Locally Optimal Soft Handoff Algorithms

Rajat Prakash and Venugopal V. Veeravalli, *Senior Member, IEEE*

***Abstract*—The design of soft handoff algorithms for cellular radio systems is considered. The design problem is posed as a tradeoff between three metrics: the rate of handoffs, the mean size of the active set, and the link quality. It is argued that the algorithm that optimizes the tradeoff between these metrics is impractical. Hence, a locally optimal (LO) handoff algorithm is derived as a practical approximation to the optimal handoff algorithm. The LO algorithm is shown to yield a significantly better tradeoff than the static threshold handoff algorithm used in second-generation code-division multiple-access (CDMA) systems. It is also shown that the dynamic threshold algorithm, which is an ad hoc algorithm proposed for third-generation CDMA systems, achieves nearly the same performance as the LO algorithm. Thus, an analytical justification is developed for the dynamic threshold algorithm. Further, handoff algorithm design is separated into independent design problems on the forward and reverse links. The forward link LO algorithm is shown to be computationally intensive but is also shown to be closely approximated by the simpler reverse link LO algorithm.**

***Index Terms*—Cellular systems, code-division multiple access (CDMA), handover.**

I. Introduction

THE problem of soft handoff arises in a cellular communication system where the mobile can communicate with multiple base stations at the same time. The set of base stations with which the mobile communicates at a given time is called the *active set*. As the mobile position and the system traffic load conditions change, the active set needs to be changed in order to maintain acceptable signal quality. This change in the active set is the soft handoff event and is governed by the soft handoff algorithm.

For code-division multiple-access (CDMA) systems employing diversity reception, the ability to be in soft handoff, i.e., the ability to have more than one base station in the active set, results in added diversity and improved signal quality [1, Section 5.5]. The consequent capacity and cell coverage gains have been analyzed by Viterbi *et al.* [2] and Sendonaris and Veeravalli [3]. However, these works do not deal with the design of algorithms to select the active set. Our focus is on the design of soft handoff algorithms to select the active set.

In contrast to soft handoff, the design of hard handoff algorithms has received much attention. The optimal hard handoff algorithm design problem was studied by Rezaiifar *et al.* [4], Asawa and Stark [5], and Veeravalli and Kelly [6]. Since optimal handoff algorithms are impractical, a locally optimal (LO) approach was suggested [6] to get suboptimal but practical hard handoff algorithms. In further work on LO algorithms, Prakash and Veeravalli [7] studied the adaptation of LO hard handoff algorithms to changing system parameters, and Akar and Mitra [8] applied the LO technique to handoff delay optimization.

Optimizing soft handoff is still largely an open problem with most of the previous work focusing on *ad hoc* analyses. An overview of recent work on soft handoff has been provided by Wong and Lim [9]. The second-generation IS-95 standard [10] recommends the use of an ad hoc static threshold soft handoff algorithm. Zhang and Holtzman [11] have provided tools to study the performance of the static threshold algorithm but have not considered techniques for handoff algorithm design. Asawa and Stark [5] have applied a limited lookahead approach to the design of soft handoff algorithms and demonstrated improvement over the static threshold algorithm. To improve the performance of the static threshold handoff algorithm, the third-generation cdma2000 [12, Section 3.2.3.3] standard recommends the *dynamic threshold* handoff algorithm. In this paper, we introduce an LO soft handoff algorithm and compare its performance with other algorithms.

The primary objective of a soft handoff algorithm is to provide good signal quality. Signal quality can be improved by including more base stations in the active set, but this comes at the cost of increased use of system resources. To lower the active set size, one option is to frequently update the active set to maintain, at each time instant, the smallest active set with sufficient signal quality. However, frequent updates or handoffs bring with them switching costs. Thus, as has been seen earlier [5], [11], a tradeoff exists among the following three metrics: the rate of active set updates, the mean size of the active set, and the average signal quality.

A handoff algorithm is said to be optimal if it attains the best tradeoff amongst the class of all handoff algorithms. The procedure for the design of optimal soft handoff algorithms is similar to that for hard handoff [5], [6] and uses a cost function that depends on many steps into the future. The design of optimal handoff algorithms requires a model for the future trajectory of the mobile. We assume that information about the future trajectory is not readily available, making it impractical to use optimal algorithms. Another drawback of optimal algorithms is that their design is computationally intractable.

In contrast to the optimal algorithms, the LO approach gives a practical solution by minimizing a one-step lookahead cost function. The LO algorithm has the advantage of not requiring any complicated models for the trajectory of the mobile. This is so because the locally optimal algorithm operates on a small

Manuscript received June 13, 2001; revised April 5, 2002. This work was supported in part by the Office of Naval Research under Grant N0014-97-1-0823 and by Nortel Networks' Global External Research (both through Cornell University) and by the National Science Foundation CAREER/PECASE under Grant CCR-0049089 (through the University of Illinois). This work was presented in part at the Vehicular Technology Conference 2000, Tokyo, Japan, May 2000.

The authors are with the School of Electrical and Computer Engineering and the Coordinated Sciences Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: rprakas1@uiuc.edu; vvv@uiuc.edu).

Digital Object Identifier 10.1109/TVT.2003.808810

time-scale, over which the motion of the mobile can be approximated by a straight line. Simulation results show that the LO algorithm offers significant improvement in performance over the ad hoc static threshold algorithm of IS-95 [10]. In contrast, the ad hoc dynamic threshold algorithm of cdma2000 [12] performs close to the LO algorithm. Furthermore, the structure of the dynamic threshold algorithm resembles that of the LO algorithm. Since the LO algorithm has a theoretical basis, the above observations give a theoretical justification for the use of the dynamic threshold algorithm.

Another issue we explore is the use of different active sets on the forward and reverse links, i.e., base stations that receive a given mobile's signal may not transmit to it, and vice versa. This asymmetry may be attractive because forward and reverse link active sets require different resources, which may have different costs, e.g., a large forward link active set contributes to forward link interference, but a large reverse link active set does not influence interference. We separate the handoff algorithm design problem into separate problems for the forward and reverse links. We show that although the structure of the forward link LO algorithm is more complex than the reverse LO link algorithm, the simpler reverse link LO algorithm gives satisfactory performance on the forward link. Thus, a separate handoff algorithm is not necessary for the forward link.

The rest of this paper is structured as follows. The channel model and the definition of the soft handoff problem are developed in Section II. Performance metrics to measure the performance of soft handoff are constructed in Section III-A. In Section III-B, the simulation environment is described and simulation results for a static threshold soft handoff algorithm are given. The locally optimal soft handoff algorithm for the reverse link is developed in Section IV-A. In Section IV-B, the performance of the LO algorithm is compared with the static and dynamic threshold handoff algorithms. In Section IV-C, the forward link LO algorithm is constructed and its performance analyzed. Conclusions and comments are presented in Section V.

II. Preliminaries

In this section, we describe the channel model and define the soft handoff problem. We adopt a discrete time model with sampling time t_s. As is usual in discrete time models, we refer to sample instant k simply as time k. We consider mobile assisted handoff, where at time k, the mobile transmits the pilot signal strength measurements of base stations in the candidate set to the controlling base station. These measurements are the primary source of information for the soft handoff algorithm. In addition to the pilot signal strength measurements, handoff may also be based on the geographic traffic load pattern. Though we do not consider handoff based on traffic load in this paper, in Section V, we comment on how traffic-load information may be incorporated in the handoff algorithm.

For a given mobile trajectory, we focus on the part where there are B base stations BS-1, ..., BS-B in the candidate set. These base stations have pilot signals strong enough to make them potential candidates for the active set. The set of base stations that transmit to the given mobile at time k is the forward active set $A_{\text{for},k}$. For example, $A_{\text{for},k} = \{2,3\}$ means that BS-2 and BS-3 are transmitting to the mobile at time k. Similarly, the reverse-link active set $A_{\text{rev},k}$ is the set of base stations that listen to the traffic signal of the given mobile. As stated earlier, we allow for $A_{\text{for},k} \neq A_{\text{rev},k}$.

Both the reverse and forward link active sets are selected based on the pilot signal strengths, which are influenced by short-term and long-term fading. Handoff algorithms cannot respond to short-term fading because handoff involves the setting up of connections between base stations and the delay in connection setup is often significantly more than the time scale of short-term fading. Thus, we assume that the handoff algorithm responds only to long-term fading and that the pilot signal strength is averaged to remove the effect of short-term fading.

Let $X_{k,i}$ (dBm) be the averaged pilot signal strength from BS-i at time k. Given the pilot source strength P_P and the transmitted traffic signal strength P_{for} (both in dBm), the traffic channel signal strength $\tilde{X}_{k,i}$ received by the mobile from base station i at time k is

$$\tilde{X}_{k,i} = X_{k,i} + P_{\text{for}} - P_P. \tag{1}$$

Thus, the pilot signal strength $X_{k,i}$ can be used to determine the signal strength $\tilde{X}_{k,i}$ on the traffic channel from each base station. On the forward link, maximal ratio combining [1, Section 5.5.3] is used to combine the signals from base stations in the active set $A_{\text{for},k}$. Given an interference level I_k (dBm) at the mobile, the forward link signal-to-interference ratio (SIR) $\gamma_{\text{for},k}$ after maximal-ratio combining is given in dB by

$$\gamma_{\text{for},k} = 10\log\left(\sum_{i\in A_{\text{for},k}} 10^{\tilde{X}_{k,i}/10}\right) - I_k. \tag{2}$$

This can be rewritten as

$$\gamma_{\text{for},k} = X_{\text{for},k} + P_{\text{for}} - P_P - I_k \tag{3}$$

where

$$X_{\text{for},k} = 10\log\left(\sum_{i\in A_{\text{for},k}} 10^{X_{k,i}/10}\right). \tag{4}$$

Here, $X_{\text{for},k}$ can be regarded as the effective pilot signal strength from all base stations in the forward active set. We have assumed a rake receiver [1, Section 8.3.1] with enough fingers for all paths to base stations in $A_{\text{for},k}$. Without this assumption, the SIR at the mobile would depend on the algorithm for selecting the paths to be combined at the rake receiver.

For the reverse link, if the combined interference and noise level at BS-i is $I_{k,i}$ (dBm), then the SIR at BS-i is given by

$$\gamma_{k,i} = X_{k,i} + P_{\text{rev}} - P_P - I_{k.i} \tag{5}$$

where P_{rev} is the transmit power of the mobile. Typically, selection diversity is used on a frame-by-frame basis on the reverse link, i.e., the signal at the base station with the largest $\gamma_{k,i}$ is more probably used for demodulation. Thus, the reverse link SIR is

$$\gamma_{\text{rev},k} = \max_{i\in A_{\text{rev},k}} (X_{k,i} - I_{k,i}) + P_{\text{rev}} - P_P. \tag{6}$$

This shows that the pilot signal strength can be used to determine the signal quality on the reverse link.

The pilot signal strength from BS-i undergoes shadow fading according to the following model [1, Section 2.4]:

$$X_{k,i} = \bar{P}_{k,i} + Z_{k,i} \tag{7}$$

$$\bar{P}_{k,i} = \mu_i - 10\eta_i \log d_{k,i} \tag{8}$$

where $\bar{P}_{k,i}$ (dBm) is the local mean pilot power and $Z_{k,i}$ (dBm) is the shadow fading at sampling time k. The local mean pilot power varies log-linearly with the distance $d_{k,i}$ from BS-i, and μ_i and η_i are constants.

To further develop the statistics of the received pilot power, we concentrate on a part of the trajectory where the mobile is moving on a straight line with fixed velocity v and use a first-order autoregressive (AR-1) model for the autocorrelation function of $\{Z_{k,i}\}$. Gudmundson [14] showed that the AR-1 model matches field measurements well. The AR-1 model has been used earlier in the analysis of handoff algorithms [4], [6].

Under the AR-1 model, the shadow fading autocorrelation function is given by

$$E\left[Z_{k,i} Z_{k+m,i}\right] = \sigma_i^2 a_i^{|m|}.$$

Here, σ_i^2 is the shadow fading *variance* and a_i is the correlation coefficient of $\{Z_{k,i}\}$, i.e.,

$$a_i = \exp\left(\frac{-vt_s}{\bar{d}_i}\right) \tag{9}$$

where $\bar{d}_i$ is the shadow fading *correlation distance*.

Under this model, it is easy to check (see, e.g., [6]) that the distribution of $X_{k+1,i}$ conditioned on $X_{k,i}$ is independent of earlier received power samples and is described completely by its conditional mean and variance

$$\begin{aligned} \mathrm{E}\left[X_{k+1,i} \mid X_{k,i} \ldots X_{1,i}\right] &= \mathrm{E}\left[X_{k+1,i} \mid X_{k,i}\right] \\ &= \bar{P}_{k+1,i} + a_i\left(X_{k,i} - \bar{P}_{k,i}\right) \end{aligned} \tag{10}$$

$$\begin{aligned} \mathrm{Var}\left[X_{k+1,i} \mid X_{k,i} \ldots X_{1,i}\right] &= \mathrm{Var}\left[X_{k+1,i} \mid X_{k,i}\right] \\ &= \left(1 - a_i^2\right)\sigma_i^2. \end{aligned} \tag{11}$$

We assume that the shadow fading processes from base stations i and j are correlated in such a way that the random variables $Z_{k,i}$ and $Z_{k,j}$ have correlation ρ for $i \neq j$.

The information available to the handoff algorithm is the information vector $\mathcal{I}_k$, which consists of all the past measured pilot strengths $X_{k,i}$, all the past active sets, and the past combined interference and noise levels at the mobile as well as the base stations, i.e.,

$$\mathcal{I}_k = \bigcup_{n=1}^{k} \left\{ \cup_{i=1}^{B} X_{n,i}, A_{\mathrm{for},n}, A_{\mathrm{rev},n}, I_n, \cup_{i=1}^{B} \{I_{n,i}\} \right\}. \tag{12}$$

In practice, however, information may be available only about the recent past (rather than the entire past), and this limited information is enough for the algorithms we study.

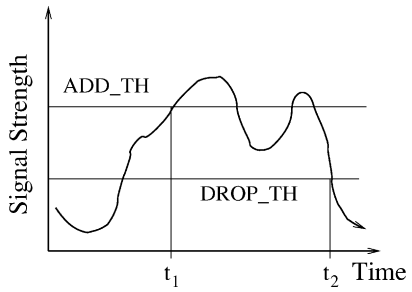

Fig. 1. Operation of the static threshold soft handoff algorithm. The mobile is in soft handoff during the time interval t_1 to t_2.

The handoff algorithm selects the active set at the next time instant using a decision function ϕ_k

$$A_{k+1} = \phi_k\left(\mathcal{I}_k\right). \tag{13}$$

The resulting handoff policy is the collection Φ of the handoff decision functions ϕ_k at all sample instances k. If the decision functions do not vary with time, the handoff policy is said to be stationary. To support different active sets on the forward and reverse links, separate decision functions $\phi_{\mathrm{for},k}$ and $\phi_{\mathrm{rev},k}$ as well as handoff policies Φ_{for} and Φ_{rev} can be defined for the forward and reverse links, respectively.

An example of a soft handoff algorithm is the static threshold handoff algorithm [10], which uses the same active set on the forward and reverse links. This algorithm is characterized by two parameters: an add threshold `ADD_TH` and a drop threshold `DROP_TH`. The difference between these two parameters is the *hysteresis level*. The working of this handoff algorithm is illustrated in Fig. 1. When the pilot from a base station goes below the drop threshold, the base station is removed from the active set. When the pilot from a base station goes above the add threshold, the base station is added to the active set. For the static threshold handoff algorithm, the decision functions ϕ_k are the same for all k.

The task of the soft handoff algorithm is to select an appropriate handoff policy Φ (if necessary separate policies Φ_{for} and Φ_{rev}) to optimize the values of the performance metrics introduced in the next section.

III. Measuring Performance

In this section, we describe performance metrics that can be used to characterize the performance of soft handoff algorithms. In Section IV, we use these metrics to compare the performance of different handoff algorithms.

A. Performance Metrics

As stated in Section I, the performance of a soft handoff algorithm can be measured by three performance metrics: rate of handoffs, average active set size, and the average signal quality. The performance metrics are functions of both the handoff policy for the link under consideration and the system parameters $\mathcal{S}$ (e.g., mobile velocity v and shadow fading correlation a_i). In case different active sets are used on the forward and reverse links, each of the three metrics will be replaced by two different metrics, one for each link.

Expressions for the three performance metrics are developed next.

1) Rate of handoffs

$$\lambda_H(\Phi,\mathcal{S}) = \mathrm{E}\left[\frac{1}{N}\sum_{k=1}^{N} \mathbb{1}_{\{A_k \neq A_{k-1}\}}\right] \tag{14}$$

where $\mathbb{1}$ is the indicator function, which takes values of one or zero depending on whether the argument is true or false. A soft handoff is said to have occurred at time k if $A_{k+1} \neq A_k$. The metric λ_H represents the switching load associated with changes in the active set.

2) Average active set size

$$\lambda_A(\Phi,\mathcal{S}) = \mathrm{E}\left[\frac{1}{N}\sum_{k=1}^{N} |A_k|\right]. \tag{15}$$

The metric λ_A represents the additional channel card and network backbone requirements of a mobile in soft handoff. During soft handoff, signals need to be carried between base stations in the active set, causing additional traffic on the network backbone. In addition, because of the extra base stations transmitting to a mobile in soft handoff, the interference level seen by other mobiles on the forward link can increase. If sophisticated power-control schemes are used to take advantage of the forward link diversity gain in soft handoff, it may be possible to reduce (instead of increase) the total interference seen by other mobiles. However, we assume that even if a large active set results in interference reduction, the resulting gains do not offset the the extra cost of increased network load and channel card usage. Thus, we use a cost function that increases with increasing active set size.

3) Average signal quality (quantified by the rate of link degradation events)

$$\lambda_{\mathrm{LD}}(\Phi,\mathcal{S}) = \mathrm{E}\left[\frac{1}{N}\sum_{k=1}^{N} \mathbb{1}_{\{\text{Link Degradation at time}k\}}\right]. \tag{16}$$

The metric λ_{LD} measures the signal quality as the fraction of time for which the link is in a degraded state. The link degradation event is defined in detail next.

First, consider the forward link degradation (FLD) event. The FLD event occurs if the SIR $\gamma_{\mathrm{for},k}$ at the mobile rake combiner output is below an SIR threshold $\tilde{\Delta}_{\mathrm{for}}$. We assume a perfect forward link power-control algorithm that sets the transmit power of all base stations to a maximum level $P_{\mathrm{for,max}}$ whenever an FLD is imminent. This is a reasonable assumption because power control operates on a much faster time-scale than handoff. Given the forward active set $A_{\mathrm{for},k}$, from (2) it follows that an FLD event occurs when the effective forward link pilot strength $X_{\mathrm{for},k}$ in (4) goes below a threshold Δ_{for}, i.e.,

$$X_{\mathrm{for},k} < \Delta_{\mathrm{for}} \tag{17}$$

where

$$\Delta_{\mathrm{for}} = \tilde{\Delta}_{\mathrm{for}} - P_{\mathrm{for,max}} + P_P - I_k. \tag{18}$$

Note that under the assumption that interference I_k does not vary with time, the threshold Δ_{for} is constant in time.

A reverse link degradation (RLD) occurs if the reverse link SIR $\gamma_{\mathrm{rev},k}$ goes below a threshold $\tilde{\Delta}_{\mathrm{rev}}$. Under a perfect reverse link power control algorithm, the mobile transmits with peak power $P_{\mathrm{rev,max}}$ whenever an RLD is imminent. Using (6), it follows that a RLD occurs when

$$\max_{i\in A_{\mathrm{rev},k}} (X_{k,i} - I_{k,i}) < \tilde{\Delta}_{\mathrm{rev}} + P_P - P_{\mathrm{rev,max}}. \tag{19}$$

If the interference levels are fixed at all base stations and do not vary with time, i.e.,

$$I_{k',i} = I_{\mathrm{rev}}, \quad i = 1,\ldots B, k' = 1,\ldots,k$$

then the RLD event reduces to

$$\left\{\max_{i\in A_{\mathrm{rev},k}} (X_{k,i}) < \Delta_{\mathrm{rev}}\right\} \tag{20}$$

where

$$\Delta_{\mathrm{rev}} = \tilde{\Delta}_{\mathrm{rev}} + P_P - P_{\mathrm{rev,max}} + I_{\mathrm{rev}}.$$

It should be noted that the link degradation event defined above is not the same as a frame error event. The frame duration is typically small (e.g., 20 ms in IS-95) and frame error events are influenced by slow as well as fast fading. This is in contrast with link degradation events, which depend only on slow fading. Even though the two events are not exactly the same, both the frame error rate and the link degradation rate are metrics of average signal quality. As argued in Section II, the handoff algorithm cannot respond to fast fading and must base its decisions only on slow fading levels. For this reason, we use the link degradation rate and not the frame error rate as a metric of signal quality.

In the case where separate handoff algorithms are used on the forward and reverse links, it can be seen that there will be three metrics for the forward link and three separate metrics for the reverse link. For the forward link, the description of the three metrics is similar to that given in (14), (15), and (16), with Φ, A_k, and LD replaced by Φ_{for}, $A_{\mathrm{for},k}$, and FLD, respectively. Metrics for the reverse link can be obtained similarly. Further, the metrics for one link will depend only on the handoff algorithm being used on that link and not on the algorithm on the other link. Thus, the analysis of the forward and reverse link handoff algorithms can be carried out separately by studying the tradeoff between the three corresponding metrics.

If the system architecture constrains the active set to be the same for the forward and reverse links, the rate of handoff and the active set size metrics will be the same for both links. The link quality metrics, however, would be different for each link. In this case, either the forward or the reverse link can be considered constraining for the system capacity, and the quality of the constraining link can be considered to be the effective link quality metric. Handoff algorithms can then be analyzed in terms of the tradeoff among the active set size, rate of handoffs, and effective link quality.

B. Tradeoff Surfaces and Handoff Algorithm Design

In this section, we define the tradeoff surface and demonstrate its role in handoff algorithm design. In Section III-A, the per-

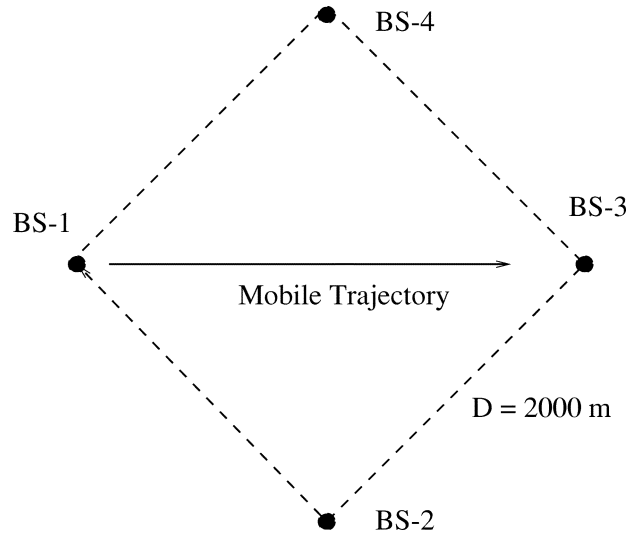

Fig. 2. Mobile trajectory.

formance metrics $(\lambda_H, \lambda_A, \lambda_{\rm LD})$ were shown to be a function of $(\Phi, \mathcal{S})$. In this section, we assume that the system parameters $\mathcal{S}$ are fixed. Then, the performance metrics are a function just of the policy Φ. Given a handoff algorithm, the handoff policy Φ depends on the handoff algorithm parameters ξ. For example, for the static threshold handoff algorithm (Fig. 1), $\xi = (\texttt{ADD\_TH}, \texttt{DROP\_TH})$ determines Φ. The tradeoff surface describes the range of performance a given handoff algorithm offers as the parameters ξ are changed.

To define the tradeoff surface, we represent the performance metrics $(\lambda_H, \lambda_A, \lambda_{\rm LD})$ as a point in a three-dimensional space (each dimension corresponding to a performance metric). The locus of operating points attained by varying ξ is defined to be the tradeoff surface for the given handoff algorithm.

Given a handoff algorithm, along with parameters ξ, the performance metrics in (14)–(16) are difficult to compute analytically. In the absence of any analytical techniques, we follow previous works [5], [11] in resorting to simulations. We consider a simulation environment where the mobile traverses a trajectory in the vicinity of $B = 4$ base stations arranged on the vertices of a square (Fig. 2), with the maximum active set size limited to three.

Fig. 3 shows a tradeoff surface for the static threshold handoff algorithm under the simulation parameters in Table I. To represent three-dimensional surfaces on paper, we show top and side views. The tradeoff between different metrics is illustrated by the tradeoff surface. For example, it is possible to reduce $\lambda_{\rm LD}$ by exploiting greater diversity at the expense of a larger λ_A. Also, it is possible to reduce λ_A at the expense of a larger λ_H by frequently updating the active set to maintain the smallest sufficient active set at each time instant.

The effect of the threshold and hysteresis levels on the operating points can be studied using the tradeoff surface. Consider points A, B, C, and D in Fig. 3. For a high add threshold (point C), both λ_A and λ_H are low, though at the expense of a high $\lambda_{\rm LD}$. Given a low add threshold (points A, D), increasing the hysteresis level reduces λ_H while maintaining a low $\lambda_{\rm LD}$ and a high λ_A. At point B, a low hysteresis level results in a low $\lambda_{\rm LD}$ at the expense of a high λ_H.

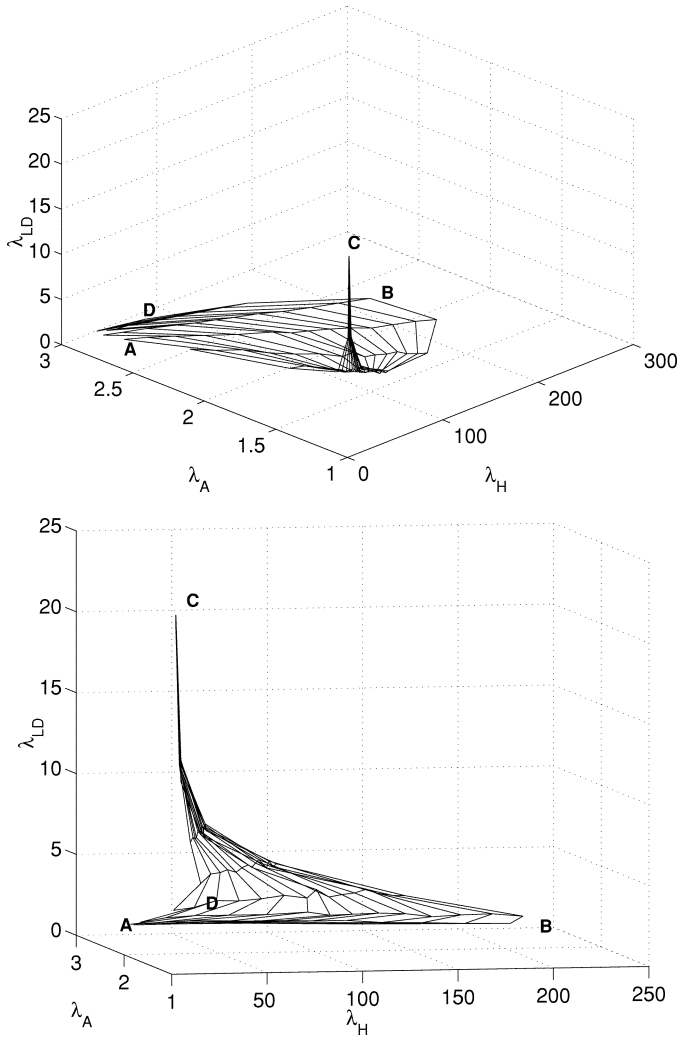

Fig. 3. Tradeoff surface for the static threshold soft handoff algorithm. Rates are measured with respect to time T required to traverse the trajectory.

TABLE I
PARAMETERS USED FOR ALL SOFT HANDOFF SIMULATIONS

D= 2000 meters,	distance from BS-1 to BS-2
μ_i= 108 dBm,	base station signal strength
η_i= 3,	path loss exponent
σ_i= 8 dB,	shadow fading std. dev.
v= 10 m/s,	mobile velocity
$\bar{d}$= 30 meters,	correlation distance
Δ= 0 dBm,	threshold of link degradation
t_s= 0.5s,	sample time
M= 20,	window length for estimation
ρ = 0.3	correlation between shadow fading from different base stations
T = 280s	time it takes to complete one trip (as shown in Fig. 2)

Tradeoff surfaces have two important uses. Given a handoff algorithm, the tradeoff surface may be used to select a desirable operating point and thus select the handoff algorithm parameters. A desirable operating point is one that achieves the given system design goals, e.g., to minimize a cost function of the three metrics, or attain a strict bound on the link quality metric

λ_{LD}. An ad hoc definition of a desirable operating point is the knee region of the tradeoff surface [13]. The other use of tradeoff surfaces is for comparing the performance of different handoff algorithms: the lower the tradeoff surface, the better the algorithm.

The definition of tradeoff surfaces we gave earlier was for handoff algorithms with a two-dimensional internal parameter ξ. For higher dimensional ξ, the locus of operating points as ξ is varied can cover a volume rather than a surface in the performance metric space. In such a case, the tradeoff surface may be defined to be the lower envelope of all attainable operating points. Alternatively, a simple heuristic would be to fix all but two parameters in ξ to obtain a tradeoff surface. By selecting different values for the fixed parameters, different tradeoff surfaces can be obtained and the "best" of these designated as the actual tradeoff surface.

IV. Design of Soft Handoff Algorithms

In the previous section, we identified techniques for measuring the performance of soft handoff algorithms. In this section, we consider the design of soft handoff algorithms. As was mentioned in Section III-A, we have independent design problems on the forward and reverse links. We begin with handoff algorithm design for the reverse link.

A. The LO Soft Handoff Algorithm for the Reverse Link

For hard handoff, the LO algorithm was developed as an alternative to the impractical optimal hard handoff algorithm [6]. In the same spirit, we first consider the optimal soft handoff algorithm and focus on the reverse link in this section. The Bayesian cost function for soft handoff has two (relative) cost parameters c_A and c_H. Parameter c_A is the cost of maintaining one extra member in the active set, while c_H is the cost of handoff. These costs are relative to a cost of one unit for a link degradation event. The Bayes cost under a policy Φ_{rev} and system parameters $\mathcal{S}$ is given by

$$J(\Phi_{\mathrm{rev}},\mathcal{S}) = \lambda_{\mathrm{LD}}(\Phi_{\mathrm{rev}},\mathcal{S}) + c_H\lambda_H(\Phi_{\mathrm{rev}},\mathcal{S}) + c_{\mathrm{A}}\lambda_A(\Phi_{\mathrm{rev}},\mathcal{S}). \quad (21)$$

The optimal soft handoff algorithm is one that minimizes the Bayes cost and can be obtained using dynamic programming (DP). To solve the DP problem, the active set at time k should be selected to minimize the cost incurred several time steps into the future. Since the cost function depends on the trajectory of the mobile, computation of the DP solution requires a (stochastic or deterministic) model for the mobile's future trajectory [4], [6]. Such a model may not be available in the system. Furthermore, numerical solution of the DP problem is difficult because the size of the state vector is large (equal to the number of entries in the candidate set). For these reasons, the optimal algorithm is impractical.

The LO algorithm overcomes the deficiencies of the optimal algorithm by minimizing an incremental cost function $J_{\mathrm{rev,incr}}$, which is in effect a one-step lookahead Bayes cost, i.e.,

$$J_{\mathrm{rev,incr}}(A_{\mathrm{rev},k+1}) = \mathbb{1}_{\{\mathrm{RLD\ at\ time}k+1\}} + c_A|A_{\mathrm{rev},k+1}| + c_H\mathbb{1}_{\{\mathrm{Reverse\ link\ handoff\ at\ time}k\}}. \quad (22)$$

The goal of the LO algorithm at time k is to select an active set to minimize the expectation of the incremental cost function, i.e.,

$$A^*_{\mathrm{rev},k+1} = \arg\min_{A_{\mathrm{rev},k+1}} \mathrm{E}\left[J_{\mathrm{rev,incr}}(A_{\mathrm{rev},k+1})\,|\mathcal{I}_k, A_{\mathrm{rev},k+1}\right]. \quad (23)$$

Thus, the LO algorithm does not need a model for the mobile's future trajectory; and, as we will show below, it is easily implementable.

The implementation of the LO algorithm involves the evaluation of the expectation in (23) for all possible $A_{\mathrm{rev},k+1}$. The number of possibilities may be large depending on the value of B. The following guidelines help to narrow down the possibilities. An incoming base station must be the strongest of those in the neighbor list, and an outgoing base station must be the weakest one in the active set. In addition, system constrains may further reduce the number of possible $A_{\mathrm{rev},k}$ choices that need to be considered. For example, hard handoff from BS-1 to BS-2, i.e., an event where the active set changes from $\{1\}$ to $\{2\}$ may be disallowed. Any such system constraints can be incorporated into the LO algorithm by considering only the valid $A_{\mathrm{rev},k}$ choices in the minimization (23).

To give a concrete form to the objective function of the minimization (23), consider each of the terms in (22) individually. The conditional expectation of the last two terms is evaluated immediately from the knowledge of $A_{\mathrm{rev},k}$ and $A_{\mathrm{rev},k+1}$. For the term corresponding to the RLD event (20), we are interested in evaluating the expectation

$$E\left[\mathbb{1}_{\{\mathrm{RLD\ at\ time}k+1\}} \mid \mathcal{I}_k, A_{\mathrm{rev},k+1}\right] = P\left(\left\{\max_{i\in A_{\mathrm{rev},k+1}}(X_{k+1,i}) < \Delta_{\mathrm{rev}}\right\} \mid \mathcal{I}_k, A_{\mathrm{rev},k+1}\right). \quad (24)$$

The distribution of $X_{k+1,i}$, conditioned on $\mathcal{I}_k$ can be obtained using (10) and (11). The probability in (24) takes the following values for different active set sizes.

1) When the active set size is one, e.g., $A_{\mathrm{rev},k+1} = \{i\}$, then (24) reduces to

$$P\left(X_{k+1,i} < \Delta_{\mathrm{rev}} \mid \mathcal{I}_k\right) = Q\left(\frac{E\left[X_{k+1,i} \mid X_{k,i}\right] - \Delta_{\mathrm{rev}}}{\sigma_i\sqrt{1-a_i^2}}\right) \quad (25)$$

where $Q(x) = (1/\sqrt{2\pi})\int_x^\infty e^{-y^2/2}dy$. The equation above is obtained in a manner similar to the LO hard handoff algorithm [6].

2) When the active set size is two, e.g., $A_{k+1} = \{i,j\}$, then (24) reduces to

$$P\left(\max(X_{k+1,i}, X_{k+1,j}) < \Delta_{\mathrm{rev}} \mid \mathcal{I}_k\right) = L\left(\frac{E\left[X_{k+1,i} \mid X_{k,i}\right] - \Delta_{\mathrm{rev}}}{\sigma_i\sqrt{1-a_i^2}}, \frac{E\left[X_{k+1,j} \mid X_{k,j}\right] - \Delta_{\mathrm{rev}}}{\sigma_j\sqrt{1-a_j^2}}, \rho\right). \quad (26)$$

Here, the L function is the cumulative distribution function for bivariate Gaussian random variables

$$L(y_1,y_2,\rho) = P(Y_1 < y_1, Y_2 < y_2), \quad \text{where } Y_1, Y_2 \sim \mathcal{N}_2(0,0,1,1,\rho).$$

Tables and evaluation methods for L functions can be found in the literature [15], [16]. We evaluate the L functions numerically.

3) When the active set size is three, e.g., $A_{k+1} = \{i, j, m\}$, (24) reduces to

$$\mathrm{P}\left(\{\max\left(X_{k+1,i}, X_{k+1,j}, X_{k+1,m}\right) < \Delta_{\mathrm{rev}}\} \mid \mathcal{I}_k\right). \quad (27)$$

This probability is evaluated numerically using the distribution function of trivariate Gaussian random variables [15, ch. 36].

The expected incremental cost need not be evaluated for larger active sets because the maximum allowed active set size in our simulations is three. In the case where larger active sets are allowed, numerical evaluation of the link degradation probability is difficult. A possible approximation can be obtained by assuming independence between the fading on different base stations.

The evaluation of the reverse link degradation probabilities in (25)–(27) requires the statistics of the Gaussian random variable $X_{k+1,i}$ conditioned on $\mathcal{I}_k$. These statistics are given by (10) and (11); however, $\bar{P}_{k,i}$ and $\sigma_i\sqrt{1-a_i^2}$ may be unknown, making it impossible to evaluate the required statistics perfectly. To get around this problem, we use $X_{k,i}$ as an estimator for $E[X_{k+1,i} \mid X_{k,i}]$ and construct the estimators $\hat{\sigma}$ and $\hat{a}$ for σ and a, respectively. These estimators were originally developed for hard handoff algorithms [7], [17], where their accuracy was studied.

For completeness, we give the following expressions that show that the estimators $\hat{\sigma}$ and $\hat{a}$ can be constructed from a size M history of pilot signal strength samples

$$\hat{\sigma} = \sqrt{C + \frac{R}{2}}$$

$$\hat{a} = \frac{C}{(\hat{\sigma})^2}$$

where C and R are given by

$$C = \frac{1}{M-1}\sum_{k=1}^{M-1}\left(X_{k,i} - \bar{X}\right)\left(X_{k+1,i} - \bar{X}\right)$$

$$R = \frac{1}{M-1}\sum_{k=1}^{M-1}\left(X_{k,i} - X_{k+1,i}\right)^2$$

with $\bar{X} = 1/M\sum_{k=1}^{M} X_{k,i}$..

Note that the nature of the LO algorithm allows for the use of any other estimators. In this paper, the focus is not on the design of estimators, but rather on the design of the handoff algorithm. Therefore, we use the estimators from [7] in our simulations.

Next, we determine the decision regions for the reverse link LO algorithm. The decision region is the region in the signal strength space where the handoff algorithm generates a handoff. A handoff algorithm is completely specified by decision regions for all possible transitions A_k to A_{k+1}. Fig. 4 shows the decision region boundary for two handoff scenarios: the addition of one base station to an active set with size one and the removal of one base station from an active set with size two. If more than two base stations are involved in the handoff, the decision region is higher dimensional and difficult to represent on paper.

To implement the LO algorithm, the decision rule (23) need not be evaluated at each decision instant. The stored decision region can be used to find the handoff decision instead. Further simplification can be achieved by storing a piecewise linear approximation to the decision region. Since the decision region boundary is smooth, such a piecewise linear approximation would be easy to construct.

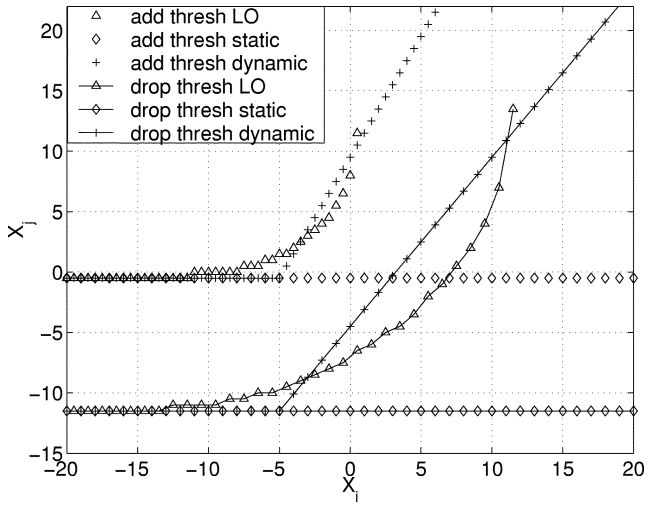

Fig. 4. Decision region boundaries for the reverse link LO algorithm and the static and dynamic threshold handoff algorithms, $c_A = 0.23$, $c_H = 0.22$, $\sigma\sqrt{1-a^2} = 5$ dB.

An important aspect of the LO reverse link decision regions is their scalability. Consider a normalized decision region corresponding to $\Delta_{\mathrm{rev}} = 0$ and $\sigma\sqrt{1-a^2} = 1$. A decision region for an arbitrary set of system parameters can be obtained by shifting the normalized decision region by Δ_{rev} and scaling it by a factor $\sigma\sqrt{1-a^2}$. Thus, it is enough to compute the decision region for one set of system parameter values and use its scaled version when the system parameters change.

The decision regions obey the following commonsense "rule": as the signal quality offered by the current active set improves, the addition of a new base station to the active set should be discouraged and the dropping of a base station from the active set should be encouraged. The LO add and drop thresholds in Fig. 4 obey this rule. In contrast, the ad hoc static threshold algorithm (decision region shown in Fig. 4) does not obey the rule; and, as shown in the next section, it results in inferior performance.

B. Reverse Link Performance

We now examine the improvement in performance that the LO algorithm offers over the static threshold algorithm. Tradeoff curves for both of the algorithms are shown in Fig. 5. The results are obtained by simulation under the environment described in Section III-B. It is immediately clear from Fig. 5 that the LO algorithm offers a significantly better tradeoff than the static threshold algorithm, i.e., for the same λ_{LD} and λ_H, the LO algorithm results in a much lower λ_A. This is because the LO algorithm follows the rule described earlier. Simulations show that the superior performance of the LO algorithm is maintained when a timer is included with the drop threshold of the static threshold algorithm.

The dynamic threshold handoff algorithm of cdma2000 [12] improves on the static threshold algorithm by adopting the heuristic rule discussed in Section IV-A. The decision region

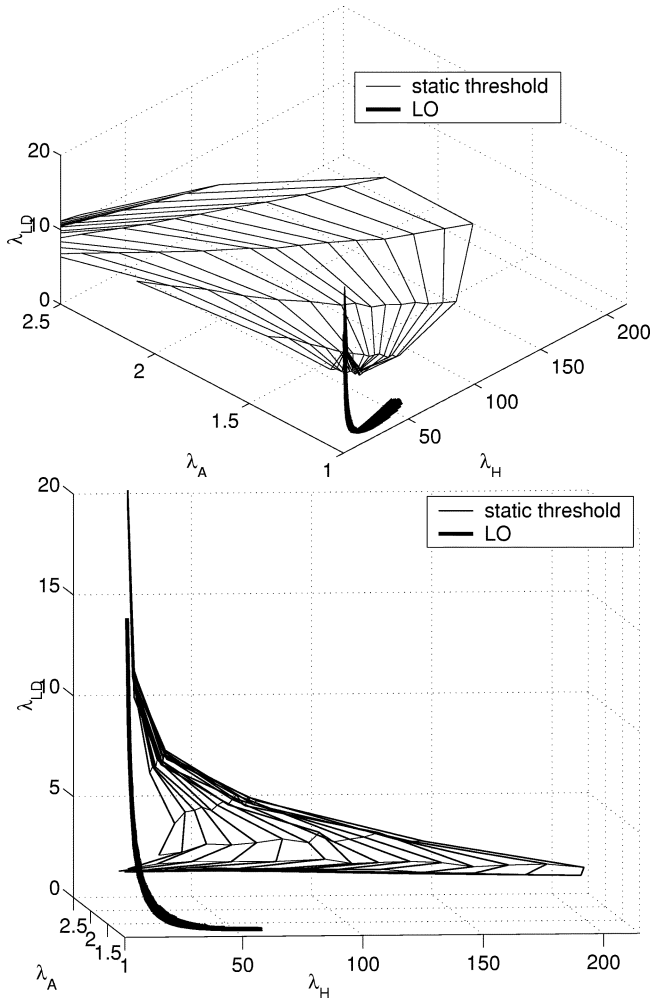

Fig. 5. Comparison of the tradeoff surfaces of the LO and static threshold soft handoff algorithms (reverse link).

of this improved algorithm is shown in Fig. 4. In contrast to the static threshold and LO algorithms, the dynamic threshold algorithm has more than two parameters (the slopes and the intercepts of the decision region boundary for both the add and the drop thresholds constitute four additional parameters). Sample tradeoff surfaces for the dynamic threshold handoff algorithm are obtained by varying only two parameters and fixing the slope and intercept parameters. Such sample surfaces are obtained for various values of the fixed parameters, and the surface that is lowest in the knee region is designated as the tradeoff surface. Our judgment about the lowest surface is somewhat subjective, but in the absence of any analytical tools, it is the only option available.

The tradeoff surfaces for the dynamic threshold and LO algorithms are compared in Fig. 6. The two algorithms result in comparable performance. The match is particularly good in the knee region of the curve, where the operating point is most likely to be selected. The reason for this match may stem from the similarity in the shapes of the decision regions of the two algorithms (Fig. 4). The dynamic threshold algorithm's decision region can be considered as a first-order approximation to the LO algorithm decision region.

The dynamic threshold algorithm was originally developed as an ad hoc improvement over the static threshold handoff algorithm, while the LO algorithm is developed here using precise analytical tools. The match in the performance of the two algorithms and the similarity of their decision regions together provide an analytical justification for the use of the dynamic threshold algorithm.

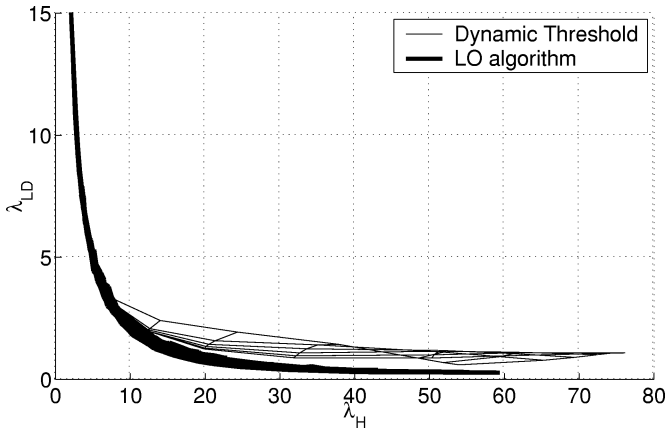

Fig. 6. A comparison of the tradeoff surfaces of the LO and dynamic threshold handoff algorithms (reverse link).

C. Forward Link LO Soft Handoff Algorithm

The forward link LO algorithm is derived as an approximation to the forward link optimal handoff algorithm. The objective of the forward link LO algorithm is to select $A_{\text{for},k+1}$ to minimize the expected value of the following forward link incremental cost

$$J_{\text{for,incr}}(A_{\text{for},\,k+1}) = \mathbb{1}_{\{\text{FLD at time } k+1\}} + c_A\,|A_{\text{for},\,k+1}| + c_H \mathbb{1}_{\{\text{Forward link handoffat time } k\}} \quad (28)$$

giving

$$A^*_{\text{for},k+1} = \arg \min_{A_{\text{for},\,k+1}} \mathrm{E}\left[J_{\text{for,incr}}\,(A_{\text{for},\,k+1}) \mid \mathcal{I}_k, A_{\text{for},k+1}\right]. \quad (29)$$

The expectation of the last two terms in (29) can be evaluated easily, but the first term corresponding to the FLD event defined in (17) presents a challenge because the power sum $X_{\text{for},k}$ in (2) is difficult to study analytically. Santucci *et al.* [18] have studied the sum of correlated lognormal random variables in the context of computing interference statistics in cellular radio. One of the methods they consider is Wilkinson's approximation method, which models the sum of log-normal random variables as a log-normal random variable with appropriately matched mean and variance (see [18] for details).

Fig. 8 compares the exact and Wilkinson approximation decision regions for the forward link LO algorithm. It can be seen that the two methods result in nearly the same decision region. Thus, Wilkinson's method can be used for simulation without significant loss of accuracy. As with the reverse link, the estimators required to implement the forward link handoff algorithm are borrowed from the hard handoff analysis [7], [17].

For implementation purposes, there is a significant difference between the reverse and forward link LO algorithms. Unlike the reverse link LO algorithm, the forward link LO algorithm's decision region does not scale with changes in $\sigma\sqrt{1-a^2}$ and Δ_{for}. This can be seen from the structure of the sum of log-normal sta-

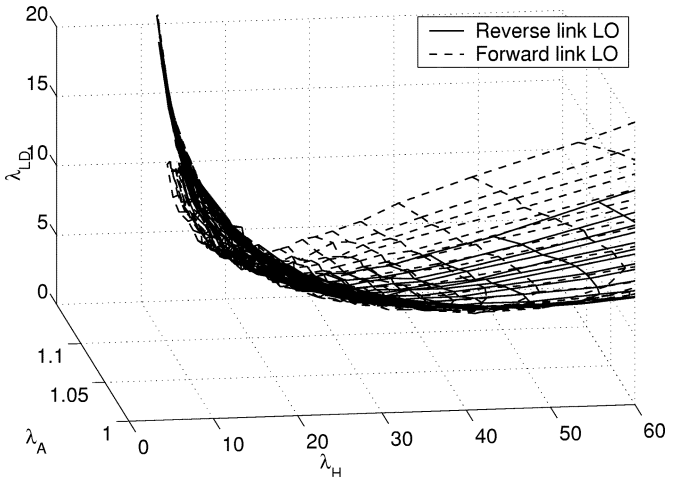

Fig. 7. Forward link performance of the reverse and forward link LO algorithms.

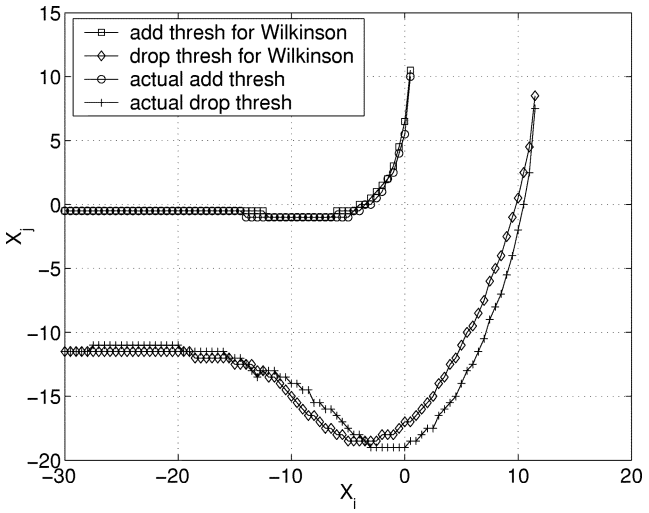

Fig. 8. Accuracy of Wilkinson's approximation ($c_A = 0.23$, $c_H = 0.22$, $\sigma\sqrt{1-a^2} = 5$ dB). The add threshold corresponds to the handoff $\{i\} \rightarrow \{i,j\}$, and the drop threshold corresponds to the transition $\{i,j\} \rightarrow \{i\}$.

tistics. Thus, implementation will involve either storage of the decision region for a large set of possible $\sigma\sqrt{1-a^2}$ and Δ_{for} values or evaluation of the statistics of the sum of log-normals each time the system parameters change.

Implementation of the forward link LO algorithm presents the following two difficulties. The exact decision region is difficult to compute, and the decision region does not scale with system parameters in a simple manner. Since the reverse link LO algorithm is relatively simple to implement, as an approximation, we consider using the reverse link LO algorithm on the forward link. We consider the reverse link LO algorithm with Δ_{rev} set equal to Δ_{for} and use the corresponding decision rule to select the forward link active set.

Under the simulation parameters of Table I the forward link performance of both the forward and reverse link LO algorithms is shown in Fig. 7. It can be seen that both the algorithms achieve nearly the same performance. For brevity, simulation results are given for only one set of parameter values (Table I). The result was also seen to hold for a range of mobile velocities (10, 20, 40 m/s) and Δ values (0 and 10 dBm).

From these results, it is demonstrated that in spite of the different handoff region shapes of the forward and reverse link LO algorithms, the algorithms actually perform quite similarly. Therefore, the simpler reverse link LO algorithm can be used to decide the forward link active set.

The following issue must be stressed. We have not shown that the decision functions Φ_{rev} and Φ_{for} can be made the same. When $\Delta_{\text{rev}} \neq \Delta_{\text{for}}$ or the relative cost parameters c_A and c_H are different on the forward and reverse links, we will actually require that $\Phi_{\text{rev}} \neq \Phi_{\text{for}}$. What we have shown is that Φ_{for} can be implemented using a decision function with the same structural form as the reverse link LO algorithm, but with the forward link parameters in place of the corresponding reverse link parameters.

To understand why the reverse link algorithm works well on the forward link, consider the decision region boundary asymptotes. For the case of handoff from $\{i\} \rightarrow \{i,j\}$, the reverse link LO decision region is shown in Fig. 4. Under the estimator of choice ($E[X_{k+1,i}|X_{k,i}] = X_{k,i}$), the boundary of the reverse link LO region satisfies

$$L\left(\frac{X_i - \Delta_{\text{rev}}}{\sigma'}, \frac{X_j - \Delta_{\text{rev}}}{\sigma'}, \rho\right) + c_A + c_H = Q\left(\frac{X_i - \Delta_{\text{rev}}}{\sigma'}\right) \tag{30}$$

where $\sigma' = \sigma\sqrt{1-a^2}$. To obtain the horizontal asymptote, we set $X_i = -\infty$, and to obtain the vertical asymptote, we set $X_j = \infty$. The resulting reverse link asymptotes are

$$\begin{aligned} X_j - \Delta_{\text{rev}} &= \sigma' Q^{-1}(1 - c_A - c_H) \\ X_i - \Delta_{\text{rev}} &= \sigma' Q^{-1}(c_A + c_H). \end{aligned} \tag{31}$$

For the forward link, the decision region is shown in Fig. 8. Though the decision region boundary depends on the complicated statistics of the sum of log-normals, the asymptotes are relatively simple to evaluate using the following fact. Let $X_{\{i,j\}} = 10\log(10^{X_i/10} + 10^{X_j/10})$ be the power sum of Gaussian random variables (r.v.'s) X_i and X_j. Then $P\{X_{\{i,j\}} < 0\}$ is equal to zero when either X_i or X_j has infinite mean and equal to $Q(E[X_i]/\sigma_i)$ when $E[X_j] = -\infty$. It can be verified from the above fact that when $\Delta_{\text{rev}} = \Delta_{\text{for}}$, the decision region asymptotes on the forward link are the same as that on the reverse link (31). This equivalence can also be verified for the handoff $\{i,j\} \rightarrow \{i\}$.

When three base stations are involved in the handoff, the decision region lies in a three-dimensional space and the asymptotes are not the same for the forward and reverse links. This, along with the difference in handoff region shapes, may account for the slight difference in performance between the LO algorithms designed for the two links.

V. Conclusions

We developed a LO soft handoff algorithm and showed that it outperforms the static threshold handoff algorithm. Further, we showed that the ad hoc dynamic threshold algorithm is a good approximation to the LO algorithm. This provides an analytical justification for the use of the dynamic threshold algorithm.

For the forward link, we developed an LO algorithm and showed that its structure is complicated. However, the simpler reverse link LO algorithm gives nearly the same performance as the forward link LO algorithm. Thus, we have shown that there

is no need to design separate LO handoff algorithms for the forward link.

An issue deserving further exploration is the adaptation of the LO soft handoff algorithm to varying system parameters, along the same lines as the adaptation of the LO hard handoff algorithm [7]. In particular, adaptation to varying traffic load is a topic of interest. The traffic load directly influences the interference levels at the mobile and the base station. The interference levels in turn influence the link degradation thresholds Δ_{for} and Δ_{rev} [see (17) and (20)]. This dependence allows for some degree of adaptation to changing traffic patterns. For example, for a heavily loaded cell, the SIR thresholds will be low, resulting in easier dropping of the cell from the active set. Similarly, a lightly loaded cell will be more likely to be added into the active set. Although handoff based on traffic conditions is not the focus of our work, the argument above shows that the LO algorithm may be able to accommodate such considerations by adjustment of the link degradation thresholds.

REFERENCES

[1] G. Stuber, *Principles of Mobile Communication*. Norwell, MA: Kluwer Academic, 1996.

[2] A. J. Viterbi *et al.*, "Soft handoff extends CDMA cell coverage and increases reverse link capacity," *IEEE J. Select. Areas Commun.*, vol. 12, pp. 1281–8, Oct. 1994.

[3] A. Sendonaris and V. V. Veeravalli, "The coverage capacity tradeoff in cellular CDMA systems with soft-handoff," in *Proc. Asilomar Conf. Sig. Sys. Comput.*, vol. 1, Pacific Grove, CA, 1997.

[4] R. Rezaiifar, A. M. Makowski, and S. P. Kumar, "Stochastic control of handoffs in cellular networks," *IEEE J. Select. Areas Commun.*, vol. 13, pp. 1348–1362, Sept. 1995.

[5] M. Asawa and W. E. Stark, "Optimal scheduling of handoffs in cellular networks," *IEEE/ACM Trans. Networking*, vol. 4, pp. 428–441, June 1996.

[6] V. V. Veeravalli and O. E. Kelly, "A locally optimal handoff algorithm for cellular communications," *IEEE Trans. Veh. Technol.*, vol. 46, pp. 603–610, Aug. 1997.

[7] R. Prakash and V. V. Veeravalli, "Adaptive hard handoff algorithms," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 2456–2455, Nov. 2000.

[8] M. Akar and U. Mitra, "Variations on optimal and suboptimal handoff control for wireless communications systems," *IEEE J. Select. Areas Commun.*, vol. 19, pp. 1173–1185, June 2001.

[9] D. Wong and T. J. Lim, "Soft handoffs in CDMA mobile systems," *IEEE Personal Commun. Mag.*, vol. 4, pp. 6–17, Dec. 1997.

[10] "Mobile station-base station compatibility standard for dual-mode wideband spread spectrum cellular system," Telecommunications Industry Association, TIA/EIA/IS-95, 1993.

[11] N. Zhang and J. M. Holtzman, "Analysis of a CDMA soft-handoff algorithm," *IEEE Trans. Veh. Technol.*, vol. 47, pp. 710–714, May 1998.

[12] "Wideband CDMA one radio transmission technology proposal," International Telecommunication Union, Radiocommunication Study Groups, TIA CDMA2000, 1998.

[13] N. Zhang and J. M. Holtzman, "Analysis of handoff algorithms using both absolute and relative measurements," *IEEE Trans. Veh. Technol.*, vol. 45, pp. 174–9, Feb. 1996.

[14] M. Gudmundson, "Correlation model for shadow fading in mobile radio systems," *Electron. Lett.*, vol. 27, no. 23, pp. 2145–2146, Nov. 1991.

[15] N. L. Johnson and S. Kotz, *Distributions in Statistics: Continuous Multivariate Distributions*. New York: Wiley, 1972.

[16] "Tables of the bivariate normal distribution function and related functions," National Bureau of Standards, ser. Applied Mathematics Series, 1959.

[17] R. Prakash, "Analysis of handoff algorithms," M.S.thesis , School of Electrical Engineering, Cornell Univ., Ithaca, NY, 1999.

[18] F. Santucci, M. Pratesi, M. Ruggeiri, and F. Graziosi, "A general analysis of signal strength handover algorithms with cochannel interference," *IEEE Trans. Commun.*, vol. 48, pp. 231–241, Feb. 2000.

Rajat Prakash received the B.Tech. degree from the Indian Institute of Technology, Kanpur, in 1997 and the M.S. degree from Cornell University, Ithaca, NY, in 1999, both in electrical engineering. He is currently pursuing the Ph.D. degree at the University of Illinois at Urbana-Champaign.

His research interests include wireless communications, information theory, communication networks, and resource allocation.

Mr. Prakash received a Nortel research assistantship at Cornell University and the Mavis Award from the University of Illinois.

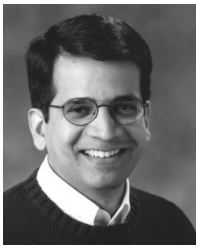

Venugopal V. Veeravalli (S'86–M'92–SM'98) received the B.Tech. degree (silver medal honors) from the Indian Institute of Technology, Bombay, in 1985, the M.S. degree from Carnegie-Mellon University, Pittsburgh, PA, in 1987, and the Ph.D. degree from the University of Illinois at Urbana-Champaign in 1992, all in electrical engineering.

He joined the University of Illinois at Urbana-Champaign in 2000, where he is currently an Associate Professor in the Department of Electrical and Computer Engineering and a Research Associate Professor in the Coordinated Science Laboratory. He was an Assistant Professor at Cornell University, Ithaca, NY, during 1996–2000. His research interests include mobile and wireless communications, detection and estimation theory, and information theory.

Dr. Veeravalli is currently an Associate Editor for IEEE TRANSACTIONS ON INFORMATION THEORY. Among the awards he has received for research and teaching are the IEEE Browder J. Thompson Best Paper Award in 1996, the National Science Foundation CAREER Award in 1998, the Presidential Early Career Award for Scientists and Engineers (PECASE) in 1999, the Michael Tien Excellence in Teaching Award from the College of Engineering, Cornell University, in 1999, and the Beckman Associate Fellowship at the Center for Advanced Study, University of Illinois, in 2002.