Optimal Control of Handoffs in Wireless Networks

Ramin Rezaiifar **Armand M. Makowski** [†]
ramin@src.umd.edu *armand@src.umd.edu*
Department of Electrical Engineering and Institute for Systems Research
University of Maryland, College Park, MD 20742

Srikanta Kumar
kumars@eecs.nwu.edu
Electrical Engineering and Computer Science, Northwestern University, Evanston, IL 60208.

ABSTRACT

A Dynamic Programming formulation is used to obtain an optimal strategy for the handoff problem in cellular radio systems. The formulation includes the modeling of the underlying randomness in received signal strengths and of the mobile's movements. The cost function is designed such that there is a cost associated with switching and a reward for improving the quality of the call. The optimum decision is characterized by a threshold on the difference between the measured power that the mobile receives from the base stations. Also we study the problem of choosing the "best" fixed threshold that minimizes the cost function. The performance of the optimal and suboptimal strategies are compared.

I Introduction

Wireless networks are experiencing rapid growth, a trend likely to continue into the foreseeable future. In both micro and macro cellular networks, a key issue for efficient operation is the problem of handoffs. A call on a portable/mobile which leaves one cell (radio coverage area) and enters a neighboring cell must be transferred to the base station of this neighboring (new) cell. Each handoff involves a signaling cost. Because of statistical fluctuations in signal strength due to fading, a call may get bounced back and forth between neighboring base stations before it is either successfully handed off, or is forced to terminate as the signal strength falls below acceptable levels. An improperly designed handoff algorithm can result in an unacceptably high level of bouncing (resulting in high signaling costs) and/or a high probability of forced termination. We argue that approaching the handoff problem in a stochastic control framework is most appropriate. We use a Markov decision process formulation, and derive optimal handoff strategies via Dynamic Programming (DP).

Typically, in a cellular mobile communication network (analog or digital), each cell is assigned a separate set of channels (frequencies, carriers, or time slots). The assigned set depends on the frequency planning strategy used for spatial reuse, and may be fixed or changed dynamically. A successful handoff entails not only the availability of a channel in the new cell (to which the mobile enters) but also an acceptable level of signal strength on the available channel.

To focus mainly on the handoff issue, we take a simple model of two adjacent cells with one channel per cell, and analyze the optimal handoff when a single mobile with an active call moves from one cell to the other. We assume that these channels are always available, distinct, and that their statistical characteristics are independent. Each channel is assumed to provide a two-way link between the respective base and the mobile (and thus we do not distinguish between frequency or time division duplex link to achieve this two-way communication). We analyze mobile controlled handoff in the sense that the signal strength on each of these channels is measured periodically at regular intervals at the mobile/portable. The signal strength so measured is subject to both path loss and shadow fading. Handoff decisions are made at these measurement instants. Multipath effects are ignored here as the correlation time is typically much smaller than the measurement interval for most cases of practical interest. Possible interference due to other calls being on a co-channel (e.g., same frequency at another base) is also ignored. Nevertheless, the results derived here form a basis for analyzing enriched models that include such interference, availability of multiple channels, and base station controlled or base-mobile negotiated handoffs. Our formulation includes modeling the movement of the mobile as well as the underlying randomness, induced by the (spatially correlated) fading environment, in the signal strengths as observed at the measurement instants.

An optimal handoff strategy should reflect the optimal tradeoff between the call quality (higher signal strength implies a higher call quality) and the signaling costs. If the handoffs could be accomplished without cost (no signaling costs), the best strategy, trivially, is for the mobile to connect to the base (channel) with higher signal strength at each instant. In the presence of non-zero signaling cost, the best handoff strategy should reflect the optimum intertemporal tradeoff (during the lifetime of the call) between the total signaling costs and the quality or signal strength achievable by the connection, instant to instant, relative to the alternative connection present. Accordingly, for purposes of optimization, we consider a cost function that entails a fixed signaling cost for each handoff, and a cost proportional to the power gain foregone when a switch to the higher power is not undertaken. We show that the optimal handoff strategy is characterized by a threshold policy, and that the threshold is defined over the signal strength difference observed on the channels. The specific cost function we use, while reflecting the necessary concerns, also simplifies the numerical computations to obtain the threshold. However, the methodology

[†]The work of these authors was supported partially through NSF Grant NSFD CDR-88-03012 and through NASA Grant NAGW277S.

is applicable to other definitions of cost.

Much of the previous research on handoffs is based on simulation studies, whereas the theoretical studies have focused on the evaluation of the expected number of handoffs for a given hysteresis strategy [3], [7]. This paper is one of the first attempts to address handoffs in a control-theoretic framework. Another contribution in that vein can be found in a recent study by Asawa and Stark [1]; these authors consider an optimization problem similar to the one presented here, but propose only an approximation for solving it.

The paper is organized as follows: In Section II, we present a general Markov decision theoretic framework for addressing the handoff issue. Section III introduces the model being used in this work to characterize the stochastic behavior of the received powers. The stochastic control problem is introduced in Section IV, where the threshold structure of the optimal policy is presented. In Section V we introduce call quality and number of handoffs as two possible measures for assessing the effectiveness of different handoff schemes. Section VI contains several numerical results and comparison between different handoff schemes. For lack of space, proofs and technical details have been omitted; they can be found in [5], [6].

A few words on the notation used throughout: For any x in $\mathbb{R}^2$, we write $\|x\|$ for its Euclidean norm, and $[X \mid Y]$ refers to any random variable (rv) which is distributed according to the conditional distribution of X given Y. We also write $X \sim \mathcal{N}(\mu, R)$ to signify that the rv X is distributed according to a Gaussian distribution with mean vector μ and covariance matrix R. For any sequence of rv's $\{\xi_t,\ t=0,1,\ldots\}$, we set $\xi^t \equiv (\xi_0, \xi_1, \ldots, \xi_t)$ for the history of the sequence up to time $t = 0, 1, \ldots$.

II The Model

We now introduce a Markov decision process formulation for the handoff problem faced by a mobile which receives signals from two distinct base stations, labeled base stations zero and one, while moving within a given geographical area.

A. *The Underlying Randomness*

We begin by describing the elements of the model which are unaffected by the mobile's control actions. This includes randomness in signal propagation and fading as well as possible randomness in the mobile's movements. The mobile moves through a region E of the plane $\mathbb{R}^2$, which we assume composed of a finite number of points in the plane. This is done in order to simplify the discussion, with the understanding that most of the developments herein applies to the case of more general regions. The mobile then travels through E according to a stochastic process $\{S_t,\ t=0,1,\ldots\}$ with S_t denoting the position in E of the mobile at the beginning of the time slot $[t, t+1)$. At time t, the strength of the received signal from base station i is denoted by P_t^i, $i = 0, 1$; it is measured in dB relative to a fixed transmitter power. For notational convenience, we write $P_t \equiv (P_t^0, P_t^1)$ and $X_t \equiv (S_t, P_t)$. The joint evolution of position and power levels $\{X_t,\ t=0,1,\ldots\}$ is modeled as a time-homogeneous Markov process with the following structure: First, we assume that the position process $\{S_t,\ t=0,1,\ldots\}$ is by itself a time-homogeneous Markov process on E with one-step transition probability matrix $Q \equiv (Q(s; s'))$ such that

$$\mathbf{P}[S_{t+1} = s_{t+1} \mid X^t = x^t] = Q(s_t; s_{t+1}). \qquad \text{(II.1)}$$

Next, we postulate

$$\begin{aligned}\mathbf{P}[P_{t+1} \le p \mid X^t = x^t, S_{t+1} = s_{t+1}] \\ = \quad G(p \mid s_t, p_t, s_{t+1}), \quad p \in \mathbb{R}^2 \end{aligned} \qquad \text{(II.2)}$$

where $G(\cdot \mid s_t, p_t, s_{t+1})$ denotes the conditional probability distribution of P_{t+1} given that the mobile is in positions s_t and s_{t+1} at time t and $t+1$, respectively, and power strengths at time t were observed at levels p_t. The assumption (II.2) attempts to model the dependence between measured power levels as rather short-term and short-range. Although not entirely accurate, (II.2) is nevertheless compatible with modeling assumptions used in previous works [3], [4], [7]; we shall return to this point in Section III.

Finally, upon combining (II.1) and (II.2), we see by a simple conditioning argument that

$$\begin{aligned}\mathbf{P}[S_{t+1} = s_{t+1}, P_{t+1} \le p \mid X^t = x^t] \\ = \quad G(p \mid s_t, p_t, s_{t+1}) Q(s_t; s_{t+1}) \end{aligned} \qquad \text{(II.3)}$$

and the process $\{X_t,\ t = 0, 1, \ldots\}$ is indeed a time-homogeneous Markov process on $E \times \mathbb{R}^2$.

The call initiated at time $t = 0$ will last a random number T of time slots. We adopt the traditional assumption that the duration of a call is adequately modeled as an exponential rv. In line with this standard assumption, in our discrete-time setup we assume that the rv T is geometrically distributed, say $\mathbf{P}[T = t+1] = \rho(1-\rho)^t$ for all $t = 0, 1, \ldots$ for some $0 < \rho < 1$. Alternatively, we may interpret ρ as the hangup probability, so that the call can be terminated in every time slot with probability ρ, and this independently of the duration of the ongoing call. The call duration T is assumed independent of the sequence $\{X_t,\ t=0,1,\ldots\}$.

B. *The Controlled System*

Fix $t = 0, 1, \ldots$. At the beginning of the time slot $[t, t+1)$, the mobile is in location S_t, the power strengths from the base stations have been measured at levels P_t^0 and P_t^1, and a decision needs to be taken so as to which base station to use for transmission during the time slot $[t, t+1)$. This action is selected on the basis of available information in a way that we now proceed to define: Let U_t denote the $\{0,1\}$–valued rv which encodes the decision taken at time t, i.e., if $U_t = i$, $i = 0, 1$, then base station i is being used during the time slot $[t, t+1)$. For reasons that will become apparent soon, we set $I_t \equiv U_{t-1}$, so that I_t denotes the base station to which the mobile is attached during the time interval $[t-1, t)$; we also define I_0 as being arbitrary.

The information available to the decision–maker is described by the rv's $\{H_t,\ t = 0, 1, \ldots\}$ which are defined recursively by $H_{t+1} \equiv (H_t, U_t, X_{t+1}, I_{t+1})$ with $H_0 \equiv (X_0, I_0)$. To determine the successive decisions on the basis of this information pattern, we introduce the following notion of a (control) policy: A policy π is a collection of mappings $\{\pi_t,\ t = 0, 1, \ldots\}$ where for each $t = 0, 1, \ldots$, π_t maps the range of H_t into $\{0, 1\}$, with the interpretation that the base station $\pi_t(h_t)$ is used during the time slot $[t, t+1)$ if $H_t = h_t$. The policy π is said to be a Markov stationary policy if there exists a single mapping $f : E \times \mathbb{R}^2 \times \{0,1\} \rightarrow \{0,1\}$ such that $\pi_t(h_t) = f(x_t, i_t)$ with x_t determined through $h_t = (h_{t-1}, u_{t-1}, x_t, i_t)$. The class of all control policies is denoted by $\mathcal{P}$.

Fix a pair (x, i) in $E \times \mathbb{R}^2 \times \{0,1\}$, and $t = 0, 1, \ldots$. For each policy π in $\mathcal{P}$, we associate a probability measure $\mathbf{P}^{\pi}_{x,i}$ such that $\mathbf{P}^{\pi}_{x,i}[X_0 = x, I_0 = i] = 1$, and

$$\begin{aligned} &\mathbf{P}^{\pi}_{x,i}[S_{t+1} = s_{t+1}, P_{t+1} \leq p, I_{t+1} = i_{t+1} \mid H_t, U_t] \\ &= \delta(i_{t+1}, U_t) G(p \mid S_t, P_t, s_{t+1}) Q(S_t; s_{t+1}) \end{aligned}$$

where we have made use of the equality $I_{t+1} = U_t$, and of the requirement that the underlying randomness be governed by (II.3), and this independently of the policy in use.

The model is fully specified if we further assume the rv T to be independent of the rv's $\{X_t, U_t,\ t = 0, 1, \ldots\}$ under $\mathbf{P}^{\pi}_{x,i}$, and this for each policy π in $\mathcal{P}$. Such specifications amount to casting this controlled system as a Markov decision process with "state" process $\{(X_t, I_t),\ t = 0, 1, \ldots\}$. We refer the reader to the monographs [2] for additional material on Markov decision processes.

III Gaussian Power Distribution Models

The conditional distribution $G(\cdot \mid s_t, p_t, s_{t+1})$ appearing in (II.2) is the component of the model that is hardest to specify. We now present a model which we use in the remainder of this paper, both for the purpose of analysis as well as for carrying out numerical experiments. This model can be viewed as a dynamic version of a static model which has been widely used to capture shadowing effects [3], [4]: Let $\{W^i(r), r \in \mathbb{R}^2\}$ denote a family of *jointly Gaussian* rv's with zero mean and variance σ_i^2, $i = 0, 1$, and with correlation structure

$$\mathbf{E}[W^i(r) W^i(r')] = \sigma_i^2 \exp(-\beta^{-1} \|r - r'\|), \quad r, r' \in \mathbb{R}^2 \tag{III.1}$$

for constants $\beta > 0$ and $\sigma_i^2 > 0$. The two families $\{W^0(r), r \in \mathbb{R}^2\}$ and $\{W^1(r), r \in \mathbb{R}^2\}$ are assumed *independent.*

Let b_i denote the location of base station i, $i = 0, 1$. In location s, the strength $P^i(s)$ of the signal produced by the base station i is then given by

$$P^i(s) \equiv A_i - B_i \log(\|s - b_i\|) + W^i(s - b_i), \quad s \in \mathbb{R}^2. \tag{III.2}$$

The constant A_i reflects the transmitter power and is a function of transmission frequency and height of the antennas, while B_i, with typical values in the range of 30-40 dB, models the path loss [4]. We find it convenient to write

$$\mu_i(s) \equiv A_i - B_i \log(\|s - b_i\|), \quad s \in \mathbb{R}^2,\ i = 0, 1. \tag{III.3}$$

The model (III.1)–(III.2) is a spatial one which specifies the distribution of power levels solely as a function of position, and does not fit into the framework of Section II. As we seek to develop a dynamic model which does fit and which is also compatible with that spatial model, we first consider the following line of reasoning: Assume that the shadowing effects are essentially static, i.e., do not vary much over the duration of a call, and are described by the static random fields $\{P^i(r), r \in \mathbb{R}^2\}$, $i = 0, 1$ – these can be thought as being generated at the beginning of time $t = 0$. It then seems reasonable to argue that the power levels at time t are those given by these static random fields evaluated at the position occupied by the mobile at time t. In other words, the power levels $\{P^i_t,\ t = 0, 1, \ldots\}$ can be obtained by "composing" the static random fields $\{P^i(r), r \in \mathbb{R}^2\}$, $i = 0, 1$ with the mobile's motion, namely

$$P^i_t \equiv P^i(S_t) = \mu_i(S_t) + W^i_t, \quad t = 0, 1, \ldots \tag{III.4}$$

where we have set $W^i_t \equiv W^i(S_t - b_i)$. The random field $\{(P^0(r), P^1(r)), r \in \mathbb{R}^2\}$, or equivalently $\{(W^0(r), W^1(r)), r \in \mathbb{R}^2\}$, is assumed *independent* of the mobile's trajectory $\{S_t,\ t = 0, 1, \ldots\}$.

Simple calculations [5] show that the Markov property (II.2) does not hold under the foregoing assumptions. Undeterred by this unfortunate state of affairs we take the position that temporal variations have short–term memory. This assumption, when coupled with additional calculations on the model (III.4), leads us to the following dynamic models [5]: We *posit* the power levels to have the form (III.4) where for each $t = 0, 1, \ldots$, the rv's W^0_{t+1} and W^1_{t+1} are *conditionally independent* given $(W^{0,t}, W^{1,t}, S^{t+1})$, and for $i = 0, 1$, the rv W^i_{t+1} is *conditionally Gaussian* given $(W^{0,t}, W^{1,t}, S^{t+1})$ with the requirement that the conditional mean γ^i_{t+1} and variance Γ^i_{t+1} depend *only* on the variables W^i_t, S_t and S_{t+1}. For the sake of concreteness we carry out the discussion in the special case

$$\gamma^i_{t+1} = W^i_t \alpha_t, \quad \text{and} \quad \Gamma^i_{t+1} = \sigma_i^2 \left(1 - \alpha_t^2\right), \tag{III.5}$$

where

$$\alpha_t = \exp(-\beta^{-1} \|S_t - S_{t+1}\|). \tag{III.6}$$

Under these assumptions, the rv's P^0_{t+1}, and P^1_{t+1} are then conditionally independent given (X^t, S_{t+1}), and for $i = 0, 1$, the rv P^i_{t+1} is conditionally Gaussian given (X^t, S_{t+1}), i.e.,

$$[P^i_{t+1} \mid X^t, S_{t+1}] \sim \mathcal{N}(\mu_i(S_{t+1}) + \gamma^i_{t+1}, \Gamma^i_{t+1}) \tag{III.7}$$

and the conditional distribution $G(\cdot \mid s_t, p_t, s_{t+1})$ is therefore Gaussian.

IV A Stochastic Optimization Problem

In order to formulate the handoff problem as a stochastic optimization problem, we need to define a cost structure which quantifies the cost associated with operating the system under any policy in $\mathcal{P}$: First we select a cost-per-stage $c : \mathbb{R}^2 \times \{0,1\} \times \{0,1\} \rightarrow \mathbb{R}$, and for every initial condition (x, i), we define the total cost function

$$J_\pi(x, i) \equiv \mathbf{E}^{\pi}_{x,i} \left[\sum_{t=0}^{T-1} c(P_t, I_t, U_t) \right], \quad \pi \in \mathcal{P}. \tag{IV.1}$$

The problem of interest is then that of finding a policy $\pi^\star$ in $\mathcal{P}$ such that

$$J_{\pi^\star}(x, i) \leq J_\pi(x, i), \quad (x, i) \in E \times \mathbb{R}^2 \times \{0,1\} \tag{IV.2}$$

for every other policy π in $\mathcal{P}$. Such a policy $\pi^\star$, when it exists, is called the optimal (handoff) policy.

To settle on a reasonable cost-per-stage c, we argue as follows: Each time the mobile unit chooses a new base station, a database in the switching center is updated to keep track of the mobile's location. Because frequent and unnecessary switches between base stations can be wasteful of system resources, the cost function must be chosen so as to create a trade off between the two possible decisions, namely switching and not switching. One particular cost-per-stage function with this property associates a cost C with switching from one base station to the other, and penalizes the action of not switching by a cost proportional to the difference in signal strength between the alternative base station and the current one. For example, if the mobile unit is connected to base 0 and the strength of the signal from the other base, namely base 1, is higher by $p^1 - p^0$, then we assign the cost $p^1 - p^0$ for not switching to base station 1. The opportunity cost $p^1 - p^0$ encourages the mobile unit to switch to the better base station, whereas the fixed switching cost C creates a trade off. The corresponding cost-per-stage function c is given by

$$c(x,i,u) = \begin{cases} C & \text{if } i \neq u \\ (-1)^i (p^1 - p^0) & \text{if } i = u, \\ & x = (s,(p^0,p^1)) \end{cases} \quad \text{(IV.3)}$$

and is used in (IV.2) throughout the discussion.

In [5], [6] we have shown that the cost function (IV.1) is well defined and finite for every policy π, and that it can be written in alternate form

$$J_\pi(x,i) = \mathbf{E}^\pi_{x,i}\left[\sum_{t=0}^{\infty}(1-\rho)^t c(X_t, I_t, U_t)\right]. \quad \text{(IV.4)}$$

Hence, the total cost problem (IV.1)-(IV.3) can be recast as an infinite horizon *discounted* cost problem with discount factor $1-\rho$. The standard machinery of DP thus applies and leads to a simple characterization of the optimal policy. In the interest of brevity, we only present the main results, with details available in [5], [6].

First, we define the value function for the problem (IV.1)-(IV.3) by

$$V(x,i) \equiv \inf_{\pi \in \mathcal{P}} J_\pi(x,i). \quad \text{(IV.5)}$$

Using the usual backward induction arguments, we can show under (III.5)-(III.6) that the value function $p \to V(s,p,i)$ is a function of the difference $z \equiv p^1 - p^0$. It then follows from the DP optimality equation that the optimal policy $\pi^\star$ is a Markov stationary policy which depends on the power level vector p *only* through the difference z in their components. In fact, it turns out that the optimal policy $\pi^\star$ can be further characterized as belonging to the following class of threshold policies: A handoff policy π is said to be a threshold policy with threshold functions $\tau_i : E \to \mathbb{R}$, $i = 0,1$, if it is a Markov stationary policy such that for every (s,z) in $E \times \mathbb{R}$, $\pi(s,z,0) = 1$ iff $z \geq \tau_0(s)$ *and* $\pi(s,z,1) = 0$ iff $z \leq \tau_1(s)$.

Proposition IV.1 *The optimal handoff policy $\pi^\star$ is a threshold policy with threshold functions $\tau_i^\star : E \to \mathbb{R}$, $i = 0,1$.*

V Average Quality of Call and Expected Number of Handoffs

Once a handoff policy (be it optimal or not) has been selected, it is of interest to compute the expected value of the quality of the call and the expected number of handoffs that the mobile experiences while the optimal policy is in effect. These two quantities constitute good measures of the effectiveness of a handoff policy. Other criteria include the expected delay in handoff which has been studied by Vijayan and Holtzman [7].

Consider a policy π in $\mathcal{P}$. Its call quality function C_π is defined as the expected cumulative strength of the signal received from the active base station under the policy π during the call session, namely

$$C_\pi(x,i) \equiv \mathbf{E}^\pi_{x,i}\left[\sum_{t=0}^{T-1} I_t P_t^1 + (1-I_t)P_t^0\right],$$

while the expected number of handoffs under the policy π is given by

$$S_\pi(x,i) \equiv \mathbf{E}^\pi_{x,i}\left[\sum_{t=0}^{T-1} \mathbf{1}[U_t \neq I_t]\right]$$

for $(x,i) \in E \times \mathbb{R}^2 \times \{0,1\}$. By an argument similar to that leading to (IV.4), both C_π and S_π can be written as discounted cost functions. As argued in [5], [6], for any Markov stationary policy π, hence for any threshold policy, we can interpret J_π, C_π and S_π as fixed points of suitably defined contractions on an appropriate Banach space of functions. This fact can be exploited in the usual manner for numerical purposes.

VI Numerical Results

In this section, we apply the ideas presented earlier to the scenario where a mobile travels in a two-dimensional region; the road divides into two different paths with 70% of the mobiles taking one path and the remaining 30% taking the other.

The discussion is carried out for the special case (III.5)-(III.6) with the numerical values $A_i = 0$, $B_i = 30$, $\sigma_i = 5$dB, $i = 0,1$, $\rho = 0.2$, $C = 6$, $\beta = 200$m; the two base stations are 2Km apart. The mobile path and the optimal thresholds are shown in Fig. 1. The thresholds are lower for the points that are closer to base 1.

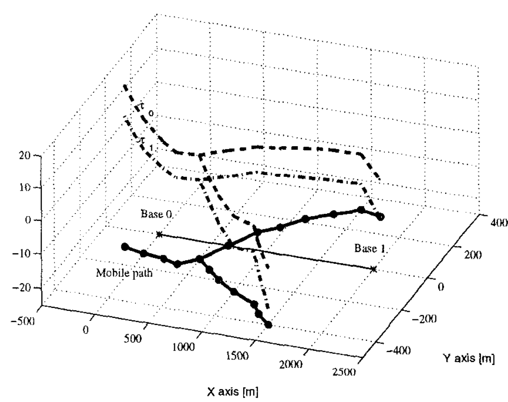

Figure 1. Mobile path together with the optimum thresholds.

Clearly, the solution of the optimization problem depends on the structure of the cost function, as well as on the choice of the various parameters that enter the cost function. One of the important parameters is the switching cost C. In what follows, we present two methods to select a reasonable value

Table 1. Values of J_π, C_π, and S_π for three handoff policies

π	J_π	C_π	S_π
Optimal	-19.01	-442.80	0.28
Sub-optimal	-18.58	-443.60	0.34
σ-threshold	-17.49	-444.15	0.63

for this parameter. Note that in the cost function presented in (IV.3), the switching cost is being compared with the improvement in the signal strength in dB. We must, therefore, determine how expensive the switching action is relative to the potential improvement achieved by switching to the better base station. Alternatively, the call quality can be computed for different values of C and based on the desired value of the average call quality, the appropriate switching cost can be obtained. In Fig. 2 we have displayed the call quality versus switching cost. Because we have normalized by setting $A_i = 0$, the constant A_i must be added to the numerical values for the average call quality in order to obtain the actual signal strength. As expected, call quality degrades with an increase in the switching cost because this increase makes the switching action more sluggish.

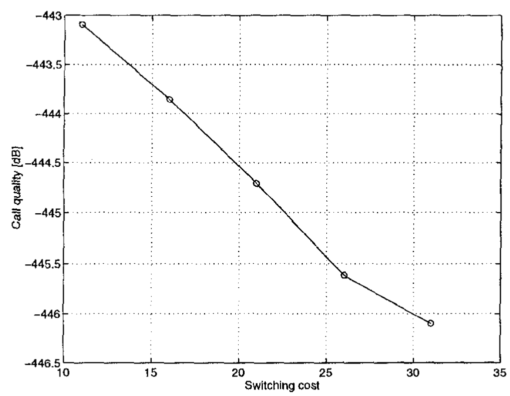

Figure 2. Call quality degrades as the switching cost increases (A_i has been set to zero).

Finally, we compare different aspects of three handoff strategies, namely, the optimal policy, the best fixed (sub-optimal) threshold policy, and a non-optimal threshold policy with thresholds equal to the value of σ. The results in Table 1 show that the optimal strategy achieves a better call quality while making fewer switches, than the other two strategies. Even the suboptimal strategy shows an improvement over the non-optimal method in both call quality and expected number of switches. It is also worth emphasizing that the optimization scheme creates a balance between call quality and the number of switches; otherwise we could improve call quality by choosing a very small threshold which has the effect of increasing the number of switches.

VII Conclusions

The problem of handoff in a cellular environment has been cast as a Markov decision problem. We exploited the well-developed machinery of DP to derive the structure of the optimal handoff policy. This contrasts with most earlier studies which focus only on analyzing the expected number of switches for a given threshold (hysteresis) policy. The optimal policy is obtained by minimizing a cost function that creates a balance between two conflicting measures, i.e., number of switches between cell sites and quality of the call.

The optimal strategy is shown to be of the threshold type, a fact which greatly facilitates its implementation. Through numerical computation we demonstrated that the optimal policy outperforms the conventional non-optimal handoff policy in both the number of switches between the cell sites and the quality of the call. This performance improvement is likely to be greater when fading variability and correlation are high. The proposed design methodology for handoff policies is also applicable for indoor wireless communication as well as for personal communication systems (PCS); in these situations the size of the cells are much smaller (microcells and picocells) and the use of an effective handoff policy is even more crucial. The results are also useful in cell engineering.

Several extensions of the model studied here will prove useful. The optimal handoff strategy depends on the mobility model. In practice, different mobiles/portables may have different patterns of movement, thus requiring different mobility models, whereas a common handoff strategy may be desired for all portables in the system. Additionally, it would be useful to extend the results of this paper to incorporate multiple channels per base stations, more than two bases, a detailed model of co-channel interference, and possible non-availability of channels. Multiple traffic classes with different objective functions and grades of service is another topic. For example, in Cellular Data Packet Delivery (CDPD) systems, data calls use channels when they are not in use by voice calls. Handoff schemes for data calls must reflect channel availability as well as required service quality (low bit error rate).

References

[1] M. Asawa and W.E. Stark, "A framework for optimal scheduling of handoff in wireless networks," in *Proceedings of IEEE Globecom '94*, San Francisco, CA, pp 1669–1673, November 1994.

[2] D.P. Bertsekas, *Dynamic Programming: Deterministic and Stochastic Models*, Prentice-Hall, Englewood Cliffs (NJ), 1987.

[3] M. Gudmundson, "Correlation model for shadow fading in mobile radio systems," in *Electron. Lett.*, vol. **27** (1991), pp. 2145–2146.

[4] W.C.Y. Lee, *Mobile Communications Design Fundamentals*, Second Edition, John Wiley & Sons, New York (NY), 1993.

[5] R. Rezaiifar, *Stochastic Optimization of Handoff in Cellular Networks*, Ph.D. Thesis, Electrical Engineering Department, University of Maryland, College Park (MD), 1995.

[6] R. Rezaiifar, A. M. Makowski, and S. Kumar, "Stochastic control of handoffs in cellular networks," *Journal on Selected Areas in Communications on "Advances in Networking,"* forthcoming (1995).

[7] R. Vijayan and J.M. Holtzman, "A model for analyzing handoff algorithms," *IEEE Transactions on Vehicular Technology* **VT–42** (1993), pp. 351–356.