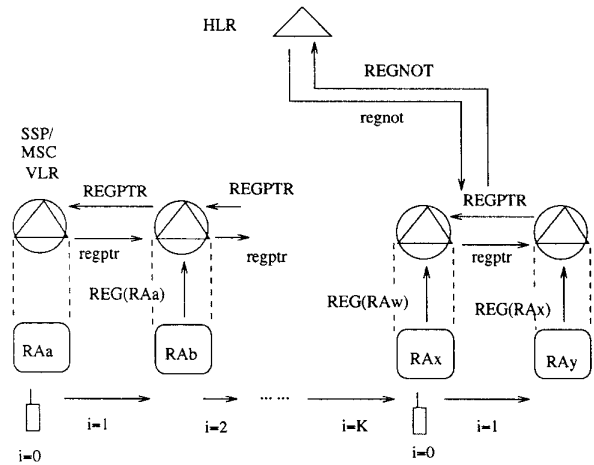


Figure 2: The *FwdMOVE()* procedure

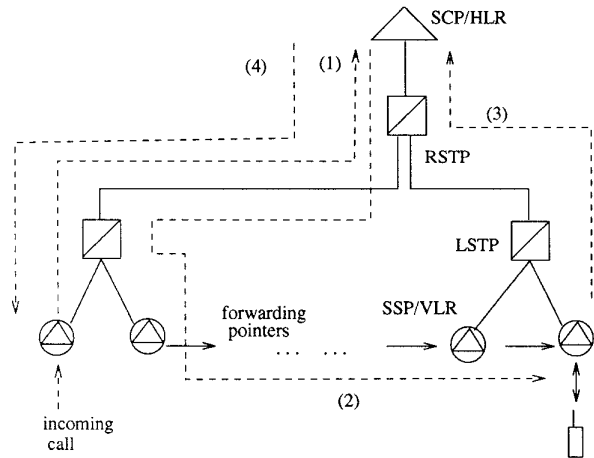


Figure 3: The *FwdFIND()* procedure

```

}
FwdFIND() {
  Call to PCS user is detected at local switch
  if called is in same RA then return
  Switch queries called's HLR
  HLR responds to caller's switch with V0
  Caller's switch queries V0
  while(Queried VLR ≠ called's current VLR)
    VLR queries next VLR in pointer chain
  /* Now called's actual VLR has been found */
  i := 0;
  Called's current VLR sends user location to HLR
  HLR sends user location to caller's switch
}

```




Figure 4: Relative total *MOVE* and *FIND* cost of forwarding: $\delta = 0.3$

4 Performance modeling

We have developed an analytic model to study the performance of per-user forwarding parameterized for different classes of users [6].

The *call-to-mobility ratio* (CMR) of a user is the average number of calls to a user per unit time, divided by the average number of times the user changes registration areas per unit time. (Note that the CMR is defined here in terms of the calls *received* by a particular user, not calls originated by the user). If λ is the mean rate of call arrivals to the user, and μ is the mean rate of RA crossings by the user, then the CMR, denoted p , is

$$p = \lambda/\mu \quad (1)$$

As we will see later, forwarding will only be applied to classes of users whose CMR is below a certain CMR threshold.

Suppose a user crosses several RAs between two consecutive phone calls. If the basic user location strategy is used, the *BasicMOVE* procedure is called every time the user moves to a new RA, in order to update the user's location at the HLR. If the forwarding strategy is used, the HLR is updated only every K th move, with forwarding pointers being set up for all other moves.

We calculate C_B and C_F , the total costs of maintaining user information and locating the user between two consecutive calls for the basic and forwarding strategies, respectively. Let m denote the cost of a single invocation of *BasicMOVE*, M be the total cost of all the *BasicMOVEs* between two calls, and F be the cost of a single *BasicFIND*. Then,

$$C_B = M + F = m/p + F \quad (2)$$




Figure 5: Relative call setup time for *FwdFIND* and *BasicFIND*: $\delta = 0.3$

Similarly, let M' be the average total cost of all *FwdMOVEs* between the two consecutive calls, and F' the average cost of the second *FwdFIND*. Then

$$C_F = M' + F' \quad (3)$$

We will be interested in two measures of performance: the ratio C_F/C_B , which reflects the average benefit of forwarding in terms of reduced total costs, and the ratio F'/F , which reflects the average additional cost paid during call setup for those users to whom call forwarding is applied. In [6] we have compared the costs of the forwarding strategy with that of the basic strategy for a specific cost model based upon the reference architecture of Fig. 1. For the moment we denote costs for various operations used in the forwarding strategy as follows.

S = Cost of setting up a forwarding pointer between VLRs during a *FwdMOVE*.

T = Cost of traversing a forwarding pointer between VLRs during a *FwdFIND*

As a first approximation we observe that updating the HLR and performing a *BasicFIND* involve the same number of messages between HLR and VLR databases, so we set $m = F$. We also assume that the cost of setting up a forwarding pointer is about twice the cost of traversing it, since twice as many messages are involved, i.e., we set $S = 2T$. Finally, we normalize $m = 1$, and $S = \delta$, with $\delta < 1$. Then, assuming Poisson call arrivals and an exponential distribution for the time that a user spends in an RA (called the *residence time*), it can be shown [6]

$$C_B = \frac{1+p}{p} \quad (4)$$

$$\frac{C_F}{C_B} = \frac{p}{1+p} \left(1 + \frac{3\delta}{2p} + \frac{1 - (K+1)\delta}{2((1+p)^K - 1)} \right) \quad (5)$$

$$\frac{F'}{F} = 1 + \frac{\delta}{2p} - \frac{\delta K}{2((1+p)^K - 1)} \quad (6)$$

In Figs. 4 and 5 we plot C_F/C_B and F'/F as functions of p for $\delta = 0.3$ and various values of K . In Fig. 4 we see that under the conditions we have assumed, for $\delta = 0.3$ and users with $p \leq 2$, forwarding can result in 20%-60% reductions in mean total network cost over the basic user location strategy. From Fig. 5 we see that this improvement comes at a price, namely increased mean call setup time. However, for $K \leq 4$, the penalty in call setup time is less than 25%.

In our performance model [6] we have also considered different values of δ , more general residence time distributions, analyzed the upper and lower performance bounds, and studied the sensitivity of the performance to residence time variance.

5 Conclusions

Depending upon the costs of traversing and setting up forwarding pointers relative to those of *MOVE* and *FIND* in the basic strategy, if pointer chains are kept short ($K < 5$), forwarding for users with low CMR ($CMR < 0.5$) can result in substantial reductions in mean total user location costs while paying reasonably low and infrequent penalties in mean call setup time. Note that the *per-user* nature of this forwarding scheme means that the maximum penalty in mean call setup time is not only paid infrequently, but only by users who receive calls infrequently; this tradeoff may be an acceptable price to pay for overall reductions in the network loads incurred in serving such users.

Acknowledgements. We thank Charles Lo for many useful and enlightening discussions. We also thank S. Mohan, A. Knapp, T. Ikuenobe, S. Levenson, J. Rizzo and R. Wolff for helpful discussions.

References

- [1] B. Awerbuch and D. Peleg. Concurrent online tracking of mobile users. In *Proc. SIGCOMM Symp. Comm. Arch. Prot.*, Oct. 1991.
- [2] EIA/TIA. Cellular radiotelecommunications intersystem operations. Technical Report IS-41 (Revision B), EIA/TIA, July 1991.
- [3] R. J. Fowler. The complexity of using forwarding address for decentralized object finding. In *Proc. SIGCOMM Symp. Comm. Arch. Prot.*, 1986.
- [4] J. Ioannidis, D. Duchamp, and G. Q. Maguire Jr. IP-based protocols for mobile internetworking. In *Proc. SIGCOMM Symp. Comm. Arch. Prot.*, pages 234-245, 1991.
- [5] Ravi Jain. A classification scheme for user location strategies in personal communications services systems. Aug. 1993. Submitted for publication.
- [6] Ravi Jain and Yi-Bing Lin. An auxiliary user location strategy employing forwarding pointers to reduce network impacts of PCS. *ACM Journal on Wireless Info. Networks*, 1995. To appear.
- [7] Ravi Jain, Yi-Bing Lin, Charles Lo, and Seshadri Mohan. A caching strategy to reduce network impacts of PCS. *IEEE J. Sel. Areas Comm.*, Oct. 1994.
- [8] C. Lo, S. Mohan, and R. Wolff. A comparison of data management alternatives for personal communications applications. Technical Memorandum TM-ARH-022333, Bellcore, June 1993.
- [9] C. N. Lo, R. S. Wolff, and R. C. Bernhardt. An estimate of network database transaction volume to support personal communications services. In *Proc. Intl. Conf. Univ. Pers. Comm.*, 1992.
- [10] Charles Lo and Richard Wolff. Estimated network database transaction volume to support wireless personal data communications applications. In *Proc. Intl. Conf. Comm.*, May 1993.
- [11] Edgar Lycksell. GSM system overview. Technical report, Swedish Telecom. Admin., Jan. 1991.
- [12] K. Meier-Hellstern and E. Alonso. The use of SS7 and GSM to support high density personal communications. In *Proc. Intl. Conf. Comm.*, 1992.
- [13] A. R. Modaresi and R. A. Skoog. Signalling System No. 7: A tutorial. *IEEE Comm. Mag.*, pages 19-35, July 1990.
- [14] S. Mohan and R. Jain. Two user location strategies for PCS. *IEEE Pers. Comm. Mag.*, 1(1), First Quarter 1994. (Premiere issue).
- [15] Michel Mouly and M. B. Pautet. *The GSM System for Mobile Communications*. M. Mouly, 49 rue Louise Bruneau, Palaiseau, France, 1992.