A Dynamic Individualized Location Management Algorithm

John Scourias* and Thomas Kunz

Department of Computer Science
University of Waterloo
Waterloo, Ontario, Canada N2L 3G1

ABSTRACT

The challenge of supporting rapidly growing numbers of mobile subscribers, while constrained by limited radio spectrum, is being met through increasingly smaller radio cells. This, however, results in increased signalling for location management procedures, which reduces the bandwidth available for user traffic. Location areas in current systems, such as GSM, consist of static and arbitrarily-defined collections of cells, which do not take into account individual subscriber mobility patterns, either in space or time. A location management algorithm is proposed which uses the mobility history of individual subscribers to dynamically create individualized location areas, based on previous movements from cell to cell. The average duration spent in each visited cell is also maintained and is used to define paging which are most likely to contain the subscriber. An activity-based mobility model was developed to test the proposed algorithm. Overall, the dynamic algorithm incurred significantly lower location management costs, in terms of signalling messages generated, for all parameters examined.

INTRODUCTION

Over the last several years, the worldwide cellular communications market has faced explosive growth. This can be attributed to several factors, including decreasing prices, improved radio coverage, and compact, lightweight terminals. Given a fixed radio spectrum allocation, as the number of subscribers increases, the size of radio cells decreases in order to accommodate the higher subscriber densities. This concept of smaller and lower-power cells is behind the new generation of digital Personal Communication Systems, such as DCS1800. Smaller cell diameters pose several challenges, however, especially concerning location management. The signalling traffic generated by the increased subscriber densities is exacerbated by smaller cells, resulting in less radio bandwidth being available for user traffic, as well as additional transmission and processing requirements on the mobile network infrastructure.

Location management is concerned with the procedures required to enable the network to maintain location information for each active mobile terminal with a registered subscriber, and to efficiently handle the establishment of incoming calls. The two fundamental procedures which comprise the basis of location management are location updates and pages. Location updating is initiated by the mobile station, and informs the network of the subscriber's current location area. Paging is initiated by the network when an incoming call arrives. Paging messages are broadcast in one or more paging areas, contained within the current location area, and inform the target user of the incoming call.

Location areas and paging areas are relatively arbitrary collections of cells. It should be noted that in current networks, location areas are equivalent to paging areas, and the terms are often used interchangeably. The relationship between location updates, pages and location area dimensioning can best be illustrated by the two extreme approaches to location management. The first approach maintains the location of the mobile subscriber at the single cell level; the size of the location area in this case is exactly one cell. Whenever the mobile station moves to a new cell, which may happen very frequently in the case of an automobile-mounted mobile station, a location update is triggered. This is clearly quite inefficient. However, paging messages need only be sent to one cell, since the exact location of the mobile station is known. At the other extreme, the location area of the mobile is comprised of all cells in the network. In this case, location updates are not required at all. However, for every incoming call arrival to a mobile subscriber, the network must page every cell in the network; this is also an unsatisfactory approach.

RELATED WORK AND MOTIVATION

Location areas in current systems are statically-defined, using aggregate statistics and traffic patterns [6]. Static location areas are geographically and temporally fixed, and apply to all subscribers regardless of individual subscriber

* Currently at Lockheed Martin Corporation, Management and Data Systems (scourias@custlan.mds.lmco.com)

mobility patterns. The problem with the current approach, and the motivation behind this paper, is the realization that not all subscribers have the same mobility patterns, spatially or temporally, and thus static and globally applicable location areas are far from optimal.

User profiles have been suggested in several proposals as a way to individualize location areas [9, 10]. However, the detailed contents of the user profiles, as well as procedures for their creation and maintenance, are not well-defined. In addition, proposed algorithms to dynamically determine individual location areas [4, 5] generally suffer from excessive complexity, which is not supportable on mobile stations. Other location management proposals [7] are based on grouping subscribers with similar behaviour. Although plausible, such grouping is fairly arbitrary since subscriber mobility characteristics may change. The key issues are that algorithms must remain flexible, efficient, and not overly complicated, even in the face of mobility behaviour changes. In developing our proposal, ideas were borrowed from transportation theory and activity pattern theories [2, 3]. In the vast majority of cases, a trip is undertaken for a particular purpose occurring at a particular destination. The purpose behind a trip is generally some activity, such as work or shopping. These activities can take place at a limited number of locations. For example, most people work at a fixed location, and schools and shopping areas are located at specific sites. The home is normally the hub from which all other trips are taken, and is generally a place where subscribers spend a substantial amount of time during the day.

Location management techniques can make use of the fact that the average subscriber has a limited number of frequently visited locations, which form part of the daily mobility pattern. Of course, there will be deviations in the daily mobility pattern, but the more fundamental activity pattern seldom changes. The proposed algorithm exploits the spatial and temporal characteristics of trips. When fixed location areas are used, particular subscribers might cross several location areas on a single trip undertaken daily. A personal, or individualized, location area would tend to include frequently traversed cells, based on information derived from the user profile, thus eliminating repeated location updates.

ALGORITHM DESCRIPTION

The proposed location management algorithm attempts to utilize the mobility history, or user profile, of the subscriber to dynamically create location areas for individual subscribers and to determine the most probable paging areas. The user profile contains the number of transitions the subscriber has made from cell to cell, and the average duration of visits to each cell. Specifically, a counter $N_{a,b}$ is kept for each cell a of the number of times the subscriber has moved from cell a to each neighbouring cell b. Also maintained is the average time T_b that the subscriber has spent in cell b. As the mobile station moves from cell a to cell b, the counter $N_{a,b}$ is incremented, and a timer t_b is started, which is stopped either when the mobile station moves out of cell b, or the mobile station is turned off. The final value of t_b is used to update the average value of T_b.

The information stored in $N_{a,b}$ and T_b for all visited cells a and b can be represented as a directed graph, where the nodes represent visited cells, and the links represent transitions between cells. The weight of a link *(a,b)* is the value of $N_{a,b}$, and the weight of node b is the value of T_b. Note that data is not required for every cell in the network. Only cells actually visited have data entries, and even from those cells, only a small but contiguous set is required for the dynamic algorithm. The issue of which cell data to overwrite in memory is similar to a cache replacement problem. One possible algorithm would be to replace rarely visited cells. Although the amount of data required is not large (*see ANALYSIS AND EVALUATION section*), the proposed algorithm performs better with the availability of more information.

A location update is triggered when a subscriber enters a new cell which is not part of the previous location area, when the mobile station is first powered on in a cell, or periodically as required by the network. Having decided a location update is necessary, the mobile station looks up the new cell in the user profile. If it is not found (i.e., if the user has never visited that cell or if the cell has been purged from memory), a classical location update is performed. For this purpose, the network is overlayed with relatively large static location areas, similar to current implementations. Note that the mobile continues to update its user profile while the mobile roams within a static location area. If the user profile contains the new cell, the list of its neighbouring cells previously visited is read, together with the number of times the subscriber has moved to those cells from the new cell (i.e., the value of $N_{a,b}$ is read for all cells b). Due to the physical layout of cells in cellular networks, there can be only a limited number of neighbouring cells, normally six. The average weight W of the links to each neighbouring cell is then calculated.

The cells corresponding to the links whose weights are greater than or equal to the average weight W are added to the new location area, in order according to the link weight. The average link weight is used to discriminate between paths which the subscriber often traverses, and paths which are seldom traversed. The mean is a relatively simple way of determining the break between frequently and infrequently used cells, in a non-arbitrary manner.

Once selected cells from the first *ring* of neighbouring cells have been added to the personal location area, the above steps are repeated using the newly selected cells, by decreasing link weight order. The average weight of the cell's outgoing links is calculated, and any nodes whose associated link weight is greater than or equal to the average are added to the personal location area, such that all cells descended from the first ring of cells are ordered by link weight. Since it is possible that one of the selected links leads back to a previous node, the algorithm ensures that duplicated cells are not added to the personal location area. These steps are repeated until the personal location area size has reached its limit, or until no other cells are left for inclusion.

The paging strategy is closely tied with the location updating strategy. A low value of T_n, the average time spent in cell n, indicates that the subscriber spends, on average, little time roaming in that cell. For example, an intermediate cell along the path from home to work may be frequently crossed, but the subscriber would seldom be found there. Cells with a high value of T_n (for example, work or home cells) indicate that the subscriber is quite likely to be found there. During a location update, all T_n values for the cells comprising the new location area are transmitted to the network, to provide current information for subsequent paging attempts.

The problem of determining which specific cells to page first, given a corresponding list of T_n values, is similar to the problem of determining which links to include in the location updating strategy. The average value of T_n among all cells in the current location area is calculated, and cells where T_n is greater than or equal to the average form the paging area used in the first round of paging. If this paging attempt is unsuccessful, all cells in the location area are paged in the second round of paging (in case the subscriber moved to the first paging area just after the first round of paging). It is possible to modify the algorithm to create more paging areas, but this would increase paging-related call setup delays, adversely affecting the quality of service.

ANALYSIS AND EVALUATION

The proposed location management algorithm was compared to several variations of the current fixed algorithms, primarily in terms of radio bandwidth efficiency. A mobility model was developed to provide simulated data for the location management algorithm comparisons. The model provides a realistic simulation of the daily movements of individual mobile subscribers, in terms of sequences of traversed radio cells). Activity patterns and durations derived from a regional trip survey [1] for the Region of Waterloo (Ontario, Canada) form the basis of the model, together with related geographic information representing cells, paths, and distances. Information such as home, work, and school locations is maintained for each subscriber. The mobility modelling algorithm generates subscriber trips over time, in terms of cells traversed and time spent in each cell. For the model, activities are selected randomly based on activity distributions and time of day, activity durations are selected randomly from duration distributions, and destinations are either random or based on subscriber-specific data.

The simulated environment consists of 45 radio cells grouped into four different fixed location area layouts, in an ad hoc manner based on geographical proximity. The fixed location area groups, or fixed algorithms, are equivalent to paging areas, and represent extreme as well as typical configurations. Fixed Group 0 has one cell per location area. Fixed Group 1 has thirteen location areas of three or four cells each, and the similar Fixed Group 2 has three to five cells in each of eleven location areas. Finally, Fixed Group 3 has five large location areas, with eight to ten cells per location area. The proposed dynamic algorithm has a parameterized maximum location area size.

The simulation output provides details about paging and location update messages generated over time by a simulated subscriber, for each of the fixed algorithms and for the proposed dynamic algorithm. The movements and messages generated by each subscriber were logged for 15 simulated days, an interval assumed at the outset to be long enough for the model to reach a steady state. The results are based on 2400 individual simulation runs, with varying dependent parameters, such as number of daily call arrivals, and random internal variables, such as home cell location. The evaluation of the proposed algorithm was based on the goals of minimal signalling overhead, minimal call setup delays, and minimal computational complexity and data storage requirements.

Minimal signalling overhead

The primary goal of location management proposals is to minimize signalling, especially over the radio interface, but also within the fixed network, in order to reduce signalling bandwidth and increase capacity. The overall average number of location updates and cells paged was calculated, as well as the total location management cost. The latter is a linear combination, used to quantify the net effect of location updates and paging messages. A location update is more expensive than a page, due to the need to establish a signalling channel, but the exact cost can be quantified in several ways. In this paper, the ratio c of the cost of location updates to pages has been defined as a choice of 5 and 10, roughly representative of the number and size of messages exchanged for each procedure.

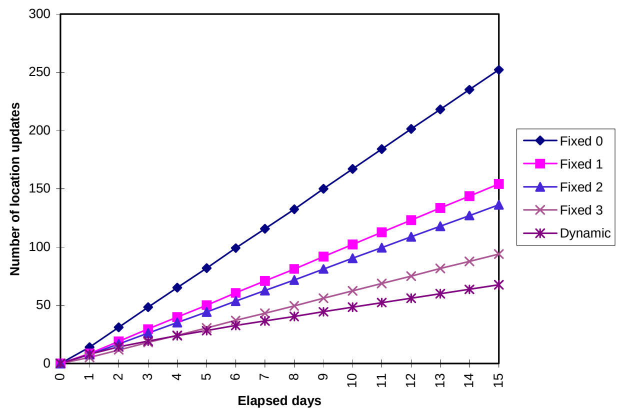

Figure 1. Overall average number of location updates over time.

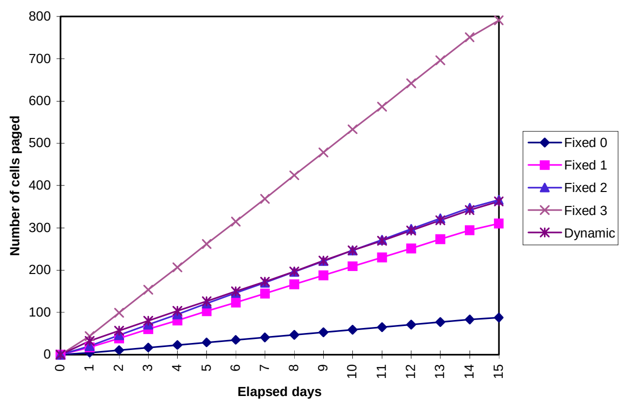

Figure 2. Overall average number of cells paged, for 6 incoming calls, over time.

After an initial 'learning curve' of three days, in which the user profile is built up, the dynamic algorithm significantly outperforms the fixed algorithms in terms of location updates, as shown in Figure 1. In terms of cells paged, the optimal solution is Fixed Group 0, with one-cell location areas. The dynamic algorithm performs similarly to Fixed Group 2, as shown in Figure 2. However, since the dynamic algorithm has the lowest number of location updates as well as a relatively low paging cost, the net effect is that the dynamic algorithm has a significantly lower total location management cost than the fixed algorithms. This was the case for both values of *c* (Figures 3 and 4), and for all levels of incoming calls (the case of 6 daily incoming calls is shown for brevity). The initial 'learning curve' remained about 3 days for all cases.

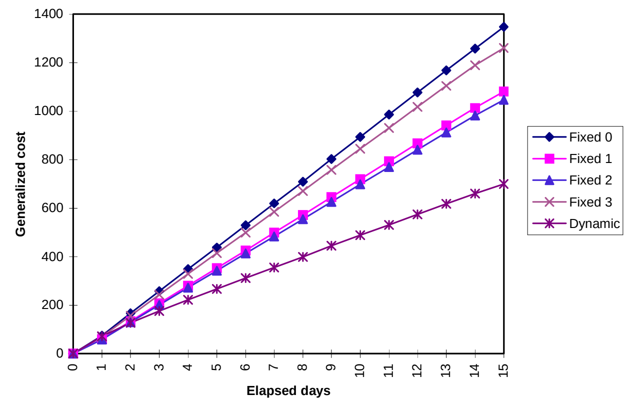

Figure 3. Overall average location management cost, for 6 incoming calls, over time, for c = 5.

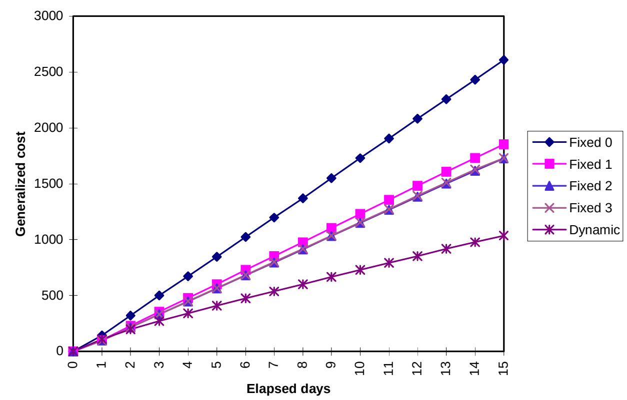

Figure 4. Overall average location management cost, for 6 incoming calls, over time, for c = 10.

Further analyses were conducted to investigate the effect of subscriber type (e.g., employed or student), work location variability (fixed or random work location), and dynamic location area size, for different incoming call levels, on the number of location updates, cells paged, and total location management cost. Although not discussed in this paper, the dynamic algorithm consistently outperformed the fixed algorithms in terms of location management cost, except in the special case of full-time employed subscribers with random work locations, and high numbers of incoming calls, for *c* equal to 5. In this case, since work location (a major daily activity) is random, the dynamic algorithm as implemented here cannot form frequently visited location areas and paging areas. Even in this case, however, the cost difference between the fixed and dynamic algorithms was less than 10%.

Minimal call setup delays

Due to the delays involved in paging (waiting for a response, possibly repeating pages due to the harshness of the radio environment, and eventual call setup), repeating the paging process for more than one or two paging areas can lead to unacceptably high delays, in terms of quality of service. The proposed paging algorithm pages at most two paging areas, and although this may in some cases result in higher paging costs, a short overall call setup time is preferable. If required, the algorithm can be easily modified to page more than two paging areas.

Minimal computational complexity and data storage

In providing its significant location management cost savings, the dynamic algorithm requires some additional logic and memory on the mobile station and network. In terms of computation, the proposed algorithm requires the mobile station to record the average duration spent in each cell, and the number of transitions from one cell to another. While roaming, a new cell identifier is compared against a list of cell identifiers (the personal location area, containing less than 20 cells), instead of comparing the stored Location Area Identifier (LAI) against the broadcast LAI.

In addition, information related to paging (i.e., the average visit duration for cells in the current location area) needs to be transferred from the mobile station to the network. This transfer would generally be done during a location update, making use of the existing signalling channel. The amount of data transferred is directly related to the size of the dynamic location area.

The size of the user profile stored at the mobile station, which determines how much cell transition and stay duration information can be used by the dynamic algorithm, was not a constraint in the analysis. A larger user profile allows the algorithm to function more effectively, by maintaining information on more cells. Assuming 4-octet integers, the memory requirements for each mobile would be 45 cells $\times$ [4-octet cell identifier + 6 labelled neighbouring cell transition counters (2 4-octet integers each) + 4-octet average stay duration], or roughly 2.5 kbytes.

On the network side, the location database (Home Location Register, or HLR) needs additional memory to store the list of cells forming the current dynamic location area, as opposed to storing only the current LAI. The associated paging-related information (average stay duration for the cells in the current dynamic location area) also needs to be stored. The cost associated with the relatively small amount of additional storage required is negligible compared to revenue from the additional radio interface capacity provided by the dynamic algorithm, since radio spectrum is a fixed resource.

SUMMARY

A location management algorithm was developed, which dynamically generates individualized location areas and paging areas based on a user profile. A novel mobility model was developed and used to compare the performance of the proposed algorithm with several variations of currently implemented fixed algorithms. The dynamic algorithm significantly outperformed the fixed algorithms in terms of total location management cost, at a small cost of additional logic and memory in the mobile station and network.

REFERENCES

[1] Tranplan Associates, "Waterloo Region Travel Survey 1987: An Overview of the Survey Findings," Regional Municipality of Waterloo, Department of Planning and Development, October 1989.

[2] Juan de Dios Ortuzar and Luis G. Willumsen, Modelling Transport, 2nd edition. John Wiley and Sons, 1994.

[3] Thomas A. Domencich and Daniel McFadden, Urban Travel Demand: A Behavioral Analysis. North-Holland, 1975.

[4] Ravi Jain et al, "A Caching Strategy to Reduce Network Impacts of PCS," IEEE Journal on Selected Areas in Communications, 12(8), November 1994.

[5] U. Madhow, M. L. Honig, and K. Steiglitz, "Optimization of Wireless Resources for Personal Communications Mobility Tracking," IEEE/ACM Transactions on Networking, 3(6), December 1995.

[6] Michel Mouly and Marie-Bernadette Pautet, The GSM System for Mobile Communications. Published by the authors, 1992.

[7] Sadaatsu Okasaka et al, "A New Location Updating Method for Digital Cellular Systems," in IEEE 41st Vehicular Technology Conference, 1991.

[8] Norbert Oppenheim, Urban Travel Demand Modeling. John Wiley and Sons, 1995.

[9] N. Shivakumar and J. Widom, "User Profile Replication for Faster Location Lookup in Mobile Environments," in ACM MOBICOM, 1995.

[10] Sami Tabbane, "An Alternative Strategy for Location Tracking," IEEE Journal on Selected Areas in Communications, 13(5), June 1995.