ACM/Baltzer MONET (5)3, 2000, pp.191-198

Demand Re-registration for PCS Database Restoration

Zygmunt J. Haas [a] and Yi-Bing Lin [b]

[a] *School of Electrical Engineering, Cornell University, Ithaca, NY, 14853, U.S.A,*
E-mail: `haas@anise.ee.cornell.edu`, URL: `http://www.ee.cornell.edu/~haas/wnl.html`
[b] *Dept. Comp. Sci. & Info. Engr., National Chiao Tung University, Hsinchu, Taiwan, R.O.C.,*
E-mail: `liny@csie.nctu.edu.tw`

In a *Personal Communications Services* (PCS) network, mobility databases such as *Home Location Register* (HLR) and *Visitor Location Register* (VLR) are utilized to support mobility management for *Mobile Stations* (MSs). If the location databases fail, the location information loss or corruption will seriously degrade the service offered to the subscribers. In this paper, we propose a new VLR failure recovery scheme called demand re-registration. In this scheme, the VLR broadcasts a re-registration request to all MSs after the VLR failure. When an MS receives the re-registration request, it sends a registration message to the VLR to recover the location record. Since all MSs will re-register after receiving the broadcasting request, traffic jam (and thus collisions) may occur. If a collision occurs, then the involved MSs must re-send the registration messages. This paper studies the performance of demand re-registration by investigating how effectively the re-registration can recover the location record for an MS before the first MS call termination occurs. Our results indicate that demand re-registration can effectively recover VLR failure.
Keywords: Demand re-registration, Failure restoration, Mobility database, Periodic re-registration, Personal Communications Services (PCS), Registration, Visitor Location Register (VLR)

1. Introduction

In a *personal communications services* (PCS) network, mobility databases such as *Home Location Register* (HLR) and *Visitor Location Register* (VLR) are utilized to support mobility management for *Mobile Stations* (MSs). Details of mobility management have been addressed in literature [10,4,7]. For the benefit of readers who are not familiar with mobility management schemes, we describe the GSM (*Global System for Mobile Communications*) mobility management procedures in Appendix A. In particular, terms such as *registration*, HLR, VLR and *location area* (LA) and defined there.

If the location databases fail, the location information loss or corruption will seriously degrade the service offered to the subscribers. Thus, fault tolerance of location databases becomes an important issue for PCS network management. This paper studies VLR failure restoration [5,1,6].

When a VLR fails, the location record of a Mobile Station (MS) is automatically restored by one of the three events: MS registration, MS call origination, and MS call termination. Details of these events are given in Appendix B, which indicates that if the first event is an MS call termination, then expensive paging operations to all LAs are required. Thus, it is desirable to recover the location record by a registration operation before other events occur.

In *periodic re-registration* [5], the GSM system may ask the MSs to periodically re-register to the VLR. With periodic location updating, there is a better chance that, after a VLR failure, the location information is recovered by the periodic location confirmation before the arrival of the first call termination. Thus, expensive MS search operation is avoided. It is important to select an appropriate re-registration frequency. If re-registration operations are too frequent, then the signaling cost may be too high. On the other hand, if the frequency is too low, re-registration may be useless. The selection of the frequency for location re-registration was studied in [6].

This paper proposes a new VLR failure recovery scheme called demand re-registration. In this scheme, the VLR broadcasts a re-registration request to all MSs after the VLR failure. The cost of this single broadcast message is inexpensive and can be ignored. When an MS receives the re-registration request (following a procedure similar to call termination paging), it sends a registration message to the VLR to recover the location record. Since all MSs will re-register after receiving the broadcasting request, traffic jam (and thus collisions) may occur. If a collision occurs, then the involved MSs must re-send the re-registration messages.

This paper studies the performance of demand re-registration by investigating how effectively the re-registration can recover the location record for an MS before the first MS call termination occurs.

2. Demand Re-registration

This section describes several measures of the demand re-registration schene, when there are N MSs in the VLR:

- $E[n|N]$: the expected number of the elapsed timeslots before an MS has successfully re-registered; this measure indicates the elapsed time between the VLR failure and when demand re-registration is successfully performed for the MS.
- $E[K|N]$: the number of re-tries until an MS has successfully re-registered
- θ_d: the demand re-registration's contribution of saving expensive paging operation; the details will follow

2.1. Derivations for $E[n|N]$ and $E[K|N]$

Consider the behavior of an MS x after a VLR failure. To initiate the failure restoration procedure, the VLR broadcasts the re-registration request. When the MS receives this request, it waits for an exponential period τ_1 before it sends out the re-registration message. This delay is used to reduce the probability of repeated collisions when more than one MS tries to send re-registration at the same time (just like the exponential backup retry strategy used in slotted ALOHA Network).

In existing TDMA-based mobile systems such as GSM, message delivery through the radio interface is performed in timeslots. Thus, τ_1 can be represented by a geometric distribution, where the MS may issue a re-registration message at a timeslot with probability α. The selection of the parameter α can be determined at the time the VLR sends the re-registration request.

Suppose that on the average, there are N MSs in the VLR area. MS x successfully sends a re-registration message if it attempts to do so (with probability α) and the remaining $N-1$ MSs do not attempt to re-register (with probability $(1-\alpha)^{N-1}$). Let $p(n|N)$ be the probability that x successfully sends the re-registration message at the nth timeslot. For $n=1$, we have

$$p(1|N) = \alpha(1-\alpha)^{N-1} \tag{1}$$

Let $q(n,k|N)$ be the probability that k MSs (excluding x) have successfully sent the re-registration messages up to timeslot n; i.e., the MS x has not completed the re-registration procedure up to timeslot n, and k MSs will not send re-registration messages again.

For $n = 1, k = 0$, $q(1,0|N)$ is the probability that no MS successfully sends the re-registration message at timeslot 1. Since $\binom{N}{1}\alpha(1-\alpha)^{N-1}$ is the probability that an MS has successfully sent the re-registration message at timeslot 1, we have

$$q(1,0|N) = 1 - \binom{N}{1}\alpha(1-\alpha)^{N-1} \tag{2}$$

Suppose that MS x does not send the re-registration message at timeslot 1 (with probability $(1-\alpha)$) and only one of the remaining MSs does (with probability $\binom{N-1}{1}\alpha(1-\alpha)^{N-2}$), then

$$q(1,1|N) = \binom{N-1}{1}\alpha(1-\alpha)^{N-1} \tag{3}$$

The probability $p(n|N)$ is derived as follows. Suppose that j MSs (excluding x) have re-registered up to timeslot $n-1$ (with probability $q(n-1,j|N)$), and MS x successfully re-registers at timeslot n (with probability $\alpha(1-\alpha)^{N-j-1}$) where $0 \le j < \min(n,N)$. We have

$$p(n|N) = \sum_{j=0}^{\min(n,N)-1} q(n-1,j|N)\alpha(1-\alpha)^{N-j-1} \tag{4}$$

If no MS has re-registered up to timeslot n, it implies that no MS has re-registered up to timeslot $n-1$ (with probability $q(n-1,0|N)$) and no MS successfully re-registers at timeslot n (with probability $1-\binom{N}{1}\alpha(1-\alpha)^{N-1}$). In other words,

$$q(n,0|N) = q(n-1,0|N)\left[1 - \binom{N}{1}\alpha(1-\alpha)^{N-1}\right] \tag{5}$$

For $0 < k < \min(n,N)$, if k MSs (excluding x) have re-registered up to timeslot n, it implies that one of the following two situations occurs:

Case I. $k-1$ MSs (excluding x) have re-registered up to timeslot $n-1$ (with probability $q(n-1,k-1|N)$), MS x does not re-register at timeslot n (with probability $1-\alpha$), and one of the remaining $N-k$ MSs successfully re-registers at timeslot n (with probability $\binom{N-k}{1}\alpha(1-\alpha)^{N-k-1}$), or

Case II. k MSs (excluding x) have re-registered up to timeslot $n-1$ (with probability $q(n-1,k|N)$), and no MS successfully re-registers at timeslot n (with probability $1-\binom{N-k}{1}\alpha(1-\alpha)^{N-k-1}$)

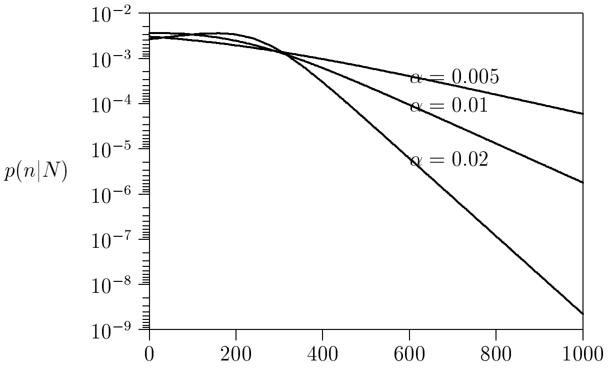

(a) $N = 100$

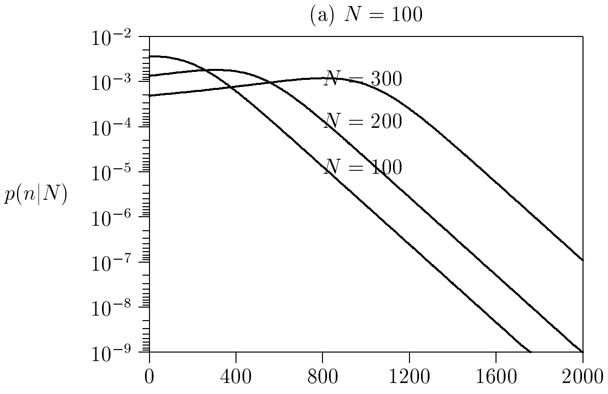

(b) $\alpha = 0.01$

Figure 1. The effect of α and N on $p(n|N)$

and we have

$$\begin{aligned} q(n, k|N) = & \, q(n-1, k-1|N) \left[\binom{N-k}{1} \alpha (1-\alpha)^{N-k} \right] \\ & + q(n-1, k|N) \left[1 - \binom{N-k}{1} \alpha (1-\alpha)^{N-k-1} \right] \end{aligned} \quad (6)$$

For $k = \min(n, N)$, **Case II** does not exist, and (6) is re-written as

$$q(n, k|N) = q(n-1, k-1|N) \left[\binom{N-k}{1} \alpha (1-\alpha)^{N-k} \right] \quad (7)$$

Based on (4), Figure 1 (a) illustrates the effect of α on $p(n|N)$. The figure indicates that the variance of $p(n|N)$ is large for a small α. Figure 1 (b) illustrates

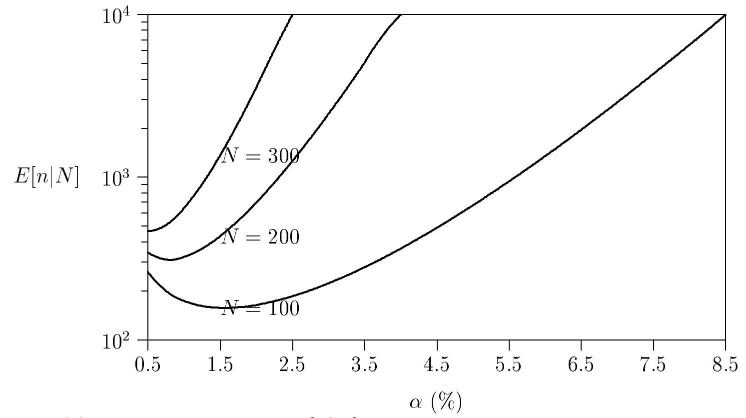

(a) The expected numer $E[n|N]$ of timeslots until an MS successfully re-registers

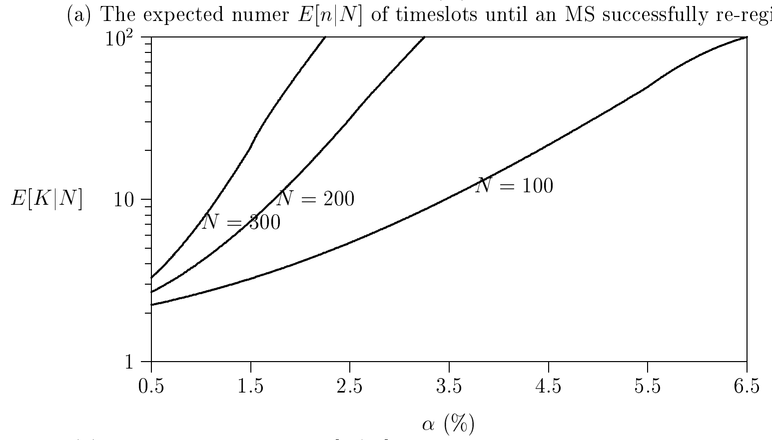

(b) The expected number $E[K|N]$ of re-tries until an MS successfully re-registers

Figure 2. The expected numbers $E[n|N]$ and $E[K|N]$

the effect of N on $p(n|N)$. The figure indicates that the variance of $p(n|N)$ is large for a large N. Figure 2 plots the expected number $E[n|N] = \sum_{n=1}^{\infty} np(n|N)$ of timeslots until an MS has successfully re-registered. The figure shows that $E[n|N]$ increases as N increases. This result is intuitive: if the number N of MSs in the VLR area is large, more collisions will occur before a specific MS has successfully re-registered. Another observation is that, as α increases, $E[n|N]$ first decreases and then increases. A small α implies that an MS will wait for a long time before it re-issues a re-registration message. Thus, as α increases, the re-try frequency increases, and $E[n|N]$ decreases. However, after some threshold value, a large α implies that more MSs will compete for the same timeslots and, thus, the number of re-collisions increases. Eventually, $E[n|N]$ increases as α becomes large. Let $E[K|N]$ be the number of re-tries until an MS has successfully re-registered (including the last successful try). Since the number of timeslots between two consecutive re-tries of an MS is $1/\alpha$ (the mean of the geometric distribution), $E[K|N]$ can be expressed as

$$E[K|N] = \frac{E[n|N]}{(1/\alpha)} + 1 = \alpha E[n|N] + 1$$

Figure 2 (b) plots $E[K|N]$ against α.

2.2. Effectiveness of Demand Re-registration

Assuming that the normal registrations of an MS form a Poisson process, the period between the VLR failure and the next registration has an exponential distribution. Since the events in a mobile system occur at timeslots, this period can be represented by a geometric distribution with probabilities γ_1. Similarly, if the call originations and call terminations are Poisson processes, then the times for the first call origination and the first call termination can be represented by geometric distributions with probabilities γ_2 and β, respectively. For our purpose, we aggregate the normal registration traffic and the call origination traffic to form a geometric distribution with probability $\gamma = \gamma_1 + \gamma_2$. The values of N, γ, β can be obtained from OA&M [8] of a mobile system. Suppose that after a VLR failure, the MS re-registration occurs at time τ_1, the first MS normal registration or call origination occurs at time τ_2, and the first MS call termination occurs at time τ_3. Then the expensive paging operation for the MS call termination (described in Appendix B) can be avoided if $\min(\tau_1, \tau_2) < \tau_3$. Note that if $\tau_2 < \tau_3$ is always true, then the expensive paging operation can be saved without demand re-registration. In other words, demand re-registration is not *effective* in this case. On the other hand, demand re-registration is *effective* if $\tau_1 < \tau_3 \leq \tau_2$, and the contribution (of saving expensive paging operation) of demand re-registration can be defined as

$$\theta_d = \frac{p_{1,d}}{p_{2,d}} = \frac{\Pr[\tau_1 < \tau_3 \leq \tau_2]}{\Pr[\min(\tau_1, \tau_2) < \tau_3]} \tag{8}$$

where

$$\begin{aligned} p_{1,d} &= \Pr[\tau_1 < \tau_3 \leq \tau_2] \\ &= \sum_{n=1}^{\infty} p(n|N) \left\{ \sum_{j=n+1}^{\infty} \beta(1-\beta)^{j-1} \left[\sum_{i=j}^{\infty} \gamma(1-\gamma)^{i-1} \right] \right\} \end{aligned}$$

$$= \sum_{n=1}^{\infty} p(n|N) \left[\frac{\beta(1-\beta)^n(1-\gamma)^n}{\beta+\gamma-\beta\gamma} \right] \tag{9}$$

and

$$\begin{aligned} p_{2,d} &= \Pr[\min(\tau_1, \tau_2) < \tau_3] \\ &= 1 - \Pr[\tau_3 \le \tau_1, \tau_3 \le \tau_2] \\ &= 1 - \sum_{n=1}^{\infty} p(n|N) \left\{ \sum_{j=1}^{n} \beta(1-\beta)^{j-1} \left[\sum_{i=j}^{\infty} \gamma(1-\gamma)^{i-1} \right] \right\} \\ &= 1 - \sum_{n=1}^{\infty} p(n|N) \left\{ \frac{\beta[1-(1-\beta)^n(1-\gamma)^n]}{\beta+\gamma-\beta\gamma} \right\} \end{aligned} \tag{10}$$

3. Performance of Demand Re-registration

Based on (9), (10) and (8), we illustrate the effect of demand re-registration. We made the following observations:

Observation 1. Effect of α: Figures 3-5 indicate that both $p_{1,d}$ and θ_d first increase and then decrease as α (the probability that an MS issues re-registration at a timeslot) increases. When α is very small, an MS does not issue/re-issue a re-registration message until a long waiting period has elapsed. It is likely that a normal registration, call origination, or call termination occurs before a successful demand re-registration operation is performed, which results in poor θ_d performance. This situation is improved as α increases. For a large α, increasing α implies a large number of re-try, which results in many collisions and thus poor θ_d performance. In other words, when α is large, θ_d decreases as α increases.

Observation 2. Effect of N: Figure 3 plots the effect of N on $p_{1,d}$ and θ_d. It is clear that the advantage of demand re-registration becomes insignificant as N increases. The number of MSs in a typical VLR is on the order of 10^4. In such a case, re-collisions of re-registrations may be significant. A simple approach can be used to limit the number of MSs (N) that need to re-register. The idea is to keep a backup of the VLR in a non-volatile storage. When the VLR is in normal operation, we may count the number of MSs entering the VLR. When the count is larger than N, these N new records are checkpointed into the backup. Suppose that a VLR failure occurs at time t_f, and that the last checkpointing occurs at time $t_c \le t_f$. Then after a VLR failure,

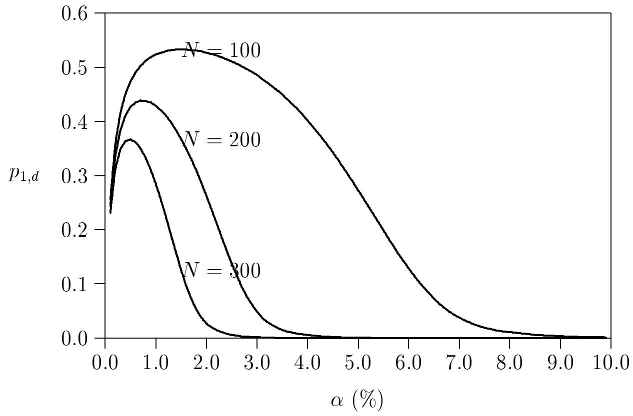

(a) The probability $p_{1,d} = \Pr[\tau_1 < \tau_3 \leq \tau_2]$

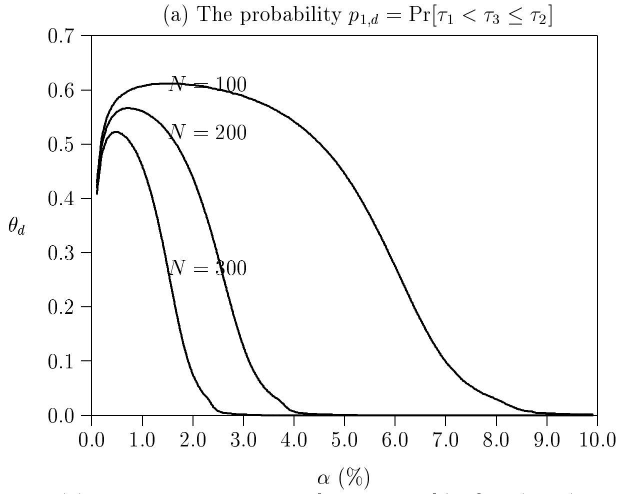

(b) The probability $\theta_d = \Pr[\tau_1 < \tau_3 \leq \tau_2] / \Pr[\min(\tau_1, \tau_2) < \tau_3]$

Figure 3. The effect of N on θ_d ($\beta = 0.0005, \gamma = 0.001$)

most MS records can be restored from the backup. The backup keeps correct locations for MSs entering the VLR before time t_c. The missing information is of those MSs entering the VLR during $(t_c, t_f]$. Thus, when the VLR broadcasts the re-registration request, the timestamp t_c can be included in the broadcast message and only MSs entering the VLR after t_c will re-register. It is clear that the number of re-registering MSs is limited to N.

Observation 3. Effects of β and γ: Since β represents the rate of call termination, it is clear that θ_d decreases as β increases (see Figure 4). The parameter γ equals to the combined rates of normal registrations and call originations. It

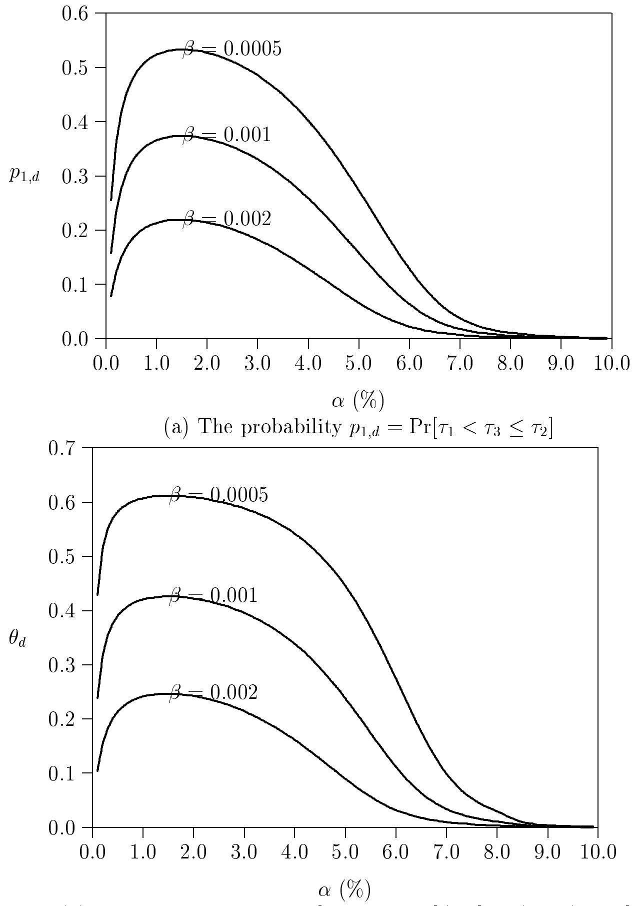

(a) The probability $p_{1,d} = \Pr[\tau_1 < \tau_3 \leq \tau_2]$

(b) The probability $\theta_d = \Pr[\tau_1 < \tau_3 \leq \tau_2]/\Pr[\min(\tau_1, \tau_2) < \tau_3]$

Figure 4. The effect of β on θ_d ($N = 100, \gamma = 0.001$)

is apparent that when γ is large, the first MS event after VLR failure is less likely to be a call termination, and the effectiveness of demand re-registration becomes remote. Thus, Figure 5 shows that θ_d increases as γ decreases.

4. Demand Re-registration vs Periodic Re-registration

Periodic re-registration has been studied in [5] where τ_2 (the period between the VLR failure and the next normal registration or the call origination) and τ_3 (the period between the VLR failure and the next call termination) are represented by continuous distributions. Furthermore, the number of the (expensive)

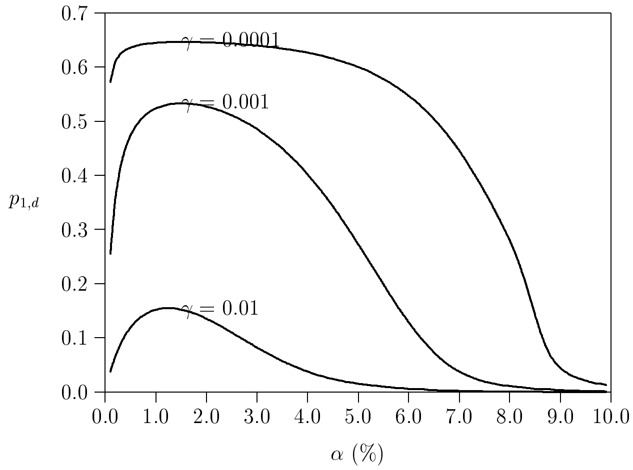

(a) The probability $p_{1,d} = \Pr[\tau_1 < \tau_3 \leq \tau_2]$

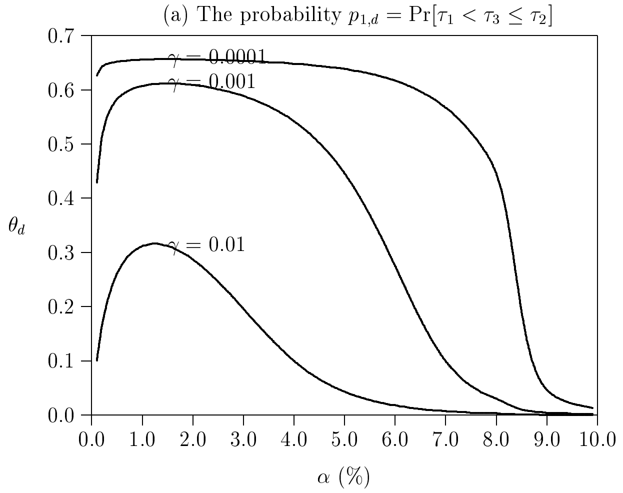

(b) The probability $\theta_d = \Pr[\tau_1 < \tau_3 \leq \tau_2]/\Pr[\min(\tau_1, \tau_2) < \tau_3]$

Figure 5. The effect of γ on θ_d $(N = 100, \beta = 0.5\gamma)$

call deliveries during the period between VLR failure and the recovery of the MS's VLR record was derived in [5], and is out of the scope of this paper. In order to compare periodic re-registration with demand re-registration in this paper, τ_2 and τ_3 are represented by geometric distributions with probabilities γ and β, respectively.

In periodic re-registration, an MS re-registers for every n_p timeslots. Let τ_p be the period between the VLR failure and the next periodic re-registration (measured in timeslots). From [5], τ_p is uniformly distributed in $[1, n_p]$. Similarly to the derivations of $p_{1,d}$ (see (9)) and $p_{2,d}$ (see (10)) for demand re-registration, we derive these probabilities for periodic re-registration as follows:

$$\begin{aligned} p_{1,p} &= \Pr[\tau_p < \tau_3 \le \tau_2] \\ &= \sum_{n=1}^{n_p} \left(\frac{1}{n_p}\right) \left\{ \sum_{j=n+1}^{\infty} \beta(1-\beta)^{j-1} \left[\sum_{i=j}^{\infty} \gamma(1-\gamma)^{i-1} \right] \right\} \\ &= \frac{\beta(1-\beta)(1-\gamma)[1-(1-\beta)^{n_p}(1-\gamma)^{n_p}]}{n_p(\beta+\gamma-\beta\gamma)^2} \end{aligned} \tag{11}$$

and

$$\begin{aligned} p_{2,p} &= \Pr[\min(\tau_p, \tau_2) < \tau_3] \\ &= 1 - \sum_{n=1}^{n_p} \left(\frac{1}{n_p}\right) \left\{ \sum_{j=1}^{n} \beta(1-\beta)^{j-1} \left[\sum_{i=j}^{\infty} \gamma(1-\gamma)^{i-1} \right] \right\} \\ &= 1 - \left[\frac{\beta}{n_p(\beta+\gamma-\beta\gamma)} \right] \\ &\quad \times \left\{ n_p - \frac{(1-\beta)(1-\gamma)[1-(1-\beta)^{n_p}(1-\gamma)^{n_p}]}{\beta+\gamma-\beta\gamma} \right\} \end{aligned} \tag{12}$$

Let θ_p be the periodic re-registration's contribution of saving expensive paging operation, then

$$\theta_p = \frac{p_{1,p}}{p_{2,p}}$$

Based on (11), (12) and (8), Figure 6 plots $p_{1,p}$ and θ_p against β and n_p. It is trivial that good θ_p performance is expected if the frequency of periodic re-registration is high (i.e., n_p is small). The figure also indicates that just like demand re-registration, θ_p decreases as β increases.

Let T_F be the interval between two VLR failures. Suppose that the size of the VLR (the average number of MSs in the VLR) is M. We utilize the checkpointing mechanism described in **Observation 2** so that the number of the re-registering MSs is limited to N. If we assume that the arrivals of MS into the VLR form a Poisson process with rate λ, then the period T_c between two checkpoints has an Erlang distribution with mean

$$E[T_c] = \frac{N}{\lambda} \ . \tag{13}$$

Let C_1 be the cost of storing a VLR record into the backup of the VLR database, C_2 be the cost of radio access for a registration operation, and C_3 be the cost of a registration. Since the registration operation consists of accessing the radio

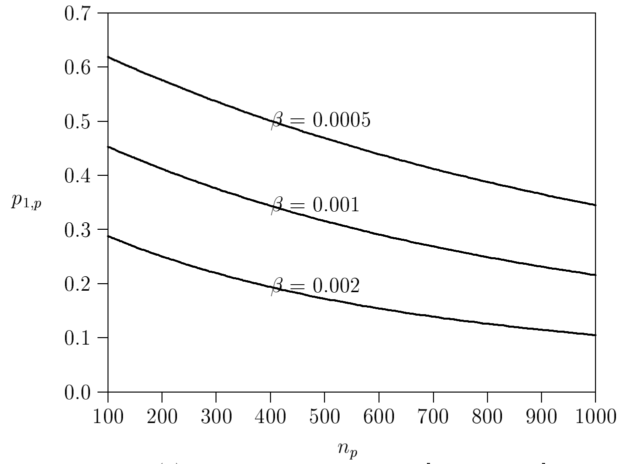

(a) The probability $p_{1,p} = \Pr[\tau_p < \tau_3 \le \tau_2]$

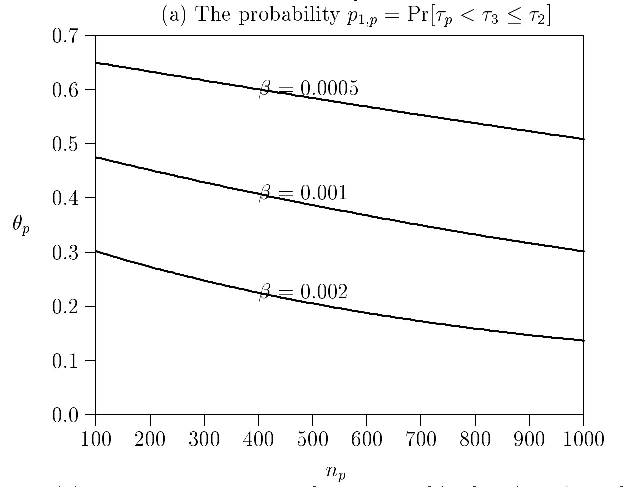

(b) The probability $\theta_p = \Pr[\tau_p < \tau_3 \le \tau_2] / \Pr[\min(\tau_p, \tau_2) < \tau_3]$

Figure 6. The Periodic Re-registration Performance ($\gamma = 0.001$)

channel (with cost C_2) and then update to the VLR database (with cost C_1), it is reasonable to assume that

$$C_3 = C_1 + C_2$$

The net cost C_d of demand re-registration for every VLR failure (the net cost of all operations performed during T_F, which are required to support demand re-registration) is

$$C_d \le C_d^*$$

$$= \left(\frac{E[T_F]}{E[T_c]}\right) NC_1 + MC_1 + N[(E[K|N] - 1)C_2 + C_3] \tag{14}$$

$$= (E[T_F]\lambda + M + N)C_1 + NE[K|N]C_2 \tag{15}$$

From (13), (14) is re-written as (15). Equation (14) consists of three components:

- During T_F, we expect $E[T_F]/E[T_c]$ checkpointing operations. Every operation writes N VLR records into the backup (with cost NC_1).
- After the VLR failure, M VLR records are restored from the backup to the VLR database (with cost MC_1).
- Finally, no more than N MSs are asked to re-register. Every of these MSs retries $E[K|N] - 1$ times (with cost $(E[K|N] - 1)C_2$) before it has successfully re-registered (with cost $C_3 = C_1 + C_2$).

Similarly, the net cost C_p of periodic re-registration is

$$C_p = M\left(\frac{E[T_F]}{n_p}\right) C_3 = \left(\frac{ME[T_F]}{n_p}\right)(C_1 + C_2) \tag{16}$$

From (15) and (16), it is clear that demand re-registration outperforms periodic re-registration if α is appropriately selected and if M is sufficiently larger than N. For example, assume that $\gamma = 0.001, \beta = 0.0005$. Suppose that $N = 100$ and $\alpha = 0.014$ are selected in demand re-registration. From Figure 3 (b), $\theta_d = 0.612$. From Figure 2 (b), $E[K|N] = 3.251$. To achieve the same θ_p performance (i.e., $\theta_p = 0.612$), $n_d = 330$ should be selected in periodic re-registration (see Figure 6 (b)). Suppose that the typical size of a VLR is $M = 10000$, and we assume that there are 1000 call deliveries to an MS between two VLR failures (this assumption favors periodic re-registration; the VLR failures should be much less frequent in reality), and $E[T_F] = 1000/\gamma = 1,000,000$. Furthermore, assume that $\lambda = 1000\gamma$. In other words, if an MS receives a call every hour, then there are 1000 MSs entering the VLR per hour. This assumption implies a large MS traffic, which favors periodic re-registration. From the above assumptions and from (15) and (16), we have

$$C_d \le C_d^* = 1,010,100C_1 + 225.1C_2 << C_p = 30,300,000(C_1 + C_2)$$

It is clear that demand re-registration outperforms periodic re-registration.

5. Conclusion

This paper proposed a new VLR failure recovery scheme called *demand re-registration*. We studied the performance of the demand re-registration scheme by investigating how effectively the re-registration can recover the location record for an MS before the first MS call termination occurs. Let θ_d be our scheme's contribution of saving the expensive paging operation. Our study demonstrated that demand re-registration may effectively recover VLR failure and the following results were observed:

- θ_d first increase and then decrease as α (the probability that an MS issues re-registration at a timeslot) increases.
- The advantage of demand re-registration becomes insignificant as N (the number of MSs that need to re-register) increases. In this paper, we proposed a simple approach to limit the number N of MS.
- θ_d decreases as β (the arrival rate of call termination) increases.
- θ_d increases as γ (the arrival rate of normal registration and call origination) decreases.

We have also compared demand re-registration with periodic re-registration. The results indicated that demand re-registration outperforms periodic re-registration.

References

[1] Chang, M.-F., Lin, Y.-B., and Su, S.-C. Improving Fault Tolerance of GSM Network. *IEEE Network*, 1(12):58–63, 1998.

[2] Cho, Y.-J., Lin, Y.-B., and Rao, C.-H. Reducing the Network Cost of Call Delivery to GSM Roamers. *IEEE Network Magazine*, 11(5):19–25, September/October 1997.

[3] ETSI/TC. Numbering, Addressing And identification. Technical Report Recommendation GSM 03.03, ETSI, 1993.

[4] ETSI/TC. Mobile Application Part (MAP) Specification, Version 4.8.0. Technical Report Recommendation GSM 09.02, ETSI, 1994.

[5] Haas, Z. J., and Lin, Y.-B. On Optimizing the Location Update Costs in the Presence of Database Failures. To appear in *ACM/Baltzer Wireless Network Journal*, 1998.

[6] Lin, Y.-B. Failure Restoration of Mobility Databases for Personal Communication Networks. *ACM-Baltzer Journal of Wireless Networks*, 1:365–372, 1995.

[7] Lin, Y.-B. Mobility Management for Cellular Telephony Networks. *IEEE Parallel & Distributed Technology*, 4(4):65–73, November 1996.

[8] Lin, Y.-B. OA&M for GSM Network. *IEEE Network Magazine*, March 1997.

[9] Lin, Y.-B. And DeVries S.K. PCS Network Signaling Using SS7. *IEEE Personal Commun. Mag.*, pages 44–55, June 1995.

[10] Mouly, M. And Pautet, M.-B. *The GSM System for Mobile Communications.* M.Mouly, 49 rue Louise Bruneau, Palaiseau, France, 1992.

Appendix

A. GSM Mobility Management Architecture

This appendix re-iterates GSM mobility management architecture and procedures given in [1] for the reader's benefit.

To exercise location tracking, a PCS service area is partitioned into several *location areas* (LAs). Every LA consists of a group of *base stations* that communicate with the MSs through radio contact. The major task of mobility management is to update the location of an MS when it moves from one LA to another. The location update procedure is referred to as *registration*, which is initiated by the MS as follows. The base stations continuously broadcast the corresponding LA addresses to the MSs. When an MS receives a different LA address, it sends a registration message to the network. The location information is stored in the PCS mobility databases called the *Home Location Register* (HLR) and the *Visitor Location Register* (VLR). For every LA, there is a corresponding VLR. When an MS visits the LA, a temporary record of the MS is created in the VLR to indicate its location. For every MS, there is a permanent record stored in the HLR, indicating the current MS location; i.e., the record stores the address of the VLR visited by the MS. We use GSM [4] as an example to illustrate the network architecture for PCS mobility management (see Figure 7). In this architecture, the base stations of a LA are connected to a *mobile switching center* or MSC (a telephone switch tailored for PCS services). Thus, an MSC covers several LAs. One or more MSCs are connected to a VLR, which exchange the location information with the VLR through the *Signaling System Number 7* (SS7) network [9]. Similarly, the VLR communicates with the HLR to exchange the location information using the SS7 messages. The GSM mobility management protocol (called *Mobile Application Part* or MAP) is implemented using the SS7 platform.

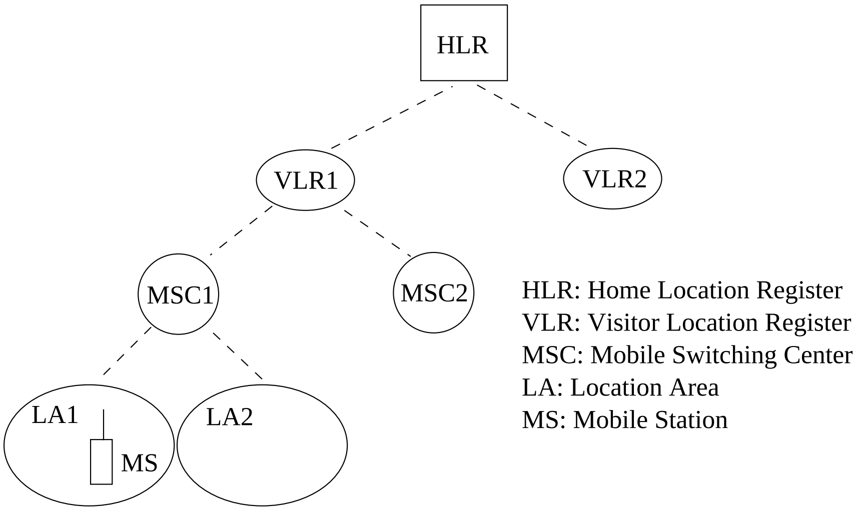

Figure 7. The PCS mobility management architecture

B. VLR Failure Restoration Procedures

For the reader's benefit, this appendix re-iterates GSM VLR failure restoration procedures given in [1]. In the VLR database, a record consists of three parts:

Mobile Station Information includes IMSI (International Mobile Subscriber Identity), MSISDN (Mobile Station ISDN Number) and TMSI (Temporary Mobile Subscriber Identity) defined in [3].

Location Information includes MSC number and the *Location Area Identity* (LAI). This information indicates the MSC and the LA where the MS resides.

Service Information includes a subset of the service information in the HLR. The VLR obtains this information from the HLR during the MS registration operation.

After a failure, the service information of a VLR record is recovered by the first contact between the VLR and the HLR (of the corresponding MS). The location information is recovered by the first radio contact between the VLR and the MS. The mobile station information is recovered either by the contact with the HLR or the MS. The VLR record restoration is initiated by one of the following three events.

MS Registration. Since the VLR record has been erased after the failure, the VLR assumes that the registration is for inter-VLR movement (note that the

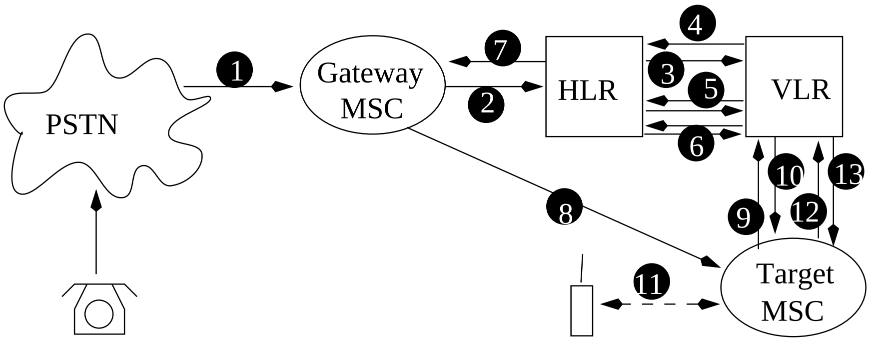

Figure 8. Call Termination Message Flow

MS may move between two LAs within an MSC, which is referred to as an intra-MSC movement). Following the normal registration procedure [2], the VLR record is recovered.

MS Call Origination. When the VLR receives the call origination request from the MSC [4], the VLR record for the MS is not found. The VLR considers the situation as a system error (the error cause is "Unidentified Subscriber"). The request is rejected, and the MS is asked to initiate the location registration procedure. After the registration, the VLR record is recovered.

MS Call Termination. The call termination message flow is illustrated in Figure 8.

Step 1. When the MSISDN is dialed by a user of the *Public Switched Telephone Network* (PSTN), the call is routed from the originating switch in the PSTN to a gateway MSC (or an ISDN exchange) by an SS7 ISUP (*ISDN User Part*) [9] IAM (*Initial Address Message*) message. The IAM message reserves the voice trunk between the originating switch and the gateway MSC.

Step 2. To obtain the routing information (i.e., the actual location of the MS), the gateway MSC interrogates the HLR by sending a routing query message. The message consists of the MSISDN of the MS (i.e, the telephone number of the MS) and other related information

Step 3. The HLR sends a query message to the VLR to obtain the MSRN (*Mobile Station Routing Number*) that indicates the MSC and the LAI of the MS. The message consists of the IMSI, the MSC number, and other related information. Note that the MSC number was maintained in both the HLR and the VLR.

The VLR searches the MS record by using the received IMSI. Since the record has been erased after the VLR failure, the search fails. The VLR creates a VLR record for the MS (note that both the service and the location information are not available in this newly created record). Then Steps 4 and 5 are executed in parallel.

Steps 4 and 7. Since the VLR does not have the routing information, it uses the MSC number provided by the HLR to create the MSRN. The number is sent back to the gateway MSC to set up the call in Step 8.

Steps 5 and 6. The VLR recovers the service information of the VLR record by exchanging a data restoration message pair with the HLR. At this point, the service information of the VLR record has been recovered. However, the location information (specifically, the LAI number) is still not available. This information will be recovered at Step 11 below.

Step 8. After the gateway MSC has received the MSRN in Step 7, an IAM message is sent to the target MSC to reserve the voice trunk between the gateway MSC and the target MSC.

Steps 9 and 10. The target MSC does not has the LA information of the MS. In order to proceed with the call setup procedure, the MSC sends a query in Step 9 to the VLR to find out the location area of the MS. Unfortunately, the VLR does not have the LAI information due to the database failure. The VLR asks the MSC to search the LA of the MS in Step 10.

Steps 11, 12 and 13. The MSC initiates paging of the MS in all LAs (Step 11). If the paging operation is successful, the current LA address of the MS is sent back from the MSC to the VLR in Step 12. At this point the location information of the VLR record is recovered. The VLR acknowledges this restoration operation in Step 13.

Note that the LA searching operation (i.e., Step 11) is an expensive operation, as every base station connected to the MSC must perform the paging operation.

Author Biographies

Zygmunt Haas received his B.Sc. in EE in 1979 and M.Sc. in EE in 1985. In 1988, he earned his Ph.D. from Stanford University and subsequently joined AT&T Bell Laboratories in Holmdel NJ as a Member of Technical Staff in Network Research Department. There he pursued research on wireless communications, mobility management, fast protocols, optical networks, and optical switching. From September 1994 till July 1995, Dr. Haas worked for the AT&T Wireless Center of Excellence in Whippany NJ, where he investigated various aspects of wireless and mobile networking, concentrating on TCP/IP networks. As of August 1995, he joined the Electrical Engineering faculty at Cornell University, currently as a tenured Associate Professor.

Dr. Haas is an author of numerous technical papers and holds twelve patents in the field of high-speed networking, wireless networks, and optical switching. He has organized several Workshops, delivered tutorials at major IEEE conferences, and serves as editor of several journals. He also was a guest editor of three IEEE JSAC issues ("Gigabit Networks," "Mobile Computing Networks," and "Ad-Hoc Networks"). Dr. Haas is a Senior Member of IEEE, a voting member of ACM, and a Vice Chair of the IEEE Technical Committee on Personal Communications. His interests include: mobile and wireless communication and networks, personal communication service, high-speed communication and protocols, and optical switching. His e-mail address is `haas@ee.cornell.edu` .

Yi-Bing Lin received his BSEE degree from National Cheng Kung University in 1983, and his Ph.D. degree in Computer Science from the University of Washington in 1990. From 1990 to 1995, he was with the Applied Research Area at Bell Communications Research (Bellcore), Morristown, NJ. In 1995, he was appointed as a professor of Department of Computer Science and Information Engineering (CSIE), National Chiao Tung University (NCTU). In 1996, he was appointed as Deputy Director of Microelectronics and Information Systems Research Center, NCTU. Since 1997, he has been elected as Chairman of CSIE, NCTU. His current research interests include design and analysis of personal communications services network, mobile computing, distributed simulation, and performance modeling.

Dr. Lin is an associate editor of IEEE Network, an associate editor of SIMULATION magazine, an area editor of ACM Mobile Computing and Communication Review, a columnist of ACM Simulation Digest, a member of the editorial board of International Journal of Communications Systems, a member of the editorial board of ACM/Baltzer Wireless Networks, a member of the editorial board of Computer Simulation Modeling and Analysis, an editor of Journal of Information Science and Engineering, Program Chair for the 8th Workshop on Distributed and Parallel Simulation, General Chair for the 9th Workshop on Distributed and Parallel Simulation. Program Chair for the 2nd International Mobile Computing Conference, the publicity chair of ACM Sigmobile, Guest Editor for the ACM/Baltzer MONET special issue on Personal Communications, and Guest Editor for IEEE Transactions on Computers special issue on Mobile Computing. Lin received 1997 Outstanding Research Award from National Science Council, ROC, and Outstanding Youth Electrical Engineer Award from CIEE, ROC. He is a senior member of IEEE. Lin's email address is `liny@csie.nctu.edu.tw` .