The Minimum Expected Cost Paging Problem for Multi-System Wireless Networks *

Abhishek Roy[1], Archan Misra[2] and Sajal K. Das[1]

[1] CReWMaN Lab, Computer Science and Engineeering Department
University of Texas at Arlington, Arlington, TX 76010, USA
aroy@cse.uta.edu, das@cse.uta.edu
[2] IBM T. J. Watson Research Center, Hawthorne, NY 10532, USA
archan@us.ibm.com

Abstract. The paper discusses the problem of optimal (minimum expected cost) paging for multi-system, fourth generation ($4G$) wireless cellular networks. A mobile node can lie simultaneously in multiple, overlapping cells. Hence, at a particular point in the paging sequence, the probability of successfully paging the mobile in a particular cell will depend not just on its marginal residence probability in that cell, but instead on the joint probability distributions. Due to this fundamental difference with a single-system cellular network, the problem of obtaining the optimal paging sequence turns out to be *NP-complete*. After proving this NP-completeness result, the paper presents a greedy, and intuitively appealing heuristic, where the paging sequence is obtained by choosing cells in the decreasing order of their conditional residence probabilities. The probability is computed *conditional on the mobile not being present in the previous cells of the paging sequence*. The paper describes how such probabilities can be computed by using a mobile-node driven, asymptotically-optimal, location update strategy and also quantifies the resultant savings in the paging cost.

1 Introduction

Location tracking in a wireless cellular network infrastructure consists of two fundamentally distinct operations: (1) location update and (2) paging. The problem of paging a mobile device within a single wireless cellular network has been extensively studied in the literature. Advances in intelligent paging strategies include the concept of customized paging profiles [10] for each mobile node ($\mathcal{MN}$), and obtaining optimal paging in terms of cell residence probabilities, both with and without constraints on paging latency [11]. The intuitive but fundamental result in [11] states that in the absence of any constraints on the paging latency, the expected number of paging messages is minimized when the various cells are paged in the decreasing order of their residence probabilities. However, all these results and algorithms consider the case of a *single* (homogeneous) cellular access infrastructure, where all the cells belong to the same cellular network. In such an infrastructure, a mobile node is always connected via only one cell at any instant of time (even though the cells may have spatial overlap), making the mobile's residence in different cells mutually exclusive (disjoint) events.

There has recently been significant interest in developing *integrated wireless infrastructures* (e.g. $4G$ wireless networks), where an $\mathcal{MN}$ seamlessly attaches to and switches between multiple independent, heterogeneous cellular networks, possibly utilizing different access technologies [3]. Such a vision hinges on the fact that mobile nodes are already being equipped with multiple radio interfaces (e.g., GSM, CDMA and Bluetooth), and can thus be simultaneously attached to multiple cellular networks. Recent work on this area (e.g., [1, 12]) has shown how an integrated location management framework, that coordinates the tracking of an $\mathcal{MN}$ in each individual sub-network[3], can outperform a framework where each sub-network manages the $\mathcal{MN}$'s location independently. The integrated location management problem is particularly challenging, since cells belonging to different sub-networks can not only be of widely varying sizes, but also exhibit arbitrary overlap and even be discontinuous (e.g., a wireless provider supporting

* The work of Abhishek Roy and Sajal K. Das is supported by NSF grants IIS-0121297, IIS-0326505 and TxTEC

[3] In this paper, each independent access network is called a sub-network. The integrated network is thus a collection of such sub-networks.

multiple 802.11-based wireless LAN hot-spots). This motivates us to investigate the problem of optimal or minimum expected cost paging algorithms for multi-system $4G$ cellular networks.

The paging problem differs from the case of a single network due to the following three reasons:

(i) The optimal paging algorithm must explicitly consider the fact that different sub-networks can impose different signaling costs for transmitting a single paging message.

(ii) Unlike the single-system scenario, the $\mathcal{MN}$'s residence in different cells of the integrated network are *no longer mutually exclusive and independent events.* Due to the arbitrary overlap between the cells belonging to different sub-networks, we shall see that the overall expected paging cost associated with any given paging sequence is a function, not of the unconditional marginal residence probabilities in each cell, but of the **the conditional residence probabilities** of the $\mathcal{MN}$ in different cells.

(iii) In a single-system network, paging is invoked only in the case of an *idle* mobile node. With multiple interfaces, however, an $\mathcal{MN}$ can be active in one sub-network (using a particular radio interface) but idle in another sub-network (using a different transceiver). In fact, to page an $\mathcal{MN}$ which is idle in one network, one can conceptually send an *indirect paging* message to the $\mathcal{MN}$ via an alternative sub-network where the $\mathcal{MN}$ is currently active. Since the alternative sub-network merely unicasts such a message to the $\mathcal{MN}$'s current location, substantial savings may be obtained by avoiding the need to search for the $\mathcal{MN}$ over the cells of the original sub-network.

The fact that the expected paging cost is dependent on the conditional residence probabilities turns out to be a key factor in designing an optimal paging strategy for such integrated multi-system networks. In fact, we shall show that the general problem of determining the optimal paging sequence is NP-complete. This is an important distinction from the single-system case, where the greedy procedure of "more-likely cells first" (paging cells in the decreasing order of their residence probabilities) results in an optimal sequence. Having demonstrated the hardness of the general optimization problem, we shall then propose an efficient heuristic that constructs the paging sequence for an $\mathcal{MN}$ proceeds in the decreasing order of the *conditional* residence probabilities (weighted appropriately by the cost of a paging message in a cell). This heuristic corresponds to a greedy algorithm, such that the next cell in the sequence is chosen to be the one where the $\mathcal{MN}$ has the highest likelihood of being found, *conditioned on its being absent in any of the cells already in the sequence.* The use of conditional residence probabilities reflects the intuition that it is useless to page the mobile node in a cell of one sub-network if it has already been unsuccessfully paged in a spatially overlapping cell (or set of cells) belonging to other sub-networks.

We shall also provide a location management architecture where this conditional probability-based paging heuristic can be practically implemented. This architecture is based on the multi-system LeZi-Update algorithm, described in [9], where the $\mathcal{MN}$ essentially sends entropy-coded updates of its entire movement sequence. A key advantage of the LeZi-Update algorithm lies in obtaining the asymptotic optimality from an information theoretic approach. Moreover, from a practical standpoint, the architecture allows the system to learn of the $\mathcal{MN}$'s movement sequence without requiring the sub-networks to expose their topology databases to one another. We shall show how this knowledge of the $\mathcal{MN}$'s movement history is adequate to first derive the $\mathcal{MN}$'s *joint residence probability distributions*, and subsequently the heuristic-based paging sequence. Simulation studies are also provided to quantify the performance savings of the proposed, greedy algorithm over a scheme that merely pages cells in the decreasing order of their marginal residence probabilities.

The paper is organized as follows: Section 2 formulates the mathematical modelling of the expected paging cost in multi-system wireless networks. The NP-completeness of this problem is proved in Section 3. In Section 4 we provide an efficient greedy heuristic, based on mobile node's conditional residence probabilities. Subsequently, the greedy algorithm in integrated location management framework is discussed in Section 5. Simulation results in Section 6 points out sufficient gains in expected paging cost in comparison to existing paging schemes. Finally Section 7 concludes the paper with pointers to future researches.

2 Modelling Expected Paging Cost in Multi-System Wireless Networks

Figure 1 shows an example of an *integrated, multi-system 4G network*, comprising a collection of independent sub-networks, like, satellite, personal communication systems (PCS) and campus area networks. For mathematical modelling, let the network consists of N sub-networks $\{S_1, S_2, \cdots, S_N\}$, where each sub-network (S_i) is a collection of (either partitioned or overlapping) cells such that C_i^j represents the j^{th} cell in S_i, for $1 \leq i \leq N$. Let $|S_i|$ denote the number of cells in S_i. For notational convenience, let each sub-network have an additional cell ϕ to capture the disconnected state (when the $\mathcal{MN}$ is outside the coverage area of the entire sub-network). Also, let PG_i represent the cost of transmitting a single paging message in a single cell of sub-network S_i.

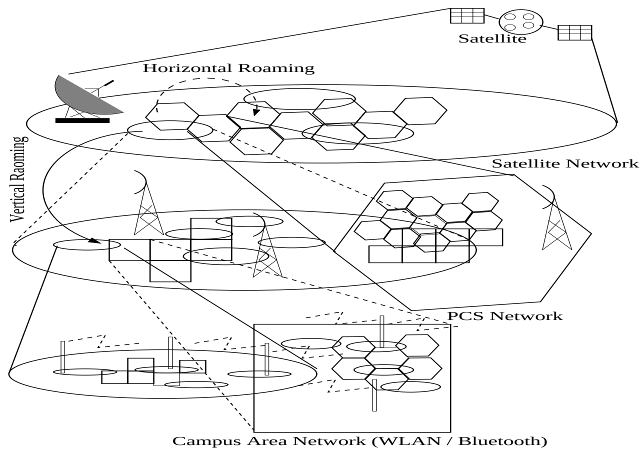

Figure 1. A Multi-System Heterogeneous Wireless Network

Due to the potentially arbitrary overlap among the cells of different sub-networks, we have shown in [9] how the $\mathcal{MN}$'s location at any point can be represented as a vector-valued random variable $\bar{X}$ of dimension N, where the i^{th} element of the vector corresponds to the current cell of sub-network S_i. For example, if $N = 3$, the vector $\bar{X} = [C_1^2, C_2^5, C_3^\phi]$ implies that the $\mathcal{MN}$ is in the 2^{nd} cell of S_1, the 5^{th} cell of S_2 and disconnected from S_3. Clearly, the user's location can then be modelled by the joint probability distribution:

$$Pr\left(\bar{X} = [x_1, x_2, \cdots, x_N]\right) = Prob\left(\mathcal{MN}\text{is located in } C_1^{x_1} \bigwedge C_2^{x_2} \bigwedge .... \bigwedge C_N^{x_N}\right), \quad (1)$$

where $\bigwedge$ denotes the intersection or overlap between any two cells. The arbitrary overlap among different cells of different sub-networks implies that the $\mathcal{MN}$'s residence in different cells are not disjoint events. Clearly, an $\mathcal{MN}$ can be simultaneously resident in both C_1^2 and C_2^5. Accordingly, given any paging sequence, the likelihood of finding the $\mathcal{MN}$ at a particular cell of the sequence will depend not just on its marginal (unconditional) probability, but instead on the joint probability of the $\mathcal{MN}$ in all previous cells of the sequence.

From a practical perspective, the joint probability will be a function of the spatial overlap between the cells. However, our goal is to devise a location management framework for a truly multi-operator $4G$ environment, where the $\mathcal{MN}$ may be associated with different providers on different interfaces. As individual sub-network operators may be unwilling to publicly share the spatial topology of their cellular network, we cannot assume the direct availability of such overlap information between cells of different sub-networks. However, the joint probability distributions of the form specified by Equation 1 is sufficient to obtain the expected paging cost associated with any particular paging sequence Seq.

To express the paging cost for any given sequence Seq, we introduce some additional notation for convenience. Let the L-length sequence be described by the Seq $= \{c_{S(1)}^1, c_{S(2)}^2, \ldots, c_{S(L)}^L\}$. In other words, the cells are indexed based on the given sequence, and $S(i)$ (different from S_i, where S_i is the i^{th} sub-network) refers to the sub-system corresponding to the i^{th} cell in the sequence. Also, let RV_i denote the unconditional event that the $\mathcal{MN}$ is in cell $c_{S(i)}^i$, and let $p(RV_i)$ denote the *marginal probability* of the $\mathcal{MN}$ being located in cell $c_{S(i)}^i$. Now, the i^{th} cell in this sequence is paged only if the $\mathcal{MN}$ is not found in any of the previous $(i-1)$ cells of Seq. Moreover, the probability of a successful paging at the i^{th} attempt must equal the probability that the $\mathcal{MN}$ is indeed in cell $c_{S(i)}^i$, but not in any of the previous cells $\{c_{S(i-1)}^{i-1}, c_{S(i-2)}^{i-2}, \ldots, c_{S(1)}^1\}$ that have been unsuccessfully paged. This is clearly given by the probability mass of the event $(\vee_{k=1}^{i-1} RV_k)' \wedge RV_i$ (where $\vee$, $\wedge$ and $'$ denote the set-theoretic union, intersection and complement operators respectively), involving the $\mathcal{MN}$'s residence in the cells $\{1, \ldots, i-1, i\}$ of Seq. For ease of notation, let $p(i|\widehat{i-1})$ denote this marginal probability, represented by:

$$p(i|\widehat{i-1}) = p(RV_i \wedge (\vee_{j=1}^{i-1} RV_k)'). \quad (2)$$

Accordingly, the expected paging cost of the sequence, Seq, is given by:

$$E[\text{Seq}] = \sum_{i=1}^{L} PG_{S(i)} \times p(i|\widehat{i-1}), \tag{3}$$

where $PG_{S(i)}$ is the paging cost of the cell $C^i_{S(i)}$. This is really the generalized form of the expression for the expected paging cost derived in [11] for a single-system network with non-overlapping cells. For a single-system network, the differential probabilities equal the marginal probabilities, since the residence events are mutually exclusive (the $\mathcal{MN}$ can be connected with the base station of only one cell at any time).

3 NP-Completeness of the Optimal Paging Sequence

Given the definition of the expected paging cost in terms of the joint residency distributions over the cells of different sub-networks, the problem of determining the optimal paging sequence, Seq_{opt}, reduces to finding the sequence that minimizes the expected paging cost defined in Equation 3. Of course, the sequence must also be guaranteed to successfully find the $\mathcal{MN}$ if it is indeed within the coverage area of any cell in any of the sub-networks. To express this mathematically, let $SEQ = \{\text{Seq}_1, \text{Seq}_2, \cdots\}$ denote the set of all plausible sequences, and the j^{th} member of this set be described by the sequence $\text{Seq}_j = \{c^1_{S(1)}(j), c^2_{S(2)}(j), \ldots, c^L_{S(L_j)}(j)\}$, where L_j is the length of Seq_j. The *optimal paging problem* **(P1)** is to find a solution Seq_{opt} such that:

$$\text{Seq}_{opt} = \underset{j \epsilon SEQ, \sum_{k=1}^{L_j} p_j(i|\widehat{i-1}) = 1}{\text{arg min j}} \sum_{i=1}^{L} PG_{S(i)} \times p_j(i|\widehat{i-1}). \tag{4}$$

In other words, we are looking for the sequence (within the set of all plausible sequences) that minimizes the associated expected paging cost, while ensuring that the sum of the conditional probabilities of the sequence is 1. Unfortunately, this general problem turns out to be NP-complete. To demonstrate this, let us, for ease of explanation, consider the simpler version, say **(P2)** of problem **(P1)** where *the PG_is are all identical and equal to 1. In other words, the paging costs in all sub-networks are identical.* The corresponding new optimal sequence is denoted by $\widetilde{Seq}_{opt}$.

Theorem 1. *The optimization problem* **(P2)** *is NP-complete.*

Proof: As in most cases of this type, we show that a special instance of this problem may be modelled as a known NP-complete problem. In particular, we show that a constrained version of generic optimization problem (**P2**) can be modelled as the *sequential cover problem* [5], which is itself a generalized version of the well known *set-covering problem* [6], which is known to be NP-complete.

Sequential-Cover Problem [5]: An instance $(Y, \mathcal{F})$ of the set covering problem consists of a finite set Y of base elements $\{y_1, y_2, \ldots, y_n\}$, and a family $\mathcal{F}$ of m subsets of $Y = \{f_1, f_2, \ldots, f_m\}$, such that every element of Y belongs to at least one subset (f_i) in $\mathcal{F}$, i.e.,

$$Y = \bigcup_{i=1}^{m} f_i. \tag{5}$$

Moreover, each element $y_i \in Y$ is associated with a probability distribution $Q = (q_1, q_2, \ldots, q_n)$ such that $\sum_{i=1}^{n} q_i = 1$. A subset $f \in \mathcal{F}$ **covers** its elements. Let z be a positive value s.t. $0 \leq z \leq 1$. The problem then is to find **a value** m', such that $0 \leq m' \leq m$, **and a sequence** $\text{Seq}^{m'} = f_{\pi(1)}, f_{\pi(2)}, \ldots, f_{\pi(m')}$ of m' subsets in $\mathcal{F}$, such that:

1.
$$\sum_{y_l \in \bigcup_{i=1}^{m'} f_{\pi(l)}} q_l \geq z, \tag{6}$$

i.e., the selected sequence has a probability mass of at least z on the underlying distribution, and

2. $\text{Seq}^{m'}$ minimizes the expected size of the search given by:

$$\sum_{y_l \in \bigcup_{i=1}^{m'}} q_l \times min\{k | y_l \in f_{\pi(k)}\} + m' \sum_{y_l \notin \bigcup_{i=1}^{m'}} q_l. \tag{7}$$

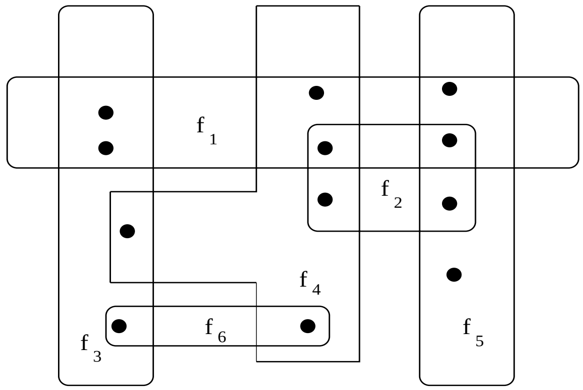

Figure 2. Solution of Set-covering problem

To see the equivalence between our multi-system paging problem and the sequential-covering problem, note that the joint probability of an $\mathcal{MN}$ in multiple simultaneous cells arises from the underlying spatial overlap of the cells. Thus (see Figure 2), the entire spatial region within which the $\mathcal{MN}$ may be located can be defined to be the set Y, where the individual elements y_i of Y correspond to the suitably small individual spatial grids (in Figure 2, the individual elements are represented by the black circles). Given this analogy, each cell can be considered to be a set of these elements, corresponding to the collection of the elements of Y that constitute the cell. Thus, each cell may be considered to be a subset f_i, such that the family of subsets $\mathcal{F}$ is nothing but the collection of all these cells. Mathematically speaking, we have $Y = \cup_{\forall i, 1 \leq i \leq N, \forall j \in S_i} C_i^j$, and $\mathcal{F} = \{C_i^j : 1 \leq i \leq N, \quad j \in 1, \ldots, |S_j|\}$. Moreover, let q_i now denote the unconditional residence probability of the $\mathcal{MN}$ in grid y_i, such that the $\mathcal{MN}$'s residence in different infinitesimal grids are mutually disjoint events. The marginal residence probability within any cell C_i^j is distributed among the elements constituting the corresponding subset f_k, such that this probability equals the sum of the probabilities of the individual elements of f_k. The probability of a user's simultaneous residence in η cells is represented by the containment of the elements in the intersection of η such subsets f_k:

$$Pr[\mathcal{MN} \in \{\bigwedge_{1 \leq i \leq N, j \in S_i} C_i^j\}] = \sum_{y_l \in \bigcap_{k=1}^{\eta} f_k} q_l$$

This analogy between the cellular paging and the sequential set covering problem is tabulated in Table 1.

We can then see that the Equation 3 associated with any sequence Seq_i can be equivalently written to be identical to Equation 7, if we set $z = 1$, such that we are looking for a sequence that gives *complete coverage*. Accordingly, with $z = 1$, we have shown that solving the optimization problem **(P2)** is equivalent to solving the *sequential cover problem* [5]. However, as shown in [5], the sequential cover problem is NP-complete. This completes the proof that determining the optimal paging sequence for arbitrary overlapping cellular coverage is also an NP-complete problem.

Table 1. Mapping of Parameters between **(P2)** and Sequential Covering

$P2$	Sequential-covering
$\bigcup_{1 < i < N, \forall j \in S_i} C_i^j$	Y
C_i^j	f_k
$Pr[\mathcal{M}N \in C_i^j]$	$\sum_{y_l \in f_k} q_l$
$\{C_i^j\}$	$\mathcal{F}$
$\widetilde{Seq}_{opt}$	$\text{Seq}^{m'}$

4 The Conditional Probability-Based Greedy Heuristic

Since the determination of the optimal paging sequence is NP-complete, we now focus on approximate algorithms or heuristics to intelligently derive a suitable paging sequence. Based on the cost structure in Equation 3, we explore one particular greedy algorithm in this paper. Before presenting the details of the greedy procedure, we first use a simple example to illustrate why a sequence based on decreasing order of marginal probabilities is sub-optimal in almost all cases.

Consider the $\mathcal{MN}$'s typical movement scenario in a heterogeneous network having two different types of sub-networks S_1 and S_2; for simplicity, assume that both have identical paging cost/message (normalized to 1). Let us assume that the cells of S_1 and S_2 are respectively represented by numbers and alphabets. A typical movement history of the $\mathcal{MN}$ may then consist of the tuples $[(C_1^1, C_2^a)$, (C_1^1, C_2^e), (C_1^1, C_2^a), (C_1^2, C_2^a), (C_1^2, C_2^d), (C_1^1, C_2^d), (C_1^1, C_2^a) , (C_1^1, C_2^f), (C_1^3, C_2^f), (C_1^5, C_2^f), $(C_1^5, C_2^b)]$. Given these 11 samples,we can compute the marginal and joint residence probabilities in various cells. Let, $p_i^j = Prob[\mathcal{MN} \in C_i^j]$ be the unconditional, marginal probability of $\mathcal{MN}$'s residence in j^{th} cell of the sub-network S_i. Thus, we have: $p_1^1 = 6/11, p_1^2 = 2/11, p_1^3 = 1/11, p_1^4 = 0/11, p_1^5 = 2/11, p_2^a = 4/11, p_2^b = 1/11$, $p_2^c = 0/11$, $p_2^d = 2/11$, $p_2^e = 1/11$, $p_2^f = 3/11$. Now, if the system pages the cells in descending order of p_i^js, the paging sequence will be the sequence $\{C_1^1$, C_2^a, C_2^f, $(C_1^2$, C_1^5, C_2^d in some random order), $(C_1^3$, C_2^b, C_2^e in some random order)$\}$.

Intuitively, this approach is not preferable, since we did not consider the joint distribution and the conditional probabilities, but only the marginal distributions. By observing the movement history of the $\mathcal{MN}$, we notice that that the $\mathcal{MN}$'s residence in cell C_2^a almost always occurs simultaneously with its residence in C_1^1. Hence, if $\mathcal{MN}$ is absent in cell C_1^1, then with high probability it is also absent in cell C_2^a. In other words, cell C_1^1 and C_2^a are almost mutually exclusive and the probability of the $\mathcal{MN}$ being in C_2^a but not in C_1^1 is given by $p(a|\widehat{1}) \approx 0$. Hence, after an unsuccessful paging attempt in cell C_1^1, it is wise to eliminate the tuples containing C_1^1 in them before deciding the next cell to be paged. Referring back to the original movement history, removal of the tuples containing C_1^1 leaves us with the samples (C_1^2, C_2^a), (C_1^2, C_2^d), (C_1^3, C_2^f), (C_1^5, C_2^f), (C_1^5, C_2^b). Subsequently, the new conditional probabilities (order-1) is given by: $p(2|\widehat{1}) = 2/5$, $p(3|\widehat{1}) = 1/5$, $p(5|\widehat{1}) = 2/5$, $p(a|\widehat{1}) = 1/5$, $p(d|\widehat{1}) = 1/5$, $p(f|\widehat{1}) = 2/5$, $p(b|\widehat{1}) = 1/5$. The system can now select the cell having the *highest conditional probability*, given that it is not in cell 1, i.e., randomly any cell from C_1^2, C_1^5, C_2^f. Suppose the cell C_2^f is chosen. Now, if $\mathcal{MN}$ is not found even in C_2^f, then we can proceed in a similar way to eliminate the tuples having C_2^f, thus reducing the search space to (C_1^2, C_2^a), (C_1^2, C_2^d) and (C_1^5, C_2^b). Now, the conditional probabilities of order-2 are computed as: $p(2|\widehat{1,f}) = 2/3$, $p(a|\widehat{1,f}) = 1/3$, $p(d|\widehat{1,f}) = 1/3$, $p(5|\widehat{1,f}) = 1/3$ and $p(b|\widehat{1,f}) = 1/3$. Hence, the system attempts to page $\mathcal{MN}$ in cell C_1^2. Thus, the new, intelligent paging sequence is: $\{C_1^1, C_2^f, C_1^2\}$, which results in a lower expected paging cost than $\{C_1^1, C_2^a, C_2^f, \ldots\}$.

This observation defines our greedy heuristic for defining the paging sequence. The mathematically concise description of the algorithm is deferred to Figure 4, which illustrates the use of our greedy algorithm within our recently proposed location management framework [9]. Intuitively speaking, at each iteration of the algorithm, the greedy procedure selects the cell where the $\mathcal{MN}$ has the highest residence probability, given that it cannot simultaneously be in any of the cells already included in the sequence.

We also use a simple example to illustrate why this heuristic is not optimal in call cases. Consider the physical layout of three cells $C1, C2$ and $C3$, as well as the joint probabilities for the various regions, illustrated in Figure 3. Then the marginal probabilities are given by $p(C1) = 0.36$, $p(C2) = 0.46$ and $p(C3) = 0.44$ It is easy to see that our greedy heuristic will result in the paging sequence $\{C2, C1, C3\}$, with an expected paging cost of $(1 \times 0.46 + 2 \times 0.2 + 3 \times 0.14) = 1.28$. However, the problem with this strategy is that it fails to account for the fact that cell $C2$ is *entirely contained* in the union of $C1$ and $C3$. Accordingly, if we first paged cell $C3$, we would then only have to page cell $C1$ and locate the mobile node without ever having to page the $\mathcal{MN}$ in $C2$. Thus, the sequence $\{C3, C1\}$ has an expected paging cost of $(1 \times 0.44 + 2 \times 0.36)) = 1.16$, which is lower than our greedy sequence $\{C2, C1, C3\}$. In general, this example illustrates the reason behind the hardness of the problem: *due to the possibility of arbitrary overlap, one must not only consider the likelihood of residency in a particular cell, but also whether that likelihood could be efficiently subsumed via a search in one or more other cells.*

Our greedy heuristic is essentially equivalent to the greedy search algorithm for the sequential cover problem, as discussed in [5]. Accordingly, we can apply the following result on the performance bound of our greedy heuristic for construction of the paging sequence.

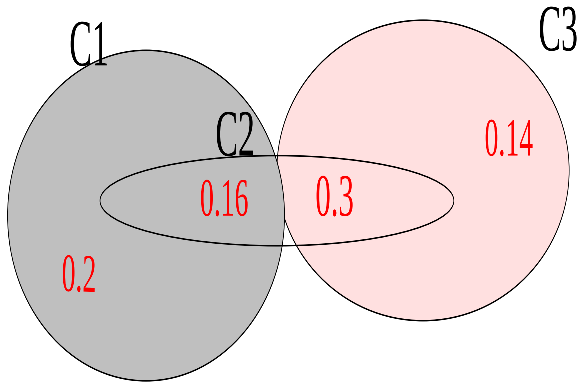

Figure 3. Sub-Optimality of our Paging Sequence Heuristic

Result 1 *The proposed greedy heuristic for problem* **(P2)** *has a worst-case approximation ratio of at least* 2.

In the next section, we shall see how the practical information theoretic location management framework [9] can be used to compute the conditional probabilities associated with a particular $\mathcal{MN}$ in a multi-system network.

5 Greedy Paging Algorithm for Integrated Location Management

Our basic framework for multi-system location management corresponds to the multi-system LeZi-Update algorithm described in [9]. The approach essentially views the movement pattern of the $\mathcal{MN}$ as a random sequence of random vectors:

$$\bar{\mathcal{X}}_n = \{\bar{X}_1, \bar{X}_2, \ldots, \bar{X}_n\}. \tag{8}$$

Here each element of the vector sequence $\bar{\mathcal{X}}_n$ is a random vector variable $\bar{X}_i$ of dimension N, corresponding to the vector valued location random variable described in Equation 1. As an example, the movement of $\mathcal{MN}$ can be represented by a sequence (string) of random vectors "$\bar{a}\bar{j}\bar{l}\bar{l}\bar{o}\bar{o}\bar{j}\bar{h}\bar{h}\bar{a}\bar{a}\bar{j}\bar{l}\bar{l}\bar{o}\bar{o}\bar{j}\bar{a}\bar{a}\bar{j}\bar{l}\bar{l}\bar{o}\bar{o}\bar{j}\bar{a}\bar{a}\bar{j}\bar{l}\bar{m}\ldots$", where each symbol in this string corresponds to a particular N-valued tuple representing the $\mathcal{MN}$'s location. The location update algorithm then uses the the Lempel-Ziv ($LZ78$) compression mechanism [13] to report updates [9, 4] in non-overlapping chunks. Due to the nature of the Lempel-Ziv entropy coding, the the network is assured that the $\mathcal{MN}$ always generates a fresh update if it encounters a new, previously unseen value of $\bar{X}$, i.e., if it moves into an area corresponding to a collection of cells that it has not previously encountered. With the sample random string enumerated earlier, the $\mathcal{MN}$ would break up its movement pattern into the distinct sub-strings (phrases) as follows: "$\bar{a}$, $\bar{j}$, $\bar{l}$, $\bar{l}\bar{o}$, $\bar{o}$, $\bar{j}\bar{h}$, $\bar{h}$, $\bar{a}\bar{a}$, $\bar{j}\bar{l}$, $\bar{l}\bar{o}\bar{o}$, $\bar{j}\bar{a}$, $\bar{a}\bar{j}$, $\bar{l}\bar{l}$, $\bar{o}\bar{o}$, $\bar{j}\bar{a}\bar{a}$, $\bar{j}\bar{l}\bar{m}$, $\ldots$,".

At the network (decoder) end, the encoded updates are first decoded and then stored as a symbol-wise context model using an efficient trie-based data structure. Our greedy paging heuristic then utilizes the information stored a dictionary (trie) which is built by decoding the encoded updates sent by the $\mathcal{MN}$. For details on the actual construction of the LeZi-Update trie in the multi-system network, see [9]. The important point for us here is that the joint distribution function of the $\mathcal{MN}$ is computed by the decoder *based solely on information that is stored in the trie.* To understand this, note that the trie is essentially a collection of symbols, where each symbol is a vector of dimension N and represents the $\mathcal{MN}$'s simultaneous location in a cell in each sub-network. The construction of the greedy paging sequence then uses the Prediction by Partial Match (PPM) style blending algorithm [13] to first construct the *unconditional* probability $Pr[\bar{\psi}]$, for each sequence $\bar{\psi}$ (also called phrase) of symbols. According to the principle of *insufficient reason* [8], the probability of an individual symbol (N-valued cell vector) is then computed based on the relative weights of symbols on these phrases. Formally, the unconditional probability $\rho(\bar{X}_k)$ of each vector-valued symbol can then be obtained as $\rho(\bar{X}_k) = \sum_{\bar{\psi}} \frac{\zeta(\bar{X}_k)}{\mathcal{L}(\bar{\psi})} \times \Pr[\bar{\psi}]$, where $\zeta(\bar{X}_k)$ denotes the number of occurrences of the symbol X_k in $\bar{\psi}$ and $\mathcal{L}(\bar{\psi})$ is the length of the sequence $\bar{\psi}$. After obtaining the unconditional vector probabilities, in order to estimate the unconditional probabilities for each *cell* within its own sub-network, we treat each element of a vector-valued symbol as having the corresponding probability mass, and then normalize these values over all the elements (cells) in the system. Thus, formally, the unconditional residence probability of cell $C^i_{S(i)}$ is given by:

$$\rho^i_{S(i)} = \sum_{k \ : \ c^i_{S(i)} \in \bar{X}_k} \rho(\bar{X}_k) \tag{9}$$

For example, consider the case where the trie-based traversal results in the following probability assignments: ($[C_1^2, C_2^1] = 0.3$, $[C_1^2, C_2^3,] = 0.15$, $[C_1^2, C_2^{\phi},] = 0.2$, and $[C_1^3, C_2^3] = 0.35$). We then obtain the individual cell residence probabilities as: ($\rho_1^2 = \frac{0.65}{1}$, $\rho_1^3 = \frac{0.35}{1}$, $\rho_2^1 = \frac{0.3}{1}$, $\rho_2^3 = \frac{0.5}{1}$, and $\rho_2^{\phi} = \frac{0.2}{1}$). Note that the joint cell probabilities (across the N sub-networks) provide enough information to construct the entire joint probability distribution function.

The computation of the paging sequence then proceeds as follows. We pick the first cell to be paged as the one having the highest unconditional probability (i.e., max $\rho^i_{S(i)}$). To construct the 2^{nd} element of the sequence, we first remove all the tuples containing the cell $c^1_{S(1)}$, leaving us with a reduced set of candidate cell tuples Set_2. For the tuples remaining in Set_2, we then compute the order-1 conditional probabilities by normalizing the probabilities of each tuple by the probability mass of the set, i.e., for each tuple $t \in Set_2$, compute $p_{t|\hat{1}} = \frac{\rho_t}{\sum_{k; \in \ Set_2} \rho_k}$. We then compute the order-1 conditional probability of each cell in Set_2 by adding up the normalized order-1 conditional probabilities of all tuples in Set_2 that contains the cell in question. For the 2^{nd} element of our greedy sequence, we pick the cell having the highest order-1 probability. We iterate over this process to construct the remaining members of the sequence. The formal description of the algorithm for computing the paging sequence via the greedy heuristic is provided in Figure 4.

Initialize $order = 0$, $j = 1$;
Initialize $set = \{C^i_{S(i)}\}$ such that $\rho^i_{S(i)} > 0$;
While ($set \neq$ null)
   For each $c^i_{S(i)}$, obtain
      $\rho^i_{S(i)}(order) = \sum_{k=1, \ c^i_{S(i)} \in X_k}^{|set|} \rho(\bar{X}_k)$;
   Set j^{th} element of Seq
      ($Seq[j]$)= $\arg \ \max \rho^i_{S(i)}(order)$;
   $order = order + 1$;
   Compute $newSet = \{\bar{X}_k\} : \bar{X}_k \in set, \ Seq[j] \notin \bar{X}_k$;
   $set = newSet$;
   For each $\bar{X}_k \in set$, set $\rho(\bar{X}_k) = \frac{\rho(\bar{X}_k)}{\sum_{l=1}^{|set|} \rho(\bar{X}_l)}$;
End-while;

Figure 4. Greedy Conditional-Probability Based Paging Sequence Computation Algorithm

Our examples on the construction of the joint distribution function was based on an underlying *movement-based update strategy*, where the $\mathcal{MN}$ generated a new symbol (vector value for the random sequence $\bar{\mathcal{X}}_n$ defined in Equation 8) only when it changed one of its N cellular coordinates. Note, however, that the joint distribution can be constructed, and the paging sequence appropriately defined, for any other symbol generation strategy, such as *timer-based* or *distance-based*. A timer-based strategy may be useful to weigh (have a higher probability mass on) combinations of cells where the $\mathcal{MN}$ residers longer, even though two distinct combinations may occur with the same frequency in the movement-based schemes. It should be emphasized that the use of the LeZi-Update based trie for computing the joint residence probabilities is independent of the underlying strategy for generating the samples on which the distribution is computed. The relative merits of different sample generation strategies is orthogonal to the main focus of this paper, and has been discussed extensively in [2].

6 Simulation Experiments

We now present simulation results obtained using a discrete-event simulation framework that we developed for studying the movement of a mobile user in a multi-system environment. A multi-system heterogeneous network topology, with different cell sizes for each sub-network, constitutes the heart of the simulation environment. Synthetic traces of user's activities are dynamically generated and fed in. The mobile user is assumed to have a set of deterministic and statistical patterns reflecting his/her life-style. Before discussing the simulation results, we describe the various parameters used in our study.

(i) The $4G$ network is assumed to consist of 3 distinct sub-network types: *satellite*, *PCS* and *WLAN*. The satellite (largest cell size) and PCS networks (intermediate cell sizes) have a cellular structure, with every cell having between 4–8 neighbors (6 on average). The 802.11 wireless LANs are generated as multiple (often disconnected) hot-spots, each with its own cellular layout. Assuming that PCS networks have the lowest update cost, the relative update costs for wireless LANs and satellite networks are respectively set to 3 and 10 times that of the PCS network. Similarly, assuming that wireless LANs have the lowest paging cost, the relative paging costs associated with the PCS and satellite networks are set to 4 and 9 times that of the wireless LANs.

(ii) The user has a home and a work-place randomly chosen from the multi-system network. This home and work-place both span multiple cells (in different sub-networks). The time spent by the user at every cell is normally distributed. The movement pattern of the user is assumed to differ between weekdays and weekends/holidays.

(iii) The call arrival process is Markov-modulated with three distinct states, each with its own Poisson arrival rate λ and normally distributed holding times with mean μ and variance σ. The states represent weekday day-times ($\lambda = 0.2$ calls/hr, $\mu = 10$ min, $\sigma = 3$ min), weekday evenings ($\lambda = 0.3$ calls/hr, $\mu = 20$ min, $\sigma = 5$ min), and weekends ($\lambda = 0.5$ calls/hr, $\mu = 30$ min, $\sigma = 7$ min).

(iv) The simulation results are based on an observation for a period of 12 weeks.

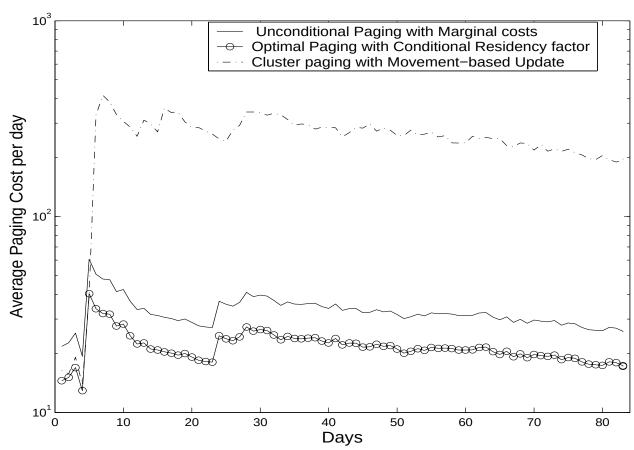

Figure 5. Improvement in Optimal Paging

We present the results of a typical study in order to compare the *paging* costs of our proposed conditional-probability based greedy heuristic with unconditional paging based on marginal costs and existing cluster paging based on movement-based location updates. In the movement-based update scheme, an update is always reported after a fixed number of cell changes. For cellular PCS networks, it has been shown that [2], as the movement threshold (the number of cell-changes between updates) increases, the update cost decreases but the paging cost increases. An average threshold of 3 was seen to achieve a reasonably fair compromise. Figure 5 demonstrates that our greedy heuristic paging strategy results in $\sim \frac{2}{3}$ and $\sim \frac{1}{10}$ of paging costs respectively, in comparison with the unconditional paging based on marginal costs and existing cluster paging based on movement-based location updates. This points out the importance of *conditional residence probabilities* arising from $\mathcal{MN}$'s simultaneous residence in various, overlapping cells, belonging to different subnetworks.

7 Conclusions

In this paper we have examined the problem of determining an optimal paging strategy in emerging multi-system $4G$ networks, where the mobile node $\mathcal{MN}$ may be simultaneously associated with multiple cells (belonging to different sub-networks). We have shown that the non-disjoint nature of the $\mathcal{MN}$'s residence events in different cells makes past results on optimal paging sequences in single-system networks inapplicable. In particular, we showed that the problem of determining the paging sequence that results in the lowest expected paging cost is NP-complete for the case of arbitrarily overlapping cellular structures.

We also showed that due to the potential for overlap, the paging cost associated with any particular sequence is a function of not the marginal residence probabilities in each cell, but instead of the conditional residence probabilities. Based on this observation, we designed an intuitively appealing greedy heuristic, where the next cell in the paging sequence is chosen to be one where the $\mathcal{MN}$ has the highest probability of residence (normalized by the paging cost in the cell), *conditioned on the mobile being simultaneously absent in all previous cells of the sequence*. This result is exploited by the profile-based, multi-system LeZi-Update framework to iteratively construct the greedy paging sequence.

Simulation experiments show that, for our chosen topology and movement statistics, our proposed greedy paging algorithm results in savings of almost 33% over an approach that merely pages cells in the decreasing order of their unconditional marginal residence probabilities. In our ongoing work, we are addressing the problem of optimal paging in the presence of constraints on the paging latency. However, as the unconstrained version of the problem illustrates, we should focus more on efficient heuristics than in attempting to find a globally optimal solution.

Acknowledgements: We gratefully acknowledge the comments of an anonymous reviewer for pointing out the incompleteness of an optimality proof in an original draft, and the assistance of Charu Agrawal of IBM Research for his suggestion that the optimization problem might be NP-complete.

References

1. I. F. Akyildiz and W. Wang, "A Dynamic Location Management Scheme for Next-Generation Multitier PCS Systems", *IEEE Transactions on Wireless Communications*, vol. 1, no. 1, pp. 178-189, Jan 2002.
2. A. Bar-Noy, I. Kessler and M. Sidi, "Mobile Users: to Update or Not to Update?", *ACM/Baltzer Wireless Networks* (WINET), vol. 1, no. 2, pp. 175-195, July 1995.
3. R. Berezdivin, R. Breinig and R. Topp, "Next-Generation Wireless Communications Concepts and Technologies", *IEEE Communications Magazine*, vol. 40, no. 3, pp. 108-116, March 2002.
4. A. Bhattacharya and S. K. Das, "LeZi-update: An Information-theoretic approach for personal mobility tracking in PCS networks," *ACM-Kluwer Wireless Networks* (WINET), vol. 8, no.2, pp.121-137, 2002.
5. E. Cohen, A. Fiat and H. Kaplan, "Efficient Sequences of Trials", *Proceedings of ACM/SIAM Symposium on Discrete Algorithms (SODA), 2003*, , pp. 737-746, Jan. 2003.
6. M. R. Garey, "Computers and intractability : a guide to the theory of NP-completeness", *W. H. Freeman Publishers*, 1979.
   T. H. Cormen, C. E. Leiserson and R. L. Rivest, "Introduction to Algorithms", *Prentice Hall Publishers.*
7. M. Feder, N. Merhav and M. Gutman, "Universal prediction of individual sequences," *IEEE Transactions on Information Theory*, vol. 38, no. 4, pp. 1258-1270, July 1992.
8. J. N. Kapur and H. K. Kesavan, *Entropy Optimization Principles with Applications*, Academic Press, 1992.
9. A. Misra, A. Roy and S. K. Das, "An Information-Theoretic Framework for Optimal Location Tracking in Multi-System 4G Wireless Networks", *To appear in IEEE InfoCom 2004*, Hong Kong, Mar. 2004.
10. G. P. Pollini and C.-L. I, "A Profile-based Location Strategy and its Performance, *IEEE Journal on Selected Areas in Communications*, vol. 15, no. 8, pp. 1415-1424, 1997.
11. C. Rose and R. Yates, "Minimizing the Average Cost of Paging under Delay Constraints", *Wireless Networks (WINET)*, vol. 1, no. 2, pp. 211-219, July 1995.
12. W. Wang and I. F. Akyildiz, "A Cost-Efficient Signaling Protocol for Mobility Application Part (MAP) in IMT-2000 Systems", *Proc. of ACM 7th Annual Intl. Conf. on Mobile Computing and Networking* (MobiCom), pp. 345-355, 2001.
13. J. Ziv and A. Lempel, "Compression of individual sequences via variable-rate coding," *IEEE Transactions on Information Theory*, vol. 24, no. 5, pp.530-536, Sept. 1978.