A Location Management Technique To Support Lifelong Numbering in Personal Communications Services *

Derek Lam, Yingwei Cui, Donald C. Cox, Jennifer Widom
Electrical Engineering & Computer Science Depts.
Stanford University, Stanford, CA 94305
{dlam, dcox}@wireless.stanford.edu, {cyw, widom}@cs.stanford.edu

This paper presents a novel *location management technique*, HOPPER, that is designed to support in a scalable and efficient manner non-geographical (lifelong) personal numbers in Personal Communications Services (PCS). Performance comparisons between our scheme and previous schemes are derived from large scale simulations using a realistic traffic modeling framework for the ten largest cities of the United States. Results show that, in addition to inherently providing non-geographical numbers, the proposed scheme significantly improves lookup performance and requires relatively little database access and network signaling resources.

1 Introduction

A Personal Communications Services (PCS) [3] network maintains per-user information about subscribers in *user profiles*. A user profile contains information such as current location of the subscriber, authentication information, billing information, and other information related to advanced services such as call blocking, etc. User profiles are stored in *databases* maintained by service providers. Subscribers are located in system-defined geographical areas called *registration zones*. Databases are assigned to serve their zones and are connected to the wireline signaling network. When a PCS user makes a call, the network performs a *location lookup* to obtain from the callee's profile the current location of the callee and other information relevant to call delivery. When a user moves from one zone to another, the network performs a *location update*, modifying information in the profile to track the user's new location. We address the *location management* problem: efficient retrieval and maintenance of user profiles concentrating on location information. To focus the problem we assume that other information, e.g., for billing purposes, is handled by other parts of the system.

The current standards, IS-41 [6] and GSM [20], manage profiles in the databases using an *HLR/VLR* scheme. In this scheme, there are two types of databases: *Home Location Registers* (HLRs) and *Visiting Location Registers* (VLRs). Subscribers are assigned an HLR in a specific registration zone to permanently store their user profiles. VLRs are used to store profiles for subscribers currently visiting a zone other than their "home" zone. When a subscriber moves to a new zone, the profile in the subscriber's HLR must be updated. During a call setup, the network must retrieve the callee's profile. If the profile is not found in the caller's local VLR, the callee's HLR must be accessed. In IS-41 and GSM, each subscriber is assigned a phone number that contains enough information for the network to derive the location of the HLR. These phone numbers are, therefore, *geographical* and are tied to specific HLRs. For example, subscribers cannot maintain their same phone numbers if they switch service providers. Furthermore, if a subscriber moves permanently to a new region, he will either need to obtain a new number or will incur the cost of having his HLR at a zone usually remote from his current location [13].

In the future, PCS subscribers will be assigned *Non Geographical Phone Numbers (NGPNs)* [12] that do not contain any information about their service provider or HLR. NGPN is an integral feature of PCS because it allows subscribers to have *lifelong* phone numbers. The main contributions in this paper are:

- We have developed a new *location management technique* (LMT), HOPPER, that uses a hierarchical organization of databases and on-line profile replication to achieve NGPN with scalability and efficiency. Hierarchical database architectures inherently support NGPNs and offer scalability [13]. Our on-line replication algorithm places copies of profiles at databases according to current call and mobility traffic to achieve efficiency: good dynamic lookup performance while requiring low database access and network signaling overheads.

- By performing extensive simulation studies, we have compared the performance of our proposal with other existing LMTs. Our dynamic performance simulations are based on realistic call/mobility models [13, 16] and a topology modeled after the ten largest cities of the United States. As will be seen, our scheme performs favorably against other schemes.

A shorter version of this paper can be found in [17]. In this extended version, we have included more simulation results to illustrate the performance of HOPPER. Interested reader can referred to [15] for more information on location management techniques in general. The rest of the paper is organized as follow. Section 2 reviews related research in this area. We present our proposed LMT, HOPPER, in Section 3. Section 4 shows simulation results comparing the performance of HOPPER against other LMTs. Finally, Section 5 concludes the paper.

2 Related Work

There are a variety of ways to provide lifelong NGPNs. One approach is a *number translation* technique, which works in

*Research partially supported by the National Science Foundation under grant# NCR-9628184 and Pacific Bell Corporation.

conjunction with the HLR/VLR scheme in the current standards. When the network needs to access an HLR, it first performs a lookup to translate the NGPN to the location of the subscriber's HLR. Once the location of the HLR is obtained, profile retrieval follows the conventional HLR/VLR protocol. Along these lines, [12] proposed using a *NGPN-to-HLR translation database* and presented a distributed hashing technique for maintaining the mapping between NGPN and HLR. [19] considered a technique that uses lookup tables and forwarding pointers to support NGPN. Our approach is to bypass number translation and directly search for profiles in the network. By organizing databases in an hierarchy and using replication techniques, we can support inherent NGPN without compromising scalability or efficiency. In Section 4 , we show performance comparisons between our proposal and the number translation scheme.

Other LMT research focuses on efficiency in PCS and on decreasing profile lookup latency. Various techniques, such as local anchoring [8], forwarding pointers [10], and caching [11], have been considered to lower the signaling impact of PCS. A technique for tracking locations of mobile users by maintaining regional, hierarchical directory servers has been proposed in [2]. Hierarchical database architecture [1, 26] have been proposed previously for local management. [1] describes a strategy for placing databases in a hierarchy that minimizes individual database loads; similar schemes are presented in [26]. Both studies dealt with only the architecture of the hierarchical scheme and did not investigate its performance.

[23] presents a network-flow algorithm to perform off-line per-user profile replication for HLR/VLR and in [13] we present a family of algorithms, HIPER, to perform off-line per-user replication for a database hierarchy. In this paper, we consider on-line replication—a natural evolution from off-line replication. Our results show that, because of its on-line nature, our proposal can achieve better dynamic lookup performance than previous off-line schemes, while at the same time requiring less memory overhead.

3 HOPPER

In this section, we describe our proposed LMT, *Hierarchical Online Parametric ProfilE Replication (HOPPER)*. We first describe a simple hierarchical LMT with no replication and then present the details of our proposal.

3.1 Basic Hierarchical Location Management

In a basic hierarchical LMT [26], the databases are organized in a multi-root tree hierarchy where the tree spanned by each of the roots belongs to one service provider (see Figure 1). Each *leaf database* serves a registration zone and stores profiles for subscribers currently located in that zone. Each *higher-level database* stores a pointer (subscriber ID + database ID) for every subscriber whose profile is currently stored in one of its descendent databases. The pointer is directed at the next lower level database that leads to the subscriber's current leaf database.

When X calls Y, the network first searches the local database at X's current zone for Y's profile. If this lookup fails, X's local database then propagates the query up the hierarchy to the first database that contains a pointer to Y's profile. The lookup query then propagates down the hierarchy following pointers until the leaf database with Y's profile is reached. In order to support multiple service providers, we require all root databases to cooperate and form a *distributed root* that contains pointers for all subscribers in the network. We assume that providers agree on a standard and allow each other to access their databases for user profiles. These assumptions are realistic given the current level of cooperation, e.g., in the European GSM cellular market.

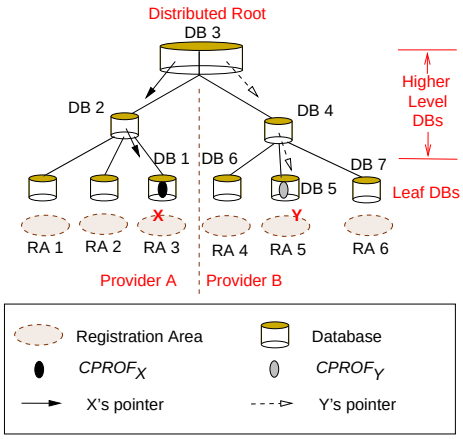

Figure 1: Hierarchical Database Architecture

When X moves from zone Z_i to zone Z_j, the PCS network must move X's profile to Z_j and delete the profile copy in Z_i. In addition, pointers in the higher-level databases spanned by the *least common ancestor (LCA)* of Z_i and Z_j, including the LCA itself, must be updated to reflect the new location of X.

We note that there is no concept of a "home database" for a subscriber as in HLR/VLR. Therefore, this architecture naturally supports lifelong NGPNs. The tradeoff is that in the hierarchical LMT location lookup may involve accessing several databases, thus potentially causing slow lookups and incuring high signaling loads. We now consider how to eliminate this disadvantage while maintaining NGPNs.

3.2 On-Line Profile Replication

We consider *replication* as a way to improve performance in a hierarchical approach. The expense of replication is increased memory and profile update costs. Replication is different from *caching* [23] in that location information in replicated profiles is kept up-to-date. Also, in contrast to replication schemes in traditional distributed databases where expensive distributed locking [22] is needed, we only require updates to be propagated to replicas. An off-line replication technique was presented in [23] for faster lookup in HLR/VLR. Call/mobility statistics are collected for each user, and then based on estimated user *local call to mobility ratio* (LCMR) [10], replicas are selectively placed at databases while satisfying certain constraints, such as maximum number of replicas per subscriber. In [13], we presented a family of LMTs called HIPER, by extending the idea of off-line replication to a hierarchy of databases. It has been shown previously in [13, 18, 23] that there is sufficient locality in calling and mobility behaviors

that replication is an efficient technique for improving performance.

In contrast to off-line schemes, our on-line replication technique uses on-line LCMR estimates and dynamically migrates replicas to appropriate databases. It offers the following advantages over off-line schemes.

- Off-line schemes estimate call/mobility traffic behavior and place replicas at locations to maximize expected benefits. On-line schemes can adapt to changes in user calling and mobility patterns by dynamically migrating replicas around the network.
- Due to practical storage constraints, it may not possible to place replicas at all locations where there are expected benefits. By dynamically adjusting the locations of replicas, on-line schemes can increase the "usefulness" of a replica and achieve the same lookup performance of off-line schemes with smaller storage requirements.

In the next subsections, we present HOPPER: a LMT that dynamically migrates replicas according to current traffic conditions. We show in Section 4 through simulations that HOPPER attains a better dynamic lookup performance with smaller storage requirements than previous proposed off-line schemes.

3.3 HOPPER Preliminaries

HOPPER's database architecture is based on the basic hierarchical scheme described in Section 3.1. HOPPER attempts to place profile replicas at leaf databases where the cost of replication does not exceed its benefits. We use LCMR as the metric for evaluating the cost and benefit of replication. We define $LCMR_{i,j} = C_{i,j}/M_i$, where $C_{i,j}$ is the number of calls to user i originating from zone j, and M_i is the number of moves for user i in a given time period. HOPPER computes on-line estimates of $LCMR_{i,j}$ and migrates replicas to databases with sufficiently high LCMR values. HOPPER requires the network to collect and maintain per-user call and mobility statistics (i.e. $C_{i,j}$ and M_i), similar to other previously proposed off-line replication and caching schemes [10, 13, 23].

To support on-line replication, we make the following additions to the basic hierarchical scheme. Let $CPROF_i$ denote the profile copy of user i in his current zone, $RPROF_i$ denote all the profile replicas for user i, and $RPROF_i[k]$ to denote the k^{th} replica of user i. For brevity, we also refer to the leaf database for zone Z_j as DB_j and the database serving the current zone of user i as $DB_{i,curr}$.

1. In addition to the usual profile information, $CPROF_i$ maintains a list containing the addresses of all databases with $RPROF_i$.
2. A location lookup for user i originating from DB_j will result, if a profile is not found locally, in $DB_{i,curr}$ sending a copy of i's profile back to DB_j. We add a token bit to the profile message. If the token bit is set, DB_j will treat the profile as a replica. Otherwise, it will just obtain the relevant information from the profile for call setup and either delete or overwrite it afterwards.
3. When a user moves from an old zone to a new zone, in addition to the location update tasks in the basic hierarchical LMT, the old $DB_{i,curr}$ must also update the location information in all replica copies indicated in $CPROF_i$. We note that consistency can be maintained if the old $DB_{i,curr}$ first updates location information in its own copy and does not delete it until some guard period after all replicas have been updated.
4. Each DB_j maintains $C_{i,j}$, the number of calls for each user i originating from Z_j.
5. Each $CPROF_i$ and each $RPROF_i[k]$ stores M_i, the number of moves for each user i in a given time period. When a user moves to a new zone, the network must update location information in all profile copies; therefore M_i can easily be maintained at each profile copy by incrementing it during a location update.
6. $CPROF_i$ contains a "token counter" that stores the number of replica tokens for user i (see Section 3.4.1).
7. Each database has an internal routing table that records its network hop distance from other databases in the network. It is expected that profile databases and network topology will not change very often. Thus, once the routing table is computed, it can be maintained in memory without much effort. The network hop distance is used in deciding where to move replicas (see Section 3.5).

3.4 HOPPER Parameterization

HOPPER is actually a family of algorithms based on its parameterization: HOPPER(N_{max}, R_{repl}, R_{del}, T_m, E_m, L) where N_{max} is the maximum number of replicas allowed for a user; R_{repl} and R_{del} determines whether a replica should be placed or deleted from a database, T_m and E_m define traffic statistics collection and on-line LCMR calculations, and L limits the amount of effort spent in on-line replication. We now discuss each of the parameters in detail.

3.4.1 Maximum Number of Replicas (N_{max})

N_{max} denotes the maximum number of replicas allowed for each user due to storage constraint considerations. We limit the number of replicas for user i by using a token counter method: $DB_{i,curr}$ is responsible for placing new replicas in the network and does that only when the token counter in $CPROF_i$ is not zero. The token counter in $CPROF_i$ is initialized to N_{max}. When a new replica is given out, the token counter is decremented. Similarly, if a replica is deleted from a database, $DB_{i,curr}$ must be notified and $CPROF_i$'s token counter will be incremented. We note that in HOPPER, replicas migrate between databases, but $CPROF_i$ controls the creation of new replicas.

3.4.2 On-line Replication Thresholds (R_{repl}, R_{del})

A common method to determine whether DB_j is a candidate for replicating i's profile is to consider $LCMR_{i,j}$. If $LCMR_{i,j}$ exceeds a certain threshold, R_{min}, DB_j is a candidate for replicating i's profile because the benefit of replication outweighs its costs. [7] proposes $R_{min} = 5$ for caching, and [23] suggests $R_{min} = 0.25$ for off-line replication in

HLR/VLR. [13] proposes two thresholds: R_{min} and R_{max} for off-line replication. R_{min} is set in such a way that if $LCMR_{i,j}$ is below it, there is definitely no gain in replicating at DB_j. Similarly, R_{max} is set so that if $LCMR_{i,j}$ exceeds it then it is always advantageous to replicate at DB_j. [13] also discusses optimal values for R_{min} and R_{max} given off-line traffic statistics and the database topology.

Unfortunately, unlike off-line schemes, it is not possible to derive "optimal" thresholds for on-line techniques without more statistical knowledge of user traffic behavior. It is also expected that on-line optimal values will vary as calling/mobility traffic patterns change. We choose a somewhat ad-hoc approach and pick the same fixed thresholds for all subscribers. As will be seen, this approach does not diminish performance. We use two thresholds: R_{repl} and R_{del}.

- If $LCMR_{i,j}$ exceeds R_{repl}, then we try to move i's profile replica to DB_j, provided that the token counter in $CPROF_i$ is not zero. The details of our *Replica Distribution Protocol* are discussed in Section 3.5.

- If DB_j currently has i's profile replica and its $LCMR_{i,j}$ falls below R_{del}, we delete i's replica from DB_j (see Section 3.5.3).

Intuitively, if R_{repl} is set too high, it takes a large $LCMR_{i,j}$ value before the system attempts to move a replica to DB_j, thus degrading dynamic lookup performance. However, if R_{repl} is set too low, databases with low $LCMR$ will compete for replicas, possibly depriving or delaying databases with high $LCMR$ from obtaining replicas. If R_{del} is set too close to R_{repl}, a database can "ping-pong" between assigning and deleting replicas, thus wasting signaling bandwidth. Conversely, if R_{del} is set too low, replicas can remain at locations with minimal benefit and waste update cost. In the next paragraph, we discuss how we pick our thresholds. From our simulation results in Section 4, we see that the suboptimal thresholds chosen still provide excellent lookup performance under realistic time-varying calling and mobility traffic conditions.

Recall from Section 3.1 that in a hierarchical LMT, multiple databases may need to be accessed during location lookups. In deciding our on-line thresholds, we consider the total number of database accesses during location lookups and updates. Consider a user i who is currently served by $DB_{i,curr}$. We also consider DB_j with an on-line $LCMR_{i,j}$ estimate of R. We interpret R as the expected number of calls for user i originating from zone j for each move of i. We let b_l be the cost (including signaling messages) of one database lookup and b_u be the cost (including signaling messages) of one database update. We also let D be the hop distance between $DB_{i,curr}$ and DB_j along the database hierarchy. Suppose there is a replica of i in DB_j. Then:

- When there is a call to i from DB_j, DB_j can always retrieve i's profile locally resulting in a lookup cost of b_l.

- Each time user i moves, the system needs to update the replica at DB_j resulting in a cost of b_u.

Now consider when there is no replica of i in DB_j.

- When there is a call to i from DB_j, location lookup involves lookups at all databases between DB_j and $DB_{i,curr}$ along the database hierarchy, and the profile is finally retrieved from $DB_{i,curr}$. The total number of lookups is $D+1$ resulting in a lookup cost of $b_l * (D+1)$.

- Since there is no replica, there is no associated replica update cost.

Therefore, we can see that by replicating in DB_j, the gain in database lookup cost for each call from DB_j to i is:

$$b_l * D \tag{1}$$

However, we tradeoff b_u update cost for each move of i. It is then easy to show that R must satisfy the following condition before replicating at j is beneficial:

$$R > \frac{b_u}{b_l * D} \tag{2}$$

We observe that D is bounded by D_{min} and D_{max}, the minimum and maximum hop distances between two databases in the hierarchy. As a result, it is always advantageous to replicate when $R > b_u/(b_l * D_{min})$. Similarly, we should never replicate when $R < b_u/(b_l * D_{max})$. We choose our thresholds to be these bounds. Thus, $R_{repl} = b_u/(b_l * D_{min})$ and $R_{del} = b_u/(b_l * D_{max})$.

3.4.3 *Traffic Statistics Collection (T_m, E_m)*

In order to compute $LCMR_{i,j}$, we need to collect and maintain $C_{i,j}$ and M_i. Recall that these statistics are stored in the leaf databases and profile copies. For any on-line scheme, the collection method must have the ability to "forget the past", so current decisions are not based on long ago activities. We choose a simple method and clear $C_{i,j}$ and M_i (i.e. set them to zero) at a fixed interval defined by parameter T_m. To avoid inaccurate $LCMR_{i,j}$ immediately after zeroing, we only use $LCMR_{i,j}$ when the total number of call and move events exceeds a threshold, E_m, i.e., $C_{i,j} + M_i > E_m$.

3.4.4 *Replica Search Distance (L)*

When $LCMR_{i,j}$ exceeds R_{repl}, the system tries to migrate a replica to DB_j. L limits the amount of signaling spent in this process. The use of L is discussed further in the Replica Distribution Protocol, covered next.

3.5 Replica Distribution Protocol

For our Replication Distribution Protocol, we first describe the general protocol in Section 3.5.1. In Section 3.5.2, we discuss how replicas are exchanged between databases and in Section 3.5.3, we cover deletions of replicas.

3.5.1 *General Protocol*

Our replica distribution protocol proceeds as follows. Suppose we have a call to user i from DB_j and there is no replica of i in DB_j. Location lookup will find $CPROF_i$ at $DB_{i,curr}$. Recall that location lookup is initiated from DB_j. In HOPPER, the following events occur.

1. Initially, DB_j sends $C_{i,j}$ along with the normal lookup message to its parent database. Eventually, the message will be received by $DB_{i,curr}$ where $CPROF_i$ is found.

2. Before $DB_{i,curr}$ sends i's profile to DB_j, it first calculates $LCMR_{i,j}$ using $C_{i,j}$ in the lookup message and M_i maintained in $CPROF_i$. There are three cases to consider:

   - **Case A:** If $LCMR_{i,j} > R_{repl}$ and $CPROF_i$'s token counter is not zero, then $DB_{i,curr}$ decrements the token counter and sends a profile to DB_j with the token bit set to one.
   - **Case B:** If $LCMR_{i,j} > R_{repl}$ but $CPROF_i$'s token counter is zero, then $DB_{i,curr}$ returns a profile with the token bit set to zero to DB_j. $DB_{i,curr}$ then executes the *Replica Exchange Protocol* (Section 3.5.2).
   - **Case C:** If $LCMR_{i,j} <= R_{repl}$, then $DB_{i,curr}$ simply sends a profile to DB_j with the token bit set to zero.

3. DB_j receives the profile from $DB_{i,curr}$. There are two cases to consider:

   - **Case A:** DB_j examines the token bit and finds that it is set to one. DB_j will store the profile as a replica and acknowledge $DB_{i,curr}$. When $DB_{i,curr}$ receives the acknowledgment, it will update its replica address list.
   - **Case B:** DB_j finds a zero token bit. DB_j will not treat the profile as a replica and will move on to its other tasks in location lookup.

3.5.2 Replica Exchange Protocol

We now describe the protocol for migrating user i's replica from one leaf database to another when the system has already distributed all N_{max} replicas of i, but finds a new database with $LCMR_{i,j} > R_{repl}$ (Event 2B in Section 3.5.1). We emphasize that the replica exchange protocol happens asynchronously with the normal location lookup procedure. As a result, replica exchanges do not delay location lookups. These events are "spawned" from Event 2B in Section 3.5.1.

1. After returning the profile to DB_j, $DB_{i,curr}$ randomly picks a database, say DB_r, from among its list of current replica addresses. If the hop distance, $d_{curr,r}$, between $DB_{i,curr}$ and DB_r, is smaller than L where L is the replica search distance parameter described previously, then $DB_{i,curr}$ sends $C_{i,j}$ (obtained from DB_j previously) to DB_r to consider "returning" its replica. If $d_{curr,r} > L$, no further event will occur. Note that the selection method is relatively simple. We discuss later in this subsection why we choose it and a possible alternate method.

2. DB_r receives the message from $DB_{i,curr}$. It compares the expected savings in database lookups between replicating at DB_j and DB_r by looking at the quantities $C_{i,j} * D_{curr,j}$ and $C_{i,r} * D_{curr,r}$ (see Equation 1). If $(C_{i,j} * D_{curr_j}) > (C_{i,r} * D_{curr,r})$, then DB_r deletes its replica of i and sends a message to inform $DB_{i,curr}$. Otherwise, no further events occur.

3. If $DB_{i,curr}$ receives notification that DB_r has deleted its replica, it then sends a profile copy to j with token bit set to one.

4. When DB_j receives the profile copy, it first acknowledges $DB_{i,curr}$ and stores the copy as a replica.

Since initially HOPPER gives out replicas in a first-come-first-serve manner, databases with relatively low $LCMR$ values can deprive databases with higher $LCMR$ values from obtaining replicas. Therefore, our replica exchange protocol attempts to migrate replicas to databases with higher $LCMR$ values to improve performance. $DB_{i,curr}$ randomly chooses a database already containing a replica in Event 1 to involve all replica databases in the migration process, as opposed to always picking the closest database. Because of random database selection, L serves to bound the maximum signaling per replica exchange. A more sophisticated selection method is to consider the entire replica address list in Event 1 and randomly pick a database from among those that satisfy the maximum signaling bound, L. However, we observe from our simulation results in Section 4 that the simple method performs very well. Thus, we choose it for its lower complexity. Also, $CPROF_i$ has control over how replicas get exchanged. During high traffic situations, e.g. when there are excessive location lookups from many different databases for user i, $CPROF_i$ can stop replica exchange altogether.

3.5.3 Deletion of Replicas

At each location update for user i, the old $DB_{i,curr}$ updates the location information in each replica. It waits for acknowledgements from all replica databases before forwarding $CPROF_i$ to the new $DB_{i,curr}$. When a replica database, DB_r, receives an update message, it updates its location information and increments its value of M_i. It also obtains $C_{i,j}$ locally to form $LCMR_{i,j}$. If $LCMR_{i,j}$ falls below R_{del}, then DB_r deletes the replica and notifies $DB_{i,curr}$ in its acknowledgement message. The old $DB_{i,curr}$ will increment $CPROF_i$'s token counter and update the replica address list before forwarding $CPROF_i$ to the new $DB_{i,curr}$.

3.6 Discussion

As described previously, HOPPER places replicas at leaf databases. While placing replicas at higher level DBs can potentially improve performance; however, this can significantly increase the complexity of the on-line replica placement algorithm. Our approach with HOPPER is to develop a low-complexity replica placement algorithm. As can be seen, HOPPER's algorithm is indeed simple to implement on-line. In fact, replica placement and deletion decisions can be made with simple threshold comparisons. Also, most of the traffic statistics and replica placement messages are carried or "piggy-backed" in existing location management messages. In the next section, we show that HOPPER can achieve excellent performance even with only placing replicas at leaf databases.

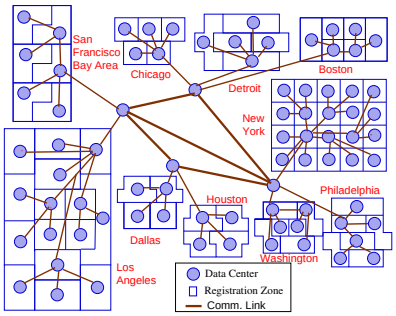

Figure 2: U.S. Simulation Geography and Signaling Topology

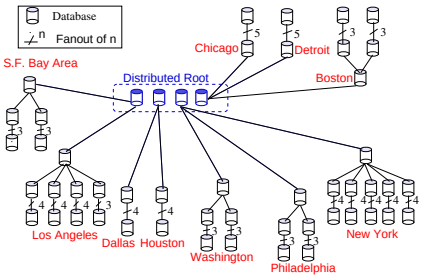

Figure 3: Database Hierarchy

4 Simulations

4.1 Simulator and Simulation Topology

We have studied the performance of various LMTs, including HOPPER, using our flexible discrete-event simulator, *Pleiades* [13, 18]. We performed our simulations on a topology modeled after the largest ten cities of the United States. Figure 2 shows our simulation geography and network signaling topology and Figure 3 shows our database topology where we consider a 4-level database hierarchy. Four databases form a distributed root (recall Section 3.1), storing pointers for all subscribers. The number of databases assigned for each city is proportional to its reported population in [25] with each one serving approximately one million subscribers. The number of registration zones is also proportional to the reported sizes of the cities. Each registration zone is approximately 400 square miles in area.

4.2 Dynamic Performance Simulation Results

For the following simulations, our goal is to investigate the performance of various LMTs under realistic traffic conditions. By using time-varying traffic models, we are able to study transient and peak performance of the LMTs.

Our call and mobility models have been discussed in [16, 18]. The call model includes various call traffic types and models call volume changes throughout the day. It also incorporates callee distribution by maintaining, for each caller, a list of the people they call most often with the probabilities of making calls to each of them. When a call is generated, the callee is selected either randomly from all users, or from the user's callee list according to its probability distribution. We have found empirically that the callee probability distribution can be modeled accurately by using *Zipf's Law* [18]. We have instantiated our call models with parameters we derived from call traffic data [24] from our university telephone exchange. We used the Metropolitan Mobility Model (METMOD) and the National Mobility Model (NATMOD) from [16] for modeling movements within and between the cities, respectively. These mobility models include various movement types and account for traffic volume changes over time. They are derived from vehicle traffic data [21], transportation surveys [5, 9, 14], and commercial airline data [4]. We now present simulation results comparing the performance of the following LMTs.

- **Basic HLR/VLR with NGPN translation.** We augment the basic HLR/VLR scheme to support NGPN by using an idealized model of the number translation technique (recall Section 2). We assume that during location lookup, a NGPN-to-HLR translation is required only when a copy of the profile is not found at the local database. We further assume that at every VLR, there is a local translation database. As a result, each NGPN-to-HLR translation incurs only one lookup and no network signaling messages. These assumptions correspond to the best possible scenario for the number translation scheme, and we ignore all updates and signaling activities associated with maintaining such a translation database at every VLR in the network. For brevity, this technique is referred to as *Simple HLR/VLR* for the rest of the paper.

- **Caching HLR/VLR with NGPN translation.** We simulate an HLR/VLR scheme with caching [10] for $R_{min} = 5$. We also augment this scheme to support NGPN by using an idealized model of the number translation technique. This technique does not constrain the maximum number of profile replicas (caches) per user. For brevity this technique is referred to as *Caching HLR/VLR* for the rest of the paper.

- **Simple Hierarchy.** We simulate the simple hierarchical LMT with no replication or caching.

- **HIPER.** We simulate HIPER(N,L) [13] which uses offline per-user replication. We set $N = 4$ to allow a maximum of four replicas per subscriber. We set $L = 3$ to allow replicas to be placed anywhere in our 4-level database hierarchy. User traffic statistics is collected every 24 hour and profile reallocations are also performed every 24 hour. This set of parameter values corresponds to the best lookup performance and results in the lowest database access and signaling loads for HIPER.

- **HOPPER.** We simulate HOPPER with $N = 4$ and $b_u = b_l = 1$. We set $R_{repl} = 0.5$ and $R_{del} = 0.125$ because $D_{min} = 2$ and $D_{max} = 8$ in our database topology, respectively. Also, $T_m = 24$ hrs, $E_m = 2$, and $L = 8$. L is set so that we are not restricting how we migrate replicas between databases.

We simulated 5,000,000 PCS subscribers, corresponding to approximately 6% of the total population of the cities. We have performed a multi-day simulation and are showing results for a representative 24 hour period because there is no significant difference between the results from different days, after start-up transients. Figure 4 shows the *local lookup percentage* versus simulation time. Local lookup percentage is defined as the percentage of all profile lookups that are satisfied at the local database. As expected, HOPPER attains the

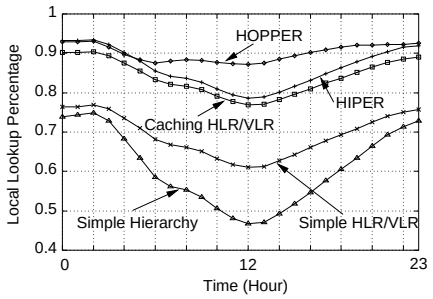

Figure 4: Local Lookup Percentage

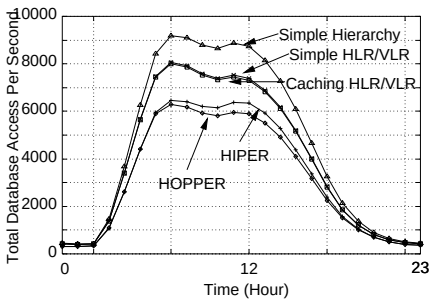

Figure 5: Global Database Access Rate

Scheme	Memory Requirements
Simple HLR/VLR	1.0
Caching HLR/VLR	2.57
HIPER	2.12
HOPPER	1.57
Simple Hierarchy	0

Table 1: Relative Memory Requirements

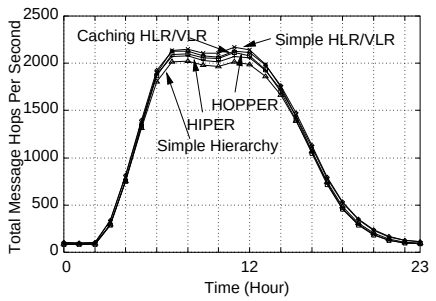

Figure 6: Global Signaling Hops

best lookup performance, averaging over 87% throughout the 24 hour period. Table 1 shows the system storage required for maintaining profile replicas. The HLR/VLR-based techniques (Simple HLR/VLR and Caching HLR/VLR) also maintain location information in HLRs. Therefore, HLR storage requirement are also included in Table 1. Effectively, Table 1 represents the additional storage requirements (in HLRs or for replicas) over what is needed to maintain information in users' current VLR or leaf databases. Entries in Table 1 have also been normalized to the requirements for Simple HLR/VLR. Simple Hierarchy has zero relative memory requirement because it only maintains information in a user's current leaf database and does not have any profile replicas. For the HLR/VLR-based techniques, we have ignored the storage requirements for the translation database. Since there is a translation database at every VLR, the actual memory requirement could be significantly higher. Even with this pessimistic assumption, HOPPER is still competitive having memory requirements not much more than HLR/VLR-based techniques. One interesting fact here is that by dynamically reallocating and deleting replicas, HOPPER is able to attain its improved lookup performance with a smaller memory requirement than off-line HIPER.

Figure 5 shows the total number of database accesses required by the different schemes. For all the schemes, we quote the total database activities from all subscribers' calling and mobility activities. For HOPPER, we have included the database activities introduced by on-line replication. For Caching HLR/VLR, we have included NGPN translation lookups. However, for HIPER, we did not include replication activities because they are expected to be done off-line. We see from Figure 5 that HOPPER incurs the smallest amount of database access, even after including on-line activities. HOPPER performs significantly better than Caching HLR/VLR in this category.

In Figure 6, we show the amount of network activity in terms of the global number of message-hops. In this category, we see that Caching HLR/VLR introduces the lowest signaling load. However, the absolute values between all five schemes are similar.

From the previous plots, it can be seen that HOPPER provides a good tradeoff in performance variables compared to the other existing schemes. In particular, HOPPER achieves very good location lookup performance. Since a somewhat ad-hoc approach has been used in determining HOPPER's replication thresholds, R_{repl} and R_{del}, we now examine the sensitivity of HOPPER's location lookup performance to the choices of these thresholds. The following plots show HOPPER's location lookup performance using the default parameters described previously with the exception of R_{repl} and R_{del}, which are varied independently. Figure 7 shows the local lookup percentage of HOPPER with $R_{repl} = 0.25, 0.5, 1.0, 2.0,$ and 4.0. R_{del} is set at its default value. It is observed that as the lookup performance of HOPPER degrades as R_{repl} is increased significantly beyond its default value. This is because for very high R_{repl} it takes longer before a database can reach the threshold for obtaining a profile replica. However, the actual variation is not large. HOPPER can still achieve a local lookup

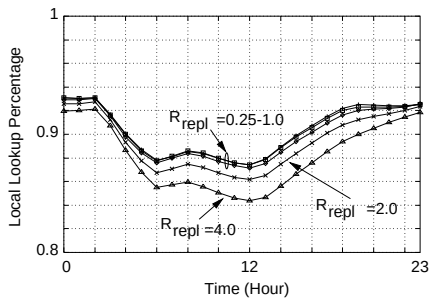

Figure 7: Local Lookup Performance of HOPPER with variations in R_{repl}

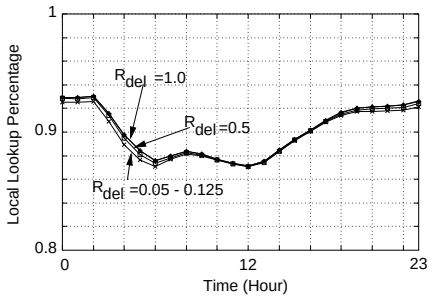

Figure 8: Local Lookup Performance of HOPPER with variations in R_{del}

percentage of over 82% throughout the 24 hour period even with $R_{repl} = 4.0$. Figure 8 shows the local lookup percentage of HOPPER with $R_{del} = 0.05, 0.125, 0.5,$ and 1.0. R_{repl} is set at its default value. It can be seen that lookup performance of HOPPER is not as sensitive to variations in R_{del}. The actual variation in local lookup percentage is less than one percent thoughout the 24 hour period. Due to space limitation, we cannot provide more simulation results comparing HOPPER and other techniques. Interested reader can referred to [15] for more details.

5 Conclusions

In this paper we have presented an efficient location management technique, HOPPER, that supports lifelong Non-Geographical Personal Numbers (NGPN). HOPPER supports NGPN by using a hierarchy of databases, and it achieves better lookup performance than previously proposed schemes by using an on-line user profile replication technique. We have performed extensive simulations on a model of the ten largest cities of the United States; we show that HOPPER achieves excellent lookup performance and has low database access requirements while incuring little signaling cost. Simulation results also show that HOPPER requires little additional memory comparing to the existing HLR/VLR-based scheme. We have also investigated how sensitive HOPPER's lookup performance is to the values of its threshold parameters. The simulation results show that HOPPER's lookup performance is not very sensitive to the values of its thresholds parameters around its carefully chosen default values. Thus, HOPPER is a competitive LMT for providing lifelong numbering given its excellent dynamic lookup performance and low database access requirements.

References

[1] V. Anantharam, M. L. Honig, U. Madhow, and V. K. Wei. Optimization of a database hierarchy for mobility tracking in a personal communications network. *Performance Evaluation*, 20:287–300, 1994.

[2] B. Awerbuch and D. Peleg. Concurrent online tracking of mobile users. In *Proceedings of SIGCOMM (SIGCOMM'91)*, pages 221–33, October 1991.

[3] D. C. Cox. Wireless personal communications: What is it? *IEEE Personal Communications*, 2(2):20–35, 1995.

[4] Origin and destination survey, t-100 domestic market data. Department of Transportation. Dataset for periods 10/91-12/91 and 10/93-12/93.

[5] ECMT. *European Transport Trends and Infrastructural Needs*. OECD publications, 2, Rue Andre-Pascal, 75775 Paris CEDEX 16, France, 1995.

[6] EIA/TIA IS-41.3 (Revision B). *Cellular Radiotelecommunications Intersystem Operations*, July 1991.

[7] H. Harjono, R. Jain, and S. Mohan. Analysis and simulation of a cache–based auxiliary location strategy for PCS. In *Proceedings of IEEE Conference on Networks and Personal Communications*, pages 1–5, March 1994.

[8] J.S.M. Ho and I.F. Akyildiz. Local anchor scheme for reducing location tracking costs in PCN. In *Proceedings of 1st ACM International Conference on Mobile Computing and Networking (MOBICOM'95)*, pages 170–180, November 1995.

[9] P. S. Hu and J. Young. *1990 NPTS Databook: Nationwide Personal Transportation Survey*. Federal Highway Administration, November 1993.

[10] R. Jain and Y. Lin. An auxiliary user location strategy employing forwarding pointers to reduce network impacts of PCS. In *Proceedings of International Conference on Communications (ICC'95)*, pages 740–744, June 1995.

[11] R. Jain, Y.-B. Lin, C. Lo, and S. Mohan. A caching strategy to reduce network impacts of PCS. *IEEE Journal on Selected Areas in Communications*, 12(8):1434–44, 1994.

[12] R. Jain, S. Rajagopalan, and L. F. Chang. Phone number portability for PCS systems with ATM backbones using distributed dynamic hashing. *IEEE Journal of Selected Areas in Communications*, 15(1):96–105, 1997.

[13] J. Jannink, D. Lam, N. Shivakumar, J. Widom, and D. C. Cox. Efficient and flexible location management techniques for wireless communication systems. In *Proceedings of 2nd ACM International Conference on Mobile Computing and Networking (MOBICOM'96)*, pages 38–49, November 1996.

[14] R. Kitamura. *Time-of-Day Characteristics of Travel: An Analysis of 1990 NPTS Data*, chapter 4. Federal Highway Administration, February 1995.

[15] D Lam. *Location Management Techniques and Teletraffic Modeling for Large Wireless Communications Networks*. PhD thesis, Stanford University, November 1997.

[16] D. Lam, D. C. Cox, and J. Widom. Teletraffic modeling for Personal Communications Services. *IEEE Communications Magazine Special Issue on Teletraffic Modeling, Engineering and Management in Wireless and Broadband Networks*, 35(2):79–87, 1997.

[17] D. Lam, Y. Cui, D. C. Cox, and J. Widom. A location management technique to support lifelong numbering in Personal Communications Services. In *Proceedings of 1997 IEEE Global Telecommunications Conference (GLOBECOM'97)*, volume 3, pages 704–710, 1997.

[18] D. Lam, J. Jannink, D. C. Cox, and J. Widom. Modeling location management for Personal Communication Services. In *Proceedings of IEEE International Conference on Universal Personal Communications (ICUPC'96)*, pages 596–601, September 1996.

[19] Y.-B. Lin. A cache approach for supporting life-time UPT number. *Wireless Networks*, 2(2):155–160, 1996.

[20] M. Mouly and M.-B. Pautet. *The GSM System for Mobile Communications*. Palaiseau, France, 1992.

[21] 1990 commute summary. Metropolitan Transportation Commission, Lotus 123 format spreadsheet, August 1990. S.F. Bay area transportation measurements.

[22] M. T. Oszu and Valduriez P. *Principles of Distributed Database Systems*. Prentice Hall, 1991.

[23] N. Shivakumar and J. Widom. User profile replication for faster lookup in mobile environments. In *Proceedings of 1st ACM International Conference on Mobile Computing and Networking (MOBICOM'95)*, pages 161–169, November 1995.

[24] Telephone call traffic data set, 3/95–9/95. S. Phillips, personal communication, October 1995. Encrypted ID numbers.

[25] International population data. U.S. Bureau of Census, International Data Base.

[26] J. Z. Wang. A fully distributed location registration strategy for universal personal communication systems. *IEEE Journal on Selected Areas in Communications*, 11(6):850–860, 1993.