Distributed Paging and Registration in Wireless Networks

Suman Das, Thierry E. Klein, Sayandev Mukherjee, Gee Rittenhouse, Louis G. Samuel, and Harish Viswanathan, Bell Laboratories — Lucent Technologies
Kin K. Leung, Imperial College, London
Haitao Zheng, Microsoft Research, Beijing

Abstract

The success of current and future wireless networks depends on their ability to provide connections to mobile terminals anywhere and at any time. It is therefore of crucial importance that wireless networks are able to quickly and efficiently locate mobile users at the time of an incoming call, which is achieved in current networks through a combination of paging and registration. In this article we present novel distributed paging and registration procedures that are naturally suited to future distributed wireless network architectures. In addition, the distributed nature of these strategies allows us to effectively balance the paging and registration traffic and the required signal processing throughout the network and alleviate any potential overloads of individual base stations.

The popularity and success of current and future wireless networks depend on their ability to provide reliable communication to mobile users at any time and anywhere. In order to provide this universal connectivity, the network has to be able to maintain and establish a connection path from source to destination at any given time, independent of the mobile's geographical location. This problem arises specifically in wireless networks due to the user mobility and the fact that there is no single physical point of attachment of the mobile to the network (as is the case in wireline networks, for example).

This broad problem of location management encompasses several aspects. First, during an active call, the network has to be able to maintain the connection even when the end user moves in the network. This aspect is usually solved by *handoff* schemes that connect the user to a different base station if the connection to that base station becomes stronger than the one to the previously chosen base station. The second aspect of location management deals with *radiolocation*, which means pinpointing the exact geographical location (as opposed to merely the closest base station) of the user. Such techniques are essential for emergency services, and vehicle and people tracking, and are not discussed in any detail in this work. The third and final aspect of location management is to be able to establish a new connection to a mobile terminal when a call destined for that mobile is initiated in the network. The main difficulty stems from the fact that the mobile may have moved from the last known location, and therefore could potentially be anywhere in the network. This problem is essentially solved by the *paging* and *registration* procedures, which work in conjunction to determine the closest base station to the mobile at the time when a connection needs to be established.

In the registration (or update) procedure, the mobile is required to send registration messages to inform the network of its location. In this article we only consider the relevant registration procedure when the mobile is powered up but not active (i.e., when the mobile is not engaged in an ongoing call). Indeed, when the mobile is active, the network is able to track the mobile through micro- and macromobility management procedures (e.g., the handoff strategies in the traditional cellular networks, or Mobile IP and Hierarchical Mobile IP in all-IP wireless networks). The registration messages could be sent periodically after expiration of a timer, whenever the mobile moves to particular locations or regions in the network, or even when the mobile has traveled a distance from the location where the last registration occurred exceeding a certain threshold. The reader is referred to [1] for a discussion of the various registration procedures. Of course, more elaborate procedures than the ones mentioned here are also possible and have been studied in the literature [2–4, references therein]. Some of these procedures are naturally distributed in nature, but to our knowledge have not been implemented in current networks.

The second important procedure is the paging procedure by which the network pages all base stations in a particular region in order to locate the mobile terminal. Once a base station receives a paging request from the network, it sends a paging message over its paging channel with the unique identifier for the mobile that needs to be located. Mobiles that are powered up are required to periodically monitor the paging channels and respond to paging messages with their identifier. Of course, it goes without saying that when mobile terminals are not powered up, it is impossible to locate them, resulting in an unsuccessful connection. The paging procedure requires a list (called the *location area*) of possible base stations at

This research was conducted when K. Leung and H. Zheng were with Bell Laboratories — Lucent Technologies.

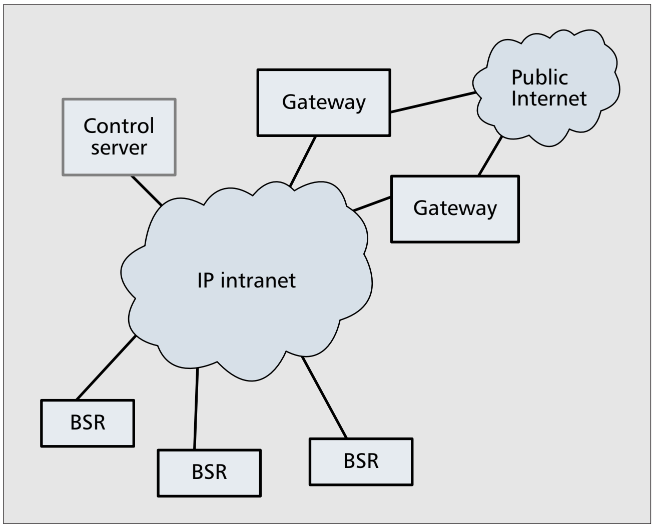

■ Figure 1. *Generic all-IP network architecture with base station routers.*

which the mobile may be located (this list may potentially even include all the base stations in the network). Paging procedures differ by the order in which the various base stations in the location area are paged (e.g., simultaneous or sequential paging of the base stations). The order in which the base stations are paged is part of the design of the paging strategy and depends on such parameters as mobile speed and direction of movement, call arrival statistics, and any a priori information about the mobile's possible location.

Two basic strategies for paging and registration illustrate the fundamental trade-offs involved. The *always-update strategy* requires that a mobile terminal send an update message upon entering a new cell. In other words, whenever the mobile detects that its signal strength to some base station has become stronger than that to the current base station, it sends a registration message to inform the network of that change in relative signal strength. Of course, for such a strategy the paging cost is minimal, as the network is always aware of the base station with the best connection to the user. On the other hand, the registration cost (and the associated power consumption for the mobile and registration traffic and signal processing) could be very large, especially if the user is highly mobile. This is certainly not an acceptable strategy for high-velocity users.

At the other extreme, in the *never-update strategy* the terminal would never send any registration messages, thereby requiring network-wide paging. The registration cost is of course zero, but the paging delay and associated traffic in the backhaul network could become unacceptably large. Therefore, sensible strategies employ a compromise between these two extreme strategies. The main issues that influence this trade-off and determine the outcome of the compromise include the cost of registration and paging of a mobile, the dissemination, recording, and storage of the location information in the network, as well as the delay in finding a particular mobile and the probability and cost of an unsuccessful paging request. This trade-off is quite extensively studied in the literature and we refer the reader, for example, to [5–7].

As mobile cellular networks evolve to carry both traditional voice and new data services, an IP-based radio access network (RAN) akin to the Internet is likely to be the most favorable approach for future RANs. The advantages of an all-IP RAN include cost efficiency from economies of scale (since the radio network has the same network elements as the Internet, allowing for use of off-the-shelf components), improved reliability, separation of the control and transport planes, convenient implementation of new services, independent scaling of control and transport, and straightforward integration of multiple networks [8–10]. A typical all-IP network architecture is illustrated in Fig. 1. The base stations, which are now routers and called base station routers (BSRs), are connected to the service provider's intranet (also referred to as the backhaul network). A gateway connects the intranet to the Internet. IP is the network protocol transporting user and control information within the intranet. The control server provides the necessary call service control. An important characteristic of the network in Fig. 1 is that a lot of the radio network functionalities are integrated with the base station functionalities and thus distributed across the network. Unlike the traditional architecture, there is no node in the RAN that is wireless-aware. The distributed network architecture therefore provides natural motivation to consider distributed paging and registration procedures.

The remainder of the article is organized as follows. We first describe currently used centralized paging and registration procedures as they are applied in today's wireless networks. We also provide details of how centralized procedures could be applied in future distributed all-IP network architectures. We then present the details of our proposed distributed paging and registration procedures. Finally, we provide some calculations to substantiate the claims that distributed paging and registration lead to effective load balancing in the network.

Centralized Paging and Registration

We now proceed to describe the most commonly used paging and registration procedures in today's wireless networks, and their extensions to the new architecture proposed in this article. The typical scenario that is employed in currently deployed 2.5-generation (2.5G) and third-generation (3G) networks is to define a location area to comprise a certain number of base stations. How to design the location area, and how many and which base stations should be part of a location area are left to the network designer and go beyond the scope of this discussion. In current networks the location areas are the same for all the mobile terminals in the network. Each terminal is required to register with the network as soon as it is powered up and again whenever it enters a new location area. By comparing the relative signal strengths of pilot signals sent by the base stations, the mobile is able to determine if and when it has moved within the vicinity of a new base station. If the mobile has information on the location areas (specifically which base stations belong to which location areas), it can determine that it has left a location area and entered a new location area, and initiate the corresponding registration message. Alternatively, if the mobile does not have location area information, the base stations are required to have such knowledge. Whenever the mobile's signal to a new base station becomes largest, the mobile registers with that base station by transmitting its mobile terminal identification number and the identification number of its previous associated base station. If the new base station determines that the mobile has moved across a location area boundary (after comparing the identification numbers with the location area information), either the mobile or the new base station, on behalf of the mobile, initiates a registration message to the network. The identification number of the new base station, as well as the corresponding information of the associated location area, is then stored in the location registry database.

In addition to the location-based registration procedure,

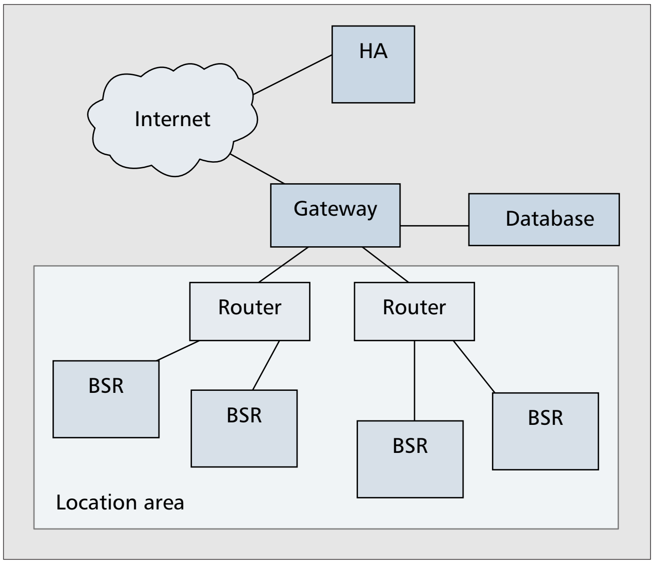

■ Figure 2. *Illustrative network configuration for distributed paging.*

mobile terminals are also required to periodically register with the network after expiration of a timer to inform the network that the mobile is still alive and that the corresponding call session state information should be maintained in the network. Furthermore, some wireless networks optionally require the mobile to register if it has traveled a minimum distance from the last known base station. We point out that the location-based registration is inherently a centralized procedure as mobiles only register with those base stations at the boundary of the location areas. In contrast, the timer and distance-based registration procedures are naturally distributed as they depend on the stochastic and independent mobility patterns of the different mobiles. However, even if the latter two registration procedures are distributed, the location information is stored in a central database that may be collocated with the radio network controller (RNC).

In currently employed location management strategies in 2.5G and 3G networks, paging of a user is restricted to the base stations in the last known location area. Typically, all the base stations within the location area are paged simultaneously. Alternatively, a particular paging order could be used if additional information on the user's location is available. Such information could, for example, be available if distance-based registration procedures are employed.

The traditional paging and registration procedures implemented by current wireless networks may very well be used in distributed network architectures, such as the one represented in Fig. 1. However, by virtue of its distributed nature, this network architecture naturally allows for additional flexibility to distribute the paging and registration functionalities, and the corresponding computational complexity and signaling load within the network. Before we proceed to describe the distributed procedures in later sections, we briefly summarize the centralized paging procedure as it would be applied in the context of a distributed architecture. We emphasize that the paging functionality is still centralized; only the radio functionalities are distributed in the network. Note that this discussion closely mirrors the procedures in current networks provided that the appropriate correspondences between network elements are considered.

In Fig. 2 we show an illustrative network configuration in which it is assumed that different BSRs are connected to a router and ultimately to a gateway, which serves as the control server depicted in Fig. 1 through a hierarchical architecture. For simplicity and to illustrate the main ideas, we consider that all the BSRs are connected to the gateway through a single router. Note that the location area may contain one or more routers and the associated BSRs. The gateway contains the location register database in which the relevant location information is stored. Note also that a gateway may be connected to more than one location area, although this is not shown in Fig. 2.

Traditionally the paging functionality resides at a central location such as the RNC. We now briefly describe the corresponding procedure, which is also given in the signal flow chart in Fig. 3. If a correspondent host in the core network initiates a call, the call is routed to the receiving terminal's home agent (HA) in its home network. Through an established binding, the HA then forwards the data packets to the terminal's foreign agent (FA), which then triggers a paging request to the centralized paging server that can reside in the gateway. In the traditional centralized architecture, the paging server would typically be located in the RNC. The paging server, upon probing the location register (LR) database, determines the last known location area and sends appropriate paging requests to all the base stations within said location area in order to locate the closest base station to the mobile terminal at that time. This could, for example, be done by a paging protocol with IP multicasting as a transport mechanism that multicasts the paging request to a subset (or all) of the base stations that are paged simultaneously. The base stations then send a layer 2 paging message to the mobile using the dedicated paging channels. The mobile responds by a page-response message to the base station from which it detected the paging request message. The base station that eventually locates the mobile terminal responds to the paging server with a layer 3 paging response message and possibly signals the mobile to perform a Mobile IP registration with the HA. Base stations that do not locate the mobile do not respond by explicit negative acknowledgment messages. At that point, the user terminal is deemed located and call setup may proceed. Note that in this scenario, both the paging functionality and LR are centralized and collocated with the gateway. We also emphasize that the paging functionality could reside in a dedicated paging server located somewhere in the network (instead of being located at the gateway). A crucial aspect of the centralized paging procedure in this architecture is that the gateway or, in general, the paging server has to be a wireless-specific and centralized network component. Therefore, it cannot simply be substituted for off-the-shelf wireline network components, entailing additional costs through specialization and lack of scalability.

Distributed Paging Procedure

In the context of future all-IP distributed network architectures, we continue to employ paging and registration procedures very similar to those employed in current systems and described earlier. However, we wish to distribute the functionalities in the network in order to avoid excessive processing at specific nodes and excessive signaling traffic on certain links in the network. In addition, the need for network entities with specific wireless-related functionalities is eliminated or at least reduced. More elaborate paging procedures can equally well be implemented in the context of this architecture.

In this scenario we continue to assume that all the base stations within the geographical region of interest are regrouped into contiguous location areas. A mobile terminal is again required to register upon entering a location area, but does not necessarily register when moving across base stations with-

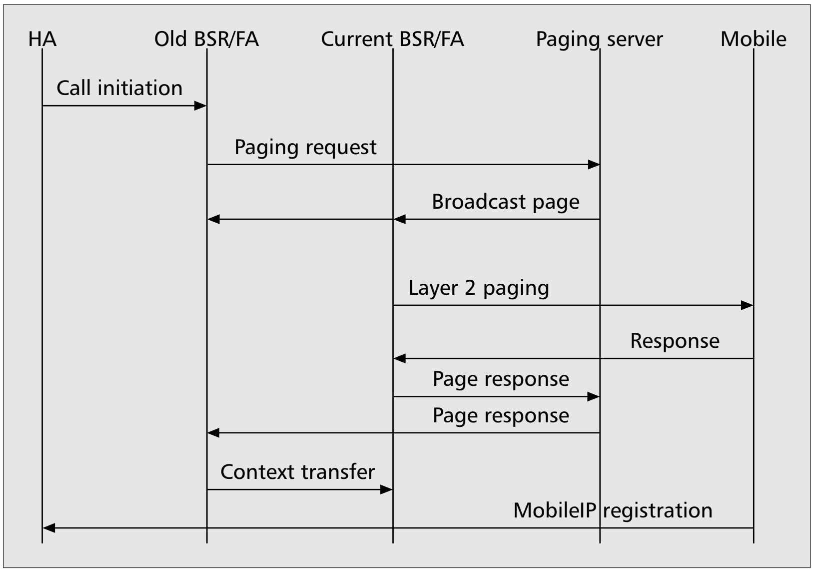

■ Figure 3. *Centralized paging signal flow.*

in the same location area. The determination of when a mobile enters a new location area can be made according to the same procedures as in the centralized architecture. Upon arrival of a new call for a particular user, the paging is then restricted to the last known location area.

In our enhanced distributed paging procedure, the LR database is still centralized, but may not have to be collocated with the gateway. However, the paging functionality is distributed and resides with the last base station BSR_{last} at which the mobile was located. BSR_{last} could be either the BSR with which the mobile registered upon first powering up in the network, or the BSR with which the mobile registered upon entering the current location area. However, if more sophisticated registration procedures are implemented in addition and on top of the location areas (e.g., timer- or distance-based procedures), BSR_{last} could equally well be any BSR in the location area.

We now describe the corresponding distributed paging procedure, for which a signal flow chart is shown in Fig. 4. The HA has a binding to the last registered FA that resides in the last known BSR, BSR_{last}, and, upon call arrival, the HA sends a paging request to BSR_{last}. The ensuing layer 2 paging procedure is unchanged from the centralized paging discussion. If BSR_{last} is unable to locate the mobile, rather than respond to a central paging server, BSR_{last} is capable of initiating paging requests to all other base stations within the same location area. The information of which base stations should be paged (i.e., the identities of all base stations within the same location area) could be stored in BSR_{last}, or alternatively obtained from the LR database. The latter is particularly attractive when location areas are not fixed and may depend on the particular user, as discussed below in the section on distributed registration. Upon locating the mobile the corresponding BSR, denoted BSR_{new}, sends a paging response message to BSR_{last}. BSR_{new} is then responsible for updating the LR database with the corresponding location information of the mobile and for signaling the mobile to perform a registration with the HA. Finally, BSR_{new} informs BSR_{last} that the mobile has been located and BSR_{last} may relinquish its paging responsibility for that mobile. Finally, BSR_{new} requests and initiates the migration of the call state and internal parameters relevant to the mobile from BSR_{last} to BSR_{new}. Note that the scheme can be extended so that the LR database is distributed and located with BSR_{last}. This would reduce the required signaling further as well as the delay in locating the mobile.

In summary, in this distributed paging procedure, the paging functionality that resides with a central entity (e.g., the RNC) in the centralized architecture is dynamically distributed in the network and resides with the last known base station for the paged mobile. Alternative approaches include choosing a specific BSR in each location area to handle all the paging functionalities for all the mobiles in that location area. However, such an approach would violate the objective that all BSRs should have the same functionalities and would therefore be cost-effective to manufacture. In addition, since the BSR at which a mobile was last registered depends on the particular mobile and changes over time, the paging load (including the signal processing and management of location information as well as the required signaling traffic in the network) can be more evenly distributed in the network by assigning the paging responsibility to the last known BSR.

Later, we provide some simple calculations to evaluate the traffic load and time required to locate a particular mobile terminal under both the centralized and the distributed paging procedures for the simple network in Fig. 2. These calculations show that the total amount of paging traffic and the total time required to locate a particular mobile are not substantially changed by the distributed paging procedure. Therefore, the complexity of the distributed and centralized paging procedures is comparable. However, while this is true for an individual user, the distributed paging procedure allows distributing the overall paging traffic (when multiple users have to be located) more evenly throughout the network. In addition, there is not a single point of failure in the network (as would be the case with a dedicated paging server), and the processing capabilities, required buffering, and traffic distribution are more evenly distributed in the network. By adapting the location area to each user, it is in fact possible, as discussed below, to balance the paging traffic in the network.

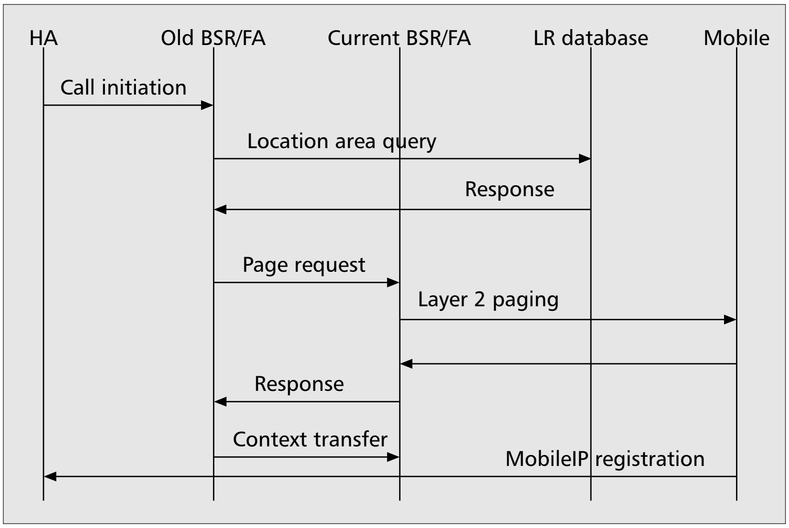

■ Figure 4. *Distributed paging signal flow.*

Distributed Registration Procedure

In this section we concentrate on enhanced distributed registration procedures. In current 2.5G and 3G networks the location areas are the same for all the users in the network and cover a certain set of base stations. This set of base stations remains the same for all users throughout the operation of the network. While such an assumption is certainly reasonable and provides the simplest solution, the calculations performed later show that it leads to an uneven distribution of the registration traffic load in the network. Specifically, only the fraction of base stations located at the boundary of the location area handle any registration traffic, while the base stations in the interior of the location area do not share in the burden.

The main idea behind distributed registration is to make sure that the registration traffic is distributed among all the base stations in the network by choosing different location areas for different users. Location areas essentially form a tiling of the geographical region, and different users could be associated with different shifted versions of the same base pattern of location areas. As an illustrative practical example, assume that the base stations in the network can be partitioned according to two different patterns of location areas, P_1 and P_2. Then each user could be assigned to either P_1 or P_2 depending, for example, on the parity of the user ID. In other words, users with even ID numbers are required to register once they cross location area boundaries as defined by P_1, whereas odd-numbered users register whenever they cross location area boundaries defined by P_2. More dynamic ways of assigning users to a different set of location areas could be determined when the user terminal is powered up and adjusted in an attempt to balance the registration traffic across the network. For example, a new user could be assigned to the set of location areas that has the fewest users assigned to it at that time. This procedure naturally leads to balancing of registration traffic across the available sets of location areas. It is quite obvious that the total amount of registration traffic in the given geographical region is not changed; however, it is no longer confined to a small fraction of the BSRs, but can be distributed in a fair and efficient manner throughout the network. The actual registration procedures are the same as those previously described and used in current wireless networks, except that the location area boundary and when the registration procedure is invoked are different for each user and can be dynamically adjusted.

Illustrative Calculations for Paging and Registration Traffic

In this section we provide some simple calculations to evaluate the paging traffic load B_{pag} and paging time T_{pag} required to locate a particular mobile terminal under both the centralized and distributed paging procedures. Finally, we derive an expression for the registration traffic B_{reg} as a function of the number of base stations and location areas in a given geographical region.

We assume for the sake of simplicity that the network configuration is the one in Fig. 2 and that the paging area (i.e., the location area) covers K routers. Further assume that N BSRs are attached to each router in a star topology. Let B_1 be the amount of bits in the paging message between the gateway (or LR database) and the router, B_2 the message size between the router and each BSR, B_3 the size of the page response message between the BSR and a router, and B_4 the message size between the router and the gateway. Of course, we would expect that $B_1 = B_4$ and that $B_2 = B_3$, but we do not make this assumption here for added generality. Similarly, we denote T_1, T_2, T_3, and T_4 the corresponding delays associated with the different messages, and the times required to process the paging and page response messages. Finally, T_p denotes the time required for layer 2 paging over the air interface.

First we concentrate on the centralized paging procedure and distinguish two possible paging scenarios. In the first scenario all the BSRs associated with all the routers are paged simultaneously. The corresponding time required to locate the mobile within the location area is then given by

$$T_{pag} = T_1 + T_2 + T_p + T_3 + T_4.$$

Similarly, the total amount of paging traffic is determined to be

$$B_{pag} = K[B_1 + NB_2] + B_3 + B_4.$$

In the second scenario the routers in the location area are paged successively until the mobile is located. Once the paging message reaches a router, all the BSRs associated with that router are paged simultaneously. Assuming that the mobile is uniformly located throughout the location area and no a priori information on the mobile's location is available, on average the procedure pages $K/2$ routers. Thus, the total average paging time is given by

$$T_{pag} = \frac{K}{2}\left[T_1 + T_2 + T_p\right] + T_3 + T_4.$$

Similarly, the average total paging traffic is determined to be

$$B_{pag} = \frac{K}{2}[B_1 + NB_2] + B_3 + B_4.$$

We note that the average total paging traffic is reduced if the routers are paged successively; however, the overall time to locate the mobile is correspondingly increased. Of course, both the paging traffic and the time required to locate the mobile could be further reduced if additional information (e.g., through a conditional probability density function, given the previous known location of the mobile and/or the distance traveled by the mobile) is available regarding the mobile's location.

Next we provide the corresponding calculations for the distributed paging procedure as described in this article. Through the concept of BSR_{last}, the paging functionality is effectively distributed among all the BSRs in the location area. Indeed, it can be anticipated that different BSRs are the last known BSRs for different mobiles at different times during the operation of the network. The consequence is that the paging traffic is more evenly distributed in the backhaul network. The calculations below, however, show that the total amount of traffic is not significantly reduced, although it is differently distributed throughout the network. For simplicity, we only consider the case when a single router is in the given location area. Our results can naturally be extended to the case of multiple routers in the same location area. Assuming that the mobile is equally likely to be located at any of the N BSRs in the location area, it is easily shown that the average paging time required to locate the mobile is given by

$$T_{pag} = T_1 + \left(2 - \frac{1}{N}\right)T_2 + \left(1 - \frac{1}{N}\right)T_p + \left(2 - \frac{1}{N}\right)T_3 + T_4.$$

Similarly, the average paging traffic in the distributed paging procedure is evaluated to be

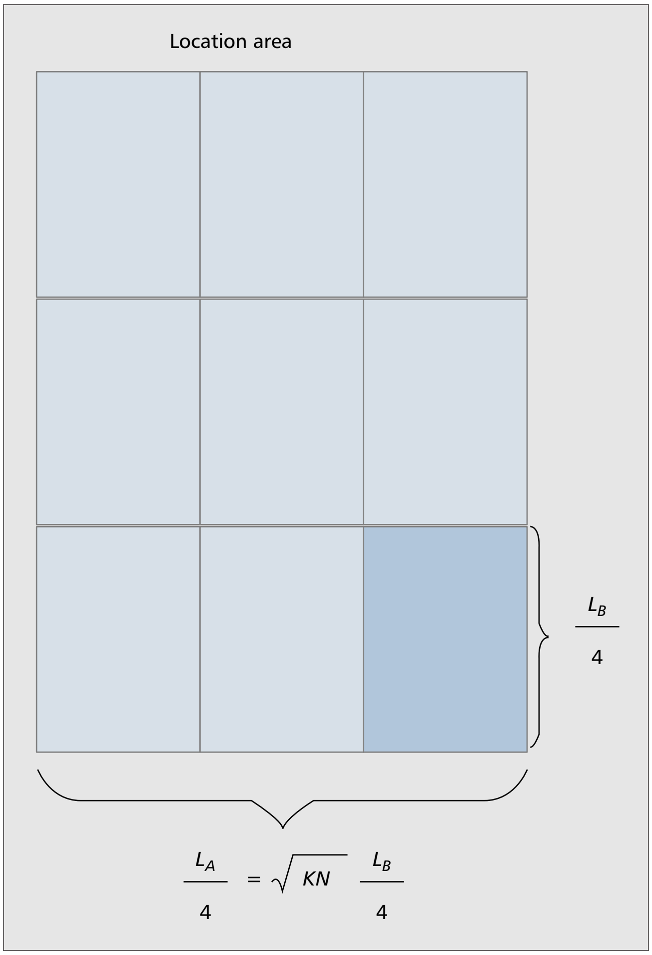

■ Figure 5. *Illustrative location area shown as a square region with* N *BSRs, each covering a square geographical region as well. The location area contains a total of* KN *BSRs.*

$$B_{pag} = B_1 + \left(N - 1 + \frac{1}{N}\right)B_2 + \left(2 - \frac{1}{N}\right)B_3 + B_4.$$

Next we turn our attention to the amount of registration traffic generated by the registration procedure described above. In these calculations, assuming that a mobile terminal registers upon entering a new location area, we show that the registration traffic is very unevenly distributed among the base stations in the location area. For simplicity and ease of exposition, assume that the area covered by a base station is a square region of perimeter L_B and that the location area covers K routers, each connected to N base stations (thus, the location area contains a total of KN base stations). Again, for simplicity we assume that the location area also forms a square region of perimeter L_A, as illustrated in Fig. 5. It is easily shown that the perimeter of the location area is then given by

$$L_A = \sqrt{KN}\,L_B.$$

Further assume that the entire geographical region of interest has M BSRs and therefore M/K distinct location areas.

Let ρ be the user density in the location area and V the average user velocity. Assuming that the users' direction of mobility is uniformly chosen and using a fluid flow mobility model, it can then be shown [11] that the rate of mobiles crossing the boundary of the location area is given by

$$R_{cross} = \frac{\rho V L_A}{\pi}.$$

Users are required to send registration messages whenever they enter (or leave) a location area. Let B_R be the size of the registration message. Then the total registration traffic in the geographical region of interest is determined as

$$B_{reg,total} = \frac{B_R \rho V M \sqrt{N} L_B}{\pi \sqrt{K}}.$$

Since users only register upon entering a new location area, it is evident that only the base stations whose cells cover the boundary of the location area handle registration traffic. In particular, all the cells in the "interior" of the location area do not share any burden of registration traffic, resulting in potential overloading of the boundary cells and uneven distribution of the signaling traffic in the backhaul network. In particular, out of the KN BSRs located in a given location area, only

$$4\left\lfloor \sqrt{KN} - 1 \right\rfloor$$

of them handle any registration traffic. Thus, the load of the registration procedures and associated signaling is very unevenly distributed in the network.

Therefore, the average (per base station) registration traffic handled by those base stations at the boundary of the location area is

$$B_{reg,BSR,cent} = \frac{B_{reg,total}}{4\left\lfloor \sqrt{KN} - 1 \right\rfloor} \approx \frac{B_R \rho V M L_B}{\pi} \frac{1}{4K},$$

where the last approximation is valid if the number of BSRs per location area, KN, is large with respect to 1. The other BSRs in the center of the location area do not carry any registration traffic. In contrast, in the distributed registration procedure, assuming that the assignment of location areas to mobiles is such that all BSRs share evenly in the registration traffic, the average per-BSR registration traffic is the same for all BSRs in the location area and is given by

$$B_{reg,BSR,dist} = \frac{B_{reg,total}}{KN} \approx \frac{B_R \rho V M L_B}{\pi} \frac{1}{K\sqrt{KN}}.$$

Conclusions

In this article we have described novel distributed paging and registration procedures for all-IP wireless networks. The main advantages of the proposed procedures are that they allow the paging functionality to be distributed throughout the network and thereby eliminate the need for any central wireless-specific network elements. The overall complexity and required signaling are about the same as in centralized procedures, but traffic is more evenly distributed in the network. On the other hand, the distributed registration procedure differs from the centralized procedure in that different users have different sets of location areas. Furthermore, the set of location areas can be dynamically assigned to users upon registration with the network.

References

[1] A. Bar-Noy, I. Kessler and M. Sidi, "Mobile Users: To Update or Not to Update?," *ACM Mobile Net. and Appls.*, vol. 1, no. 1, Feb. 1995, pp. 49–56.
[2] I. Akyildiz and J. Ho, "Dynamic Mobile User Location Update for Wireless PCS Networks," *ACM Wireless Networks*, vol. 1, no. 2, Feb. 1995, pp. 187–96.
[3] S. Sen, A. Bhattacharya, and S. Das, "A Selection Location Update Strategy for PCS Users," *ACM Wireless Networks*, vol. 5, no. 5, Oct. 1999, pp. 313–26.

[4] J. Ho and I. Akyildiz, "Mobile User Location Update and Paging under Delay Constraints," *ACM Wireless Networks*, vol. 1, no. 4, Feb. 1995, pp. 413–25.
[5] C. Rose and R. Yates, "Minimizing the Average Cost of Paging under Delay Constraints," *ACM Wireless Networks*, vol. 1, no. 2, Feb. 1995, pp. 211–19.
[6] C. Rose, "Minimizing the Average Cost of Paging and Registration: A Timer-Based Method," *ACM Wireless Networks*, vol. 2, no. 2, June 1996, pp. 109–16.
[7] U. Madhow, M. Honig, and K. Steiglitz, "Optimization of Wireless Resources for Personal Communications Mobility Tracking," *IEEE/ACM Trans. Net.*, vol. 3, no. 6, Dec. 1995, pp. 698–707.
[8] D. Wisely, P. Eardley, and L. Burness, *IP for 3G*, Wiley, 2002.
[9] S. Uskela, "Key Concepts for Evolution Toward Beyond 3G Networks," *IEEE Wireless Commun.*, vol. 10, no. 1, Feb. 2003, pp. 43–48.
[10] M. S. Corson *et al.*, "A New Paradigm for IP Based Cellular Networks," *IT Pro*, Nov.–Dec., 2001, pp. 20–29.
[11] R. Ramjee *et al.*, "HAWAII: A Domain-Based Approach for Supporting Mobility in Wide-area Wireless Networks," *IEEE/ACM Trans. Net.*, vol. 10, no. 3, 2002, pp. 396–410.

Biographies

SUMAN DAS (sumand@lucent.com) received his undergraduate degree in computer science and engineering from the Indian Institute of Technology, Kharagpur, in 1994, and his M.S. and Ph.D. degree from Rice University, Houston, Texas, in 1997 and 2000, respectively. Since then he has been a member of technical staff in the Wireless Technology Research Department, Bell Laboratories — Lucent Technologies, Murray Hill, New Jersey. His primary research interests are in algorithms and architectures for signal processing and communication, and design and analysis of future wireless networks.

THIERRY E. KLEIN (tek@lucent.com) is a member of technical staff in the Networking Infrastructure Research Department in the Wireless Research Lab at Bell Laboratories — Lucent Technologies, Murray Hill, New Jersey. He received B.S. and M.S. degrees in mechanical engineering from the Université de Nantes, France, in 1993 and 1994, respectively, and his electrical engineer's degree in automatics from Ecole Centrale de Nantes in 1995 (ranked first in class). In October 2000 he received a Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge. Since January 2001 he has been with Bell Laboratories, where his research interests include information and communication theory, mobility management and resource allocation in wireless networks, as well as end-to-end data performance analysis and cross-layer optimizations.

KIN K. LEUNG (kkleung@ieee.org) received his Ph.D. degree from the University of California, Los Angeles, in 1985. After graduation he worked at AT&T (Bell) Labs and Bell Laboratories — Lucent Technologies until he joined Imperial College, University of London, United Kingdom in 2004. His research interests include broadband wireless networks, protocols, and network performance. He has received several awards for his contributions in these areas, and is an editor for *IEEE Transactions on Communications* and *IEEE Transactions on Wireless Communications*.

SAYANDEV MUKHERJEE [M'92] (sayan@lucent.com) received a Bachelor of Technology degree in electrical engineering from the Indian Institute of Technology, Kanpur, in 1991, and M.S. and Ph.D. degrees in electrical engineering from Cornell University, Ithaca, New York, in 1994 and 1997, respectively. Since 1996 he has been a member of technical staff at the Wireless Research Laboratory, Bell Laboratories — Lucent Technologies. His research interests include connectivity and capacity in multihop networks, wireless system simulations, and intelligent resource allocation in wireless systems.

GEE RITTENHOUSE (gee@lucent.com) received his undergraduate degree in physics from the University of California, Los Angeles. Then in 1993 he received his Ph.D. in electrical engineering and computer science from the Massachusetts Institute of Technology. He joined Bell Laboratories as a member of technical staff in 1993, where he worked on high-speed circuits using X-ray lithography for optical networking applications. He later joined the Wireless Research Laboratory at Bell Laboratories, where his research focused on RF front-end radio architectures and cellular system engineering. In 2000 he was promoted to director of the Wireless Technology Research Department and led several projects including MIMO system development, network optimization, wireless IP networks, and fourth-generation wireless systems. In 2002 he received the Bell Labs Fellow award. In 2003 he was promoted to vice president, wireless research. In this new role he is responsible for strategy, technology innovation, and the integration of wireless research into Lucent. He has numerous publications and patents in the areas of wireless systems and circuits.

LOUIS (SAM) SAMUEL (lsamuel@lucent.com) is currently a research technical manager at Lucent Technologies. He received his M.Eng. and Ph.D. from Queen Mary and Westfield College, University of London in 1995 and 1999, respectively. Since November 1998 he has been with Bell Labs Research, where he has been involved in the development of advanced protocols and network architectures for wireless communications.

HARISH VISWANATHAN (harishv@lucent.com) received a B.Tech. degree from the Department of Electrical Engineering, Indian Institute of Technology, Chennai, in 1992, and M.S. and Ph.D. degrees from the School of Electrical Engineering, Cornell University, in 1995 and 1997, respectively. He is currently a Distinguished Member of Technical Staff at Bell Laboratories — Lucent Technologies. His research interests include information theory, communication theory, wireless networks, and signal processing.

HAITAO ZHENG (htzheng@microsoft.com) received her B.S. degree (with highest honor) from Xian Jiaotong University in July 1995, and her M.S.E.E. and Ph.D. degrees in electrical and computer engineering from the University of Maryland, College Park, in May 1998 and July 1999, respectively. She joined the Wireless Research Laboratory at Bell Laboratories — Lucent Technologies as a member of technical staff in August 1999. Recently, she joined Microsoft Research Asia as a project lead and researcher. She was admitted to the highly gifted class of Xi'an Jiaotong Univeristy, P.R. China at 15. She received the 2002 Bell Laboratories President's Gold Award for outstanding level of innovation, technical excellence, and business impact, and the 1998–1999 George Harhalakis Outstanding Graduate Student Award for outstanding contributions in cross-disciplinary research from the Institute of System Research, University of Maryland, College Park. She has served as a guest editor for *JSAC* and *EURASIP* special issues, and as a TPC member of many conferences. Her research area includes wireless communications and networking, and multimedia computing.