Two-State Pattern-Recognition Handoffs for Corner-Turning Situations

K. Daniel Wong, *Member, IEEE,* and Donald C. Cox, *Fellow, IEEE*

***Abstract*—Handoff algorithms are used in wireless cellular systems to decide when and to which base station to handoff. Traditional handoff algorithms generally cannot keep both the average number of unnecessary handoffs and the handoff decision delay low. They do not exploit the relative constancy of path loss and shadow fading effects at any given location around a base station. However, handoff algorithms with both a negligible number of unnecessary handoffs and a negligible decision delay can be realized by exploiting this information. One example is the set of handoff algorithms using pattern recognition introduced in previous work. In this paper, we describe how pattern-recognition handoff algorithms can be applied to the problem of turning a corner. This can be used as part of an integrated pattern-recognition handoff algorithm or together with a traditional handoff algorithm, in which case the pattern recognition handles only the special cases like turning a corner.**

***Index Terms*—Handoffs, handovers, land mobile radio cellular systems, land mobile radio propagation factors, pattern recognition.**

I. Introduction

HANDOFF algorithms in existing cellular systems are either of the hard handoff or the soft handoff variety. This paper discusses an improvement in hard handoff algorithms using pattern recognition.[1] These handoff algorithms will be referred to as traditional handoff algorithms (THOs). THOs generally use relative signal strength as a main component of the handoff decision process. The relative signal strength of a base station is the difference in signal power level associated with that base station from that associated with the serving base station. Whenever the relative signal strength of a given base station rises above a threshold Δ dB, handoff is performed to that base station. Because of fading effects, the relative signal strength of base stations other than the serving base station can be positive for brief periods of time when it may not really be necessary to hand off to those base stations. Handoff in such cases wastes network resources and is undesirable. Therefore, the threshold Δ is typically set at a few decibels, introducing hysteresis into the handoff process, to reduce the frequency of unnecessary handoffs. While hysteresis reduces the frequency of unnecessary handoffs, it also increases the decision delay. Large decision delay is undesirable, like frequent unnecessary handoffs. The smaller Δ is, the more frequent the unnecessary handoffs; but the larger Δ is, the larger the decision delay. There is therefore an inherent tradeoff between number of unnecessary handoffs and decision delay with THOs [1]. In other words, THOs cannot be *accurate* and *timely* at the same time.

We have introduced new handoff algorithms [2] to overcome the shortcomings of THOs discussed above. These new algorithms use statistical pattern-recognition techniques to exploit previous knowledge about the propagation characteristics of the environment, unlike THOs. Our pattern-recognition handoff algorithms (PRHOs) train on the signal strength measurements at the locations in which handoffs may be desired, and so acquire specific knowledge about the propagation characteristics at these locations. In particular, the signal strength measurements are composed of relatively deterministic components (path loss and shadow fading) and random components (Rayleigh fading) whose averages are the deterministic components. The reasonable assumption is made that each time a user travels along the same stretch of road, the measured signal power from a particular base station [without loss of generality, we assume a mobile-assisted handoff control scheme] has the same unique deterministic component but a different random component. Pattern classes can therefore be set up to yield a one-to-one correspondence with a user being at or near these locations. Matching these patterns could accurately signal the need for a handoff, without incurring much delay. Hence, pattern-recognition techniques can be applied to provide more accurate and timely handoff decisions than is possible with THO. Handoff algorithms with few unnecessary handoffs and low decision delay can thus be realized.

Handoffs that are consistently both accurate and timely can result in higher capacity and better overall link quality than what is available with today's systems. With continual demands for more system capacity, there is a trend toward smaller cells, also known as microcells. Handoffs are more critical in systems with smaller cells, because for a given average user speed, handoff rates tend to be inversely proportional to cell size. A serious problem in microcellular systems is that the "corner turning" problem can be quite severe, as will be explained in Section II. Hence, we examine the use of PRHO to assist in such situations.

Previous work on the use of pattern-recognition techniques for handoff-related purposes includes [3] and [4]. The pattern classes in [4] are associated with *serving base stations*. As long as a user's measured signal strength patterns are classified as belonging to the pattern class associated with its serving base

Manuscript received October 11, 1999; revised June 27, 2000. This work was supported in part by Nortel and in part by the Stanford University Center for Telecommunications. This work was presented in part at Globecom'98.

K. D. Wong is with Telcordia Technologies, Red Bank, NJ 07701 USA (e-mail: dwong@research.telcordia.com; kuokwong@alumni.princeton.edu).

D. C. Cox is with the Center for Telecommunications, Stanford University, Stanford, CA 94305 USA (e-mail: dcox@nova.stanford.edu).

Publisher Item Identifier S 0018-9545(01)01220-8.

[1]It should not be difficult to extend these the algorithms to work with soft handoffs.

station, it does not hand off to another base station. As soon as the patterns are classified as belonging to the pattern class associated with another base station, the user hands off to the other base station. This differs from our work in that our pattern classes are associated with *stretches of street*. Reference [3] uses hidden Markov models for pattern recognition of signal strength patterns for user location. It also mentions that the proposed pattern-recognition techniques could possibly be used for various system functions like handoff. More recent work on pattern recognition for handoff algorithms is described in [5]. This work uses probabilistic neural networks and attempts much more extensive tracking of user location than do the algorithms in this paper.

The rest of this paper is structured as follows. Section II describes the problem sometimes faced in turning a corner, when radio link conditions may change rapidly and drastically. Section III reviews PRHO concepts. Our solution to the corner turning problem using PRHO methods is introduced in Section IV. Performance and simulation results of the new algorithm are discussed in Section V. Concluding remarks are made in Section VI.

II. The Corner-Turning Problem

Traditional cellular systems have cell radii on the order of a few kilometers. However, the radio spectrum is limited, and so there is only a limited number of users that can be served per cell. One way to increase the system capacity is to decrease the size of the cells. This permits more users to be supported per unit area. This is known as a microcellular system, and it makes sense where there is a high concentration of users, such as in a large city. As [6] puts it,

> The greatest single factor in enhancing spectral efficiency of a network is not complex multiple access techniques, efficient speech and channel coding, modulation, powerful protocols, etc., but the mass deployment of microcells. By this simple technique we can repeatedly and efficiently reuse the precious spectrum.

There are two main reasons why we have special interest in the microcellular propagation environment. First, because of the increasing popularity of cellular services, the trend is to move toward microcellular systems to provide increasing capacity. An example of such a system already installed, and with a rapidly growing subscriber base, is the Japanese Personal Handyphone System [7]. Secondly, smaller cells generally mean more handoffs are needed, all else being equal. This has important implications for handoff algorithms, which will need to perform better in order to handle the greater responsibility placed on them. For example, to keep the same probability of a call's being dropped because of handoff failure sometime during the duration of a call, the per-handoff probability of droppage needs to decrease.

A typical microcellular environment consists of tall buildings touching each other along both sides of streets, as found in the downtown area of many large cities [8]. The location of base stations along streets with antenna heights well below building tops results in the channeling or guiding of radio energy along the street [9], which we call the *waveguide effect* because of the resemblance to wave propagation in waveguides [10]. For example, the base-station antennas (and maybe even the base stations themselves) might be mounted on flag poles or lamp posts along streets. This kind of microcellular environment presents two challenging problems to handoff algorithms. First, significant coupling loss (on the order of 20–25 dB) occurs around corners at intersections [11], [12]. Large signal loss also accompanies the movement of pedestrian users into buildings along the streets. One implication of this phenomenon is that a communication link can go from adequate (or good) to bad as a user turns a corner. Handoff algorithms must be alert for such changes, for better or for worse. In order to make accurate handoff decisions, the handoff algorithms must monitor not only the serving base station but also potential candidate serving base stations. Secondly, the drop in signal strength may not only be severe—it may also be quick. It may occur within a few meters of the corner. Hence, quick handoff decisions are crucial in this situation. Otherwise, the communication links will suffer severe degradation and calls may even be dropped.

The corner-turning scenario poses a great difficulty to traditional handoff algorithms, which have basically two degrees of freedom—they can adjust the averaging window or the hysteresis margin. Reference [13] notes that "effective handoff algorithms for this scenario should use short temporal averaging window and a large hysteresis, so that rapid changes in the mean signal strength are detected and unnecessary handoffs are prevented." The authors also note: "Unfortunately, temporal averaging with a short fixed window length gives optimal handoff performance for only a single velocity." Reference [14] looks at the problem of averaging interval in more depth. However, we are not aware of any satisfactory solution to the challenge posed by the corner-turning scenario, using traditional or other handoff algorithms. The closest attempt might be that of Verdone and Zanella [15], which is based on using the instantaneous gradient of the serving base station signal strength measurements as a handoff trigger. The reason this is used is that the gradient is expected to decrease sharply when there is a sudden, large drop in signal strength, e.g., when turning a corner. We have simulated this algorithm and found that it does not satisfy our requirements. We found that either the number of handoffs increases (compared with THO) or the delay increases, but not both at the same time, no matter what threshold for the gradient is used. Hence, like THO, this algorithm does not escape the tradeoff between number of unnecessary handoffs and handoff delay. It will be shown that the proposed PRHO can escape this tradeoff. Verdone and Zanella's algorithm does not exploit memory like PRHO does.

III. Fundamentals of PRHO

Signal strength patterns are collections of signal strength measurements. The patterns consists of measurements of consecutive samples of signal strength measured from one or more base stations as a user is moving along streets in a cellular system. The structure in such signal strength patterns, and in particular their repeatability, motivates the use of signal strength patterns to provide user location information, which may be useful as input to the handoff decision process.

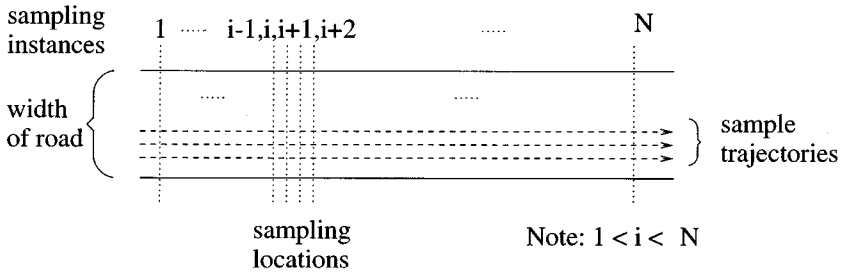

Fig. 1. Sample trajectories of a measurement vehicle used in training to obtain samples of signal strength to estimate optimal handoff location-destination pairs. The sampling locations are also illustrated.

The fundamental basis of our approach is discussed in Section III-A. The assumptions we make and the terminology we use in discussing signal strength patterns are introduced in Sections III-B and III-C. The sliding window pattern classification of signal strength patterns is developed in Section III-D, where our template-based statistical pattern-recognition technique is introduced.

A. The Basis of Our Approach

Large-scale fading effects in cellular or microcellular systems typically consist mainly of reflection, refraction, and scattering from objects like buildings and trees, which are relatively stationary over periods of days or weeks. Therefore, for any given location in a cellular system, the local mean signal power level is relatively fixed. Furthermore, users in cellular or microcellular systems in cities tend to move in very regular, determined paths, namely, along lanes of streets or along sidewalks. These two factors result in considerable *regularity* of structure and *repeatibility* in the signal strength measurements (of beacon signals of base stations), which are made by users in cities.

Consider a situation in which a user moves along a particular stretch of street, and then another user a little while later moves along the same stretch of street. Both users make signal strength measurements from several base stations around them as they move along the street. Let the measurements made by the first user be consecutive samples of signal strength denoted by $\{\gamma_{0,i}\}$ for $i = 0, 1, \ldots, N-1$ and the measurements made by the second user be denoted by $\{\gamma_{1,i}\}$ for $i = 0, 1, \ldots, N-1$. If the sample points are aligned in space, as those in any two sample trajectories in Fig. 1 are, then $\{\gamma_{0,i}\}$ and $\{\gamma_{1,i}\}$ can be modeled as random vectors, both of which have the same N-dimensional joint distributions. Consider a third set of measurements $\{\gamma_{2,i}\}$ made along any other stretch of street.

Given an appropriately chosen metric space, it is expected that the measurement vectors possess two fundamental properties: 1) the distance between $\{\gamma_{0,i}\}$ and $\{\gamma_{1,i}\}$ is small and 2) the distances between $\{\gamma_{0,i}\}$ and $\{\gamma_{2,i}\}$ and between $\{\gamma_{1,i}\}$ and $\{\gamma_{2,i}\}$ are both significantly larger than the distance between $\{\gamma_{0,i}\}$ and $\{\gamma_{1,i}\}$. This regularity and repeatability in the signal strength measurements is not exploited in current cellular systems. We have investigated ways to exploit it using statistical pattern-recognition techniques.

B. Assumptions about the Signal Strength Samples

Consecutive signal strength samples of the small-scale fading are approximately independent if spaced more than half a wavelength apart [16], []. Since our algorithm is designed to work with independent samples of the small-scale fading, it is desirable for the spacing between consecutive samples to be greater than half a wavelength. However, the spacing should not be very much greater than half a wavelength, in order not to lose useful data. The spacing between consecutive samples should therefore be fixed and small and have dimensions of length. Note that this is *spatial sampling*, not the more common *temporal sampling*, where the spacing between consecutive samples has dimensions of time. Another reason for using spatial sampling is that it makes the pattern classification simpler, since samples are easier to "align" in doing pattern matching. If the speed of the subscriber is variable, we assume that this can be estimated so that spatial sampling can be performed by varying the rate of sampling proportionately.

To reduce the large variation in signal strength caused by small-scale fading, the measured samples are averaged. The term *signal strength indicator* (SSI) will describe the preprocessed sample values (spatial samples averaged over a few samples). SSI will be in dB unless otherwise stated. SSI can be related to the uplink and measured by a base station or related to the downlink and measured by a user. We will use "SSI associated with a base station" to mean either of these cases when the difference is irrelevant to a particular discussion. Note that we do *not* assume that small-scale fading is "averaged out." Although a convenient assumption to make, it is not a good model of reality [17]. Note that SSI can be used as the basis of signal strength patterns because it shares the structural regularity properties of unaveraged signal strength measurements, but with less variability.

C. Definitions, Notation, and Terminology

A *pattern* is "a distinctive arrangement of structural elements" [18]. For the case of signal strength patterns, the structural elements are the SSIs, and the distinctive arrangements are a result of the structural characteristics of local mean signal power level (Section III-A) that we exploit. We define a *feature matrix* to be an $\mathcal{M} \times N_p$ matrix of SSIs with the following properties: each row contains consecutive SSIs associated with one base station, and the SSIs in any two different rows are measured at approximately the same time. The space in which feature matrices lie is the *feature space*. Feature matrices are a compact and convenient representation of signal strength patterns. The two fundamental properties of the signal strength patterns (Section III-A) enable pattern classes to be defined according to similarity or closeness according to some

distance metric. Members of a *pattern class*, then, belong to a set of patterns each of whose feature matrix lies within a region in feature space. Our *pattern classification* problem consists of identifying the region in feature space into which a given pattern falls.

Each pattern class is associated, in one-to-one correspondence, with three objects: a stretch of street, a set of base stations, and a pattern action. Members of each class can be represented by feature matrices, which contain SSIs from certain base stations within the corresponding set, made along that corresponding stretch of street. The unique pattern class specification of a specific pattern class specifies the following.

1) *Pattern Stretch.* This is a stretch of street (N_p samples long) for which the system wants to be informed whenever a user traverses it. For example, this pattern may signal the need for a handoff-related action. Representative members of the class (template patterns of signal strength measurements) are obtained by making measurements along the pattern stretch. The components of the template patterns are the SSI vectors measured from surrounding base stations while a user is traveling along that pattern stretch. The list of base stations included, and the ordering of the base stations, is specified in the *included base station set*, defined below. For convenience, *locating a pattern* at a particular location will sometimes be used to mean defining a pattern class whose pattern stretch is at that location. Pattern stretches are inherently directional, going from the *stretch start* to the *stretch end*.
2) *Pattern Action.* This is the action associated with the pattern class, such as "Change billing rate" or "Reduce transmit power." The action may also be handoff-related, such as "Handoff from base station 0 to base station 1." In the case of multistate handoff decision algorithms (such as the two-state algorithm we are introducing for the corner-turn scenario), the action may be "Transition from State 1 to State 2."
3) *Included Base Station Set.* This is a set of $\mathcal{M}$ base stations. Vectors of SSI measurements from these $\mathcal{M}$ base stations made while traversing the pattern stretch can be arranged into a feature matrix of the pattern class as follows:

$$\mathbf{P} = \begin{bmatrix} \text{SSI vector of 1st base station} \\ \text{SSI vector of 2nd base station} \\ \vdots \quad \vdots \quad \vdots \\ \text{SSI vector of } \mathcal{M}\text{th base station} \end{bmatrix}. \tag{1}$$

More useful concepts and terminology follow.

1) *Pattern Instance.*[2] A pattern instance is a feature matrix obtained by observing the SSI vectors from the included base station list over a pattern stretch. Types of pattern instances include the following.
   a) *Training Pattern Instance* or *Template*. This is a pattern instance that is observed in training. If there are several templates (e.g., L, where $L > 1$) for a particular class, they will be indexed by the *instance number* $(1, \ldots, L)$.
   b) *Sliding Window Pattern Instance*. This is a pattern instance within a block of parallel sliding windows, as illustrated in Fig. 2 by the sliding window pattern instance $\mathbf{P_0}$. Fig. 2 is described in Section III-D.
2) *Pattern Set.* This is the set of all pattern classes associated with a particular base station. These pattern classes are the ones whose pattern actions apply to users for whom the particular base station is the serving base station. Distributed pattern classification, with independent processing in the different base stations, can be implemented, as discussed in Section III-D. Each base station handles the patterns in its own pattern set. The size of the pattern set of base station x is denoted by $[\Omega]_x$.

Elements of pattern matrices are represented by $\mathbf{P}^{(c)}_{l,m,n}$, where $c \in 1, \ldots, \Omega$ is the pattern number, $l \in 1, \ldots, L$ is the index of the pattern instance (or $l = 0$ if it is not a training pattern instance), $m \in 1, \ldots, \mathcal{M}$ is the base station index, and $n \in 1, \ldots, S$ is the sample index ordered from the stretch start to the stretch end. Then SSI vectors are denoted by $\mathbf{P}^{(c)}_{l,m}$ and pattern instances are denoted by $\mathbf{P}^{(c)}_{l}$. Where necessary to distinguish between patterns or variables referring to different pattern sets, we write $[]_q$ around the pattern or variable, where q is the pattern set index. Examples are $[\mathbf{P}^{(c)}_{l,m,n}]_q$ and $[\Omega]_q$.

D. Sliding Window Pattern Recognition

Each base station oversees its pattern set of Ω patterns. The pattern set is different for each base station, which handles pattern classification for only the pattern classes in its pattern set, independent of the other base stations. The included base-station set of all Ω patterns in a base station's pattern set is the same set of $\mathcal{M}$ base stations (however, the set may be different from that of the patterns of other base stations), but the patterns are otherwise independent, with their own pattern stretches and actions. Several training pattern instances of each class are used as templates. Reference [2] discusses training in more detail.

Fig. 2 illustrates how sliding window pattern-recognition works in the system operation phase. For each user, a sliding window of SSI measurements is kept at the serving base station. At regular intervals, the sliding window is updated, as the oldest samples from the $\mathcal{M}$ base stations are discarded and a fresh set of samples is transmitted from the user terminal to replace them. The contents of the sliding window are a sliding window pattern instance, and an attempt is made to classify it as belonging to one of the Ω pattern classes or none of them (the null classification). The null classification will normally be very likely, especially if Ω is small, because most of the time the user will then not be in one of the locations corresponding to a pattern class.

Decision rules like the nearest neighbor decision rule [19] can be used for classification. The classification method we use has a nearest neighbor decision rule with a null classification option. The null classification is made whenever distance to the nearest neighbor exceeds a *match threshold*. Otherwise, the class of the nearest neighbor is selected. The selected distance metric between points in pattern space is a normalized squared Euclidean

[2]This is also known as a *sample* in pattern-recognition terminology. But we use this term to avoid confusion between samples in this sense and samples in the sense of individual signal strength measurement samples.

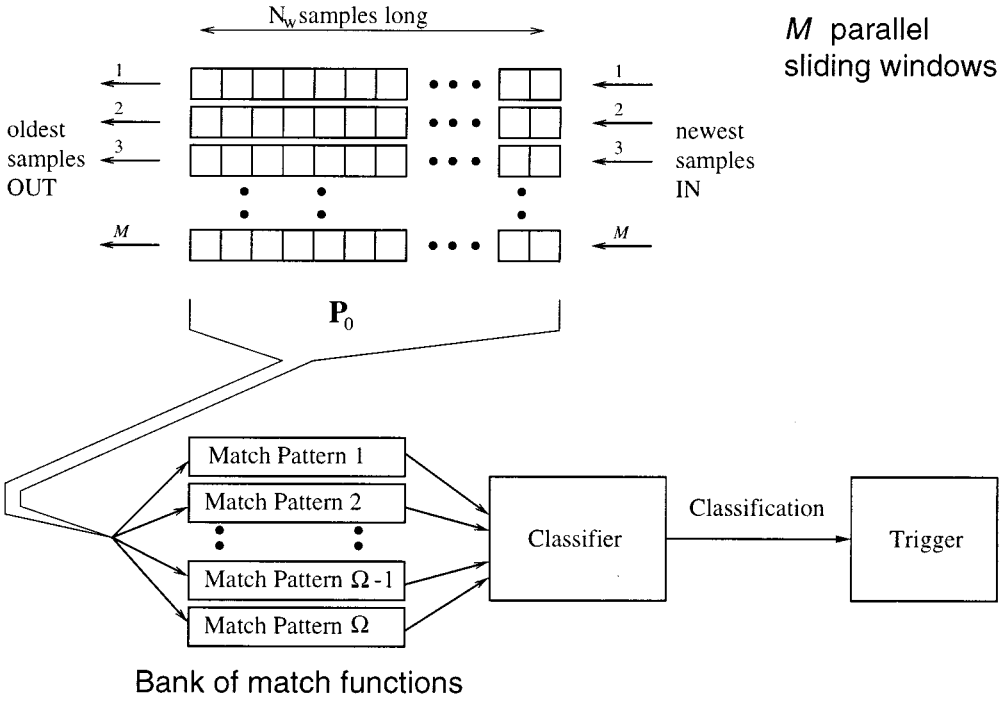

Fig. 2. Sliding window pattern-recognition block diagram.

distance. The distance of a sliding window pattern instance $\mathbf{P}_0$ to a pattern class $\mathbf{P}^{(c)}$ is taken to be the minimum of the normalized squared Euclidean distances between $\mathbf{P}_0$ and each of the L members (templates) of $\mathbf{P}^{(c)}$

$$M\left(\mathbf{P}^{(c)}, \mathbf{P}_0\right) = \frac{1}{N_p \mathcal{M}} \min_{l \in 1, 2, \cdots, L} \sum_{m=1}^{\mathcal{M}} \sum_{n=1}^{N_p} \cdot \left(\mathbf{P}^{(c)}_{l,m,n} - \mathbf{P}_{0,m,n}\right)^2. \quad (2)$$

The distance computations may be performed by *match functions*. Fig. 2 illustrates a bank of match functions, one for each of the Ω pattern classes used as input to the classifier. If the sliding window pattern instance is classified as belonging to a particular pattern class, instead of the null class, the information is passed to the trigger. The trigger implements the pattern action associated with that pattern class.

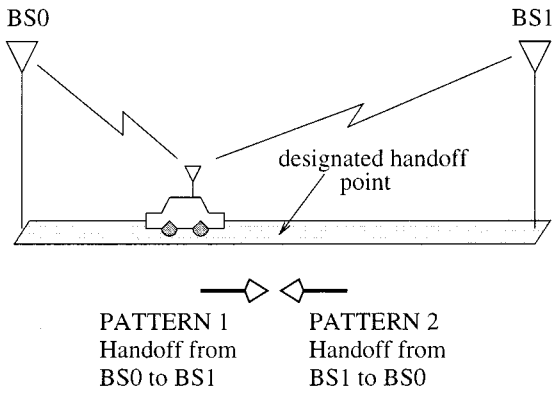

Fig. 3. The canonical handoff problem setup for PRHO. Without the patterns in the picture, this would be the canonical handoff scenario, which is often used in the literature as an unofficial testbed for handoff algorithm performance. The user moves in a straight line between two base stations, and one handoff somewhere near the midpoint of the two base stations is desired.

IV. A Two-State Handoff Algorithm

The trigger block in Fig. 2 can in general perform many functions. For example, it may trigger a state transition or a handoff to a given base station. What happens when a particular pattern class is matched is specified by the pattern action of the matched pattern class. We have simulated both THO and PRHO in a canonical straight-line handoff scenario, as illustrated in Fig. 3. This handoff scenario is useful for performance comparisons with THO because the performance of THO in this scenario has been treated in the literature (for example, [1] and [20]). A version of PRHO we have tried in this scenario has two patterns, as shown in Fig. 3. "Pattern 1" and "Pattern 2" refer to two pattern classes, one each in the pattern sets of BS0 and BS1. "Pattern 1," pointing from BS0 to BS1, belongs to the pattern set of BS0 and is located to the left of, and pointing to, the desired handoff point ("designated handoff point"). "Pattern 2," on the other hand, points the opposite direction, belongs to the pattern set of BS1, and is to the right of, and pointing to, the desired handoff point. The pattern action of "Pattern 1" is simply to handoff to BS1, while that of "Pattern 2" is to handoff to BS0. Hence, a one-state decision process is used. In general, PRHO can be used with a multiple-state decision process, as will be discussed shortly. Nevertheless, we have found that even the basic PRHO design, which we call vanilla PRHO, can achieve high accuracy (practically no unnecessary handoffs) and timeliness (negligible delay) while THO cannot avoid the tradeoff between unnecessary handoffs and decision delay regardless of hysteresis threshold. Our results are discussed in [2].

These are encouraging results. However, PRHO has the potential to handle even more challenging handoff decision

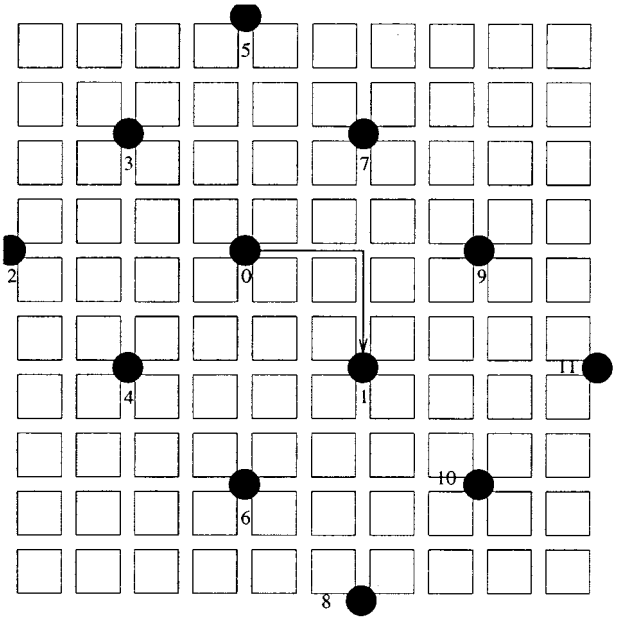

Fig. 4. Base-station arrangement for computer simulations. The squares are city blocks and the circles are base stations. The arrow between BS0 and BS1 indicates the path of motion of the user. When the serving base station is BS0, the included base station set is {0, 1, 2, 3, 4, 5, 6, 7, 9}.

problems well. An example of such a handoff scenario is the turn-around-the-corner handoff scenario shown in Fig. 4 and introduced in Section II. The system is a microcellular system in the downtown area of a city. The user enjoys a good link with the original serving base station right until it reaches the corner. After the user turns around the corner, the signal strength from the original serving base station can rapidly drop by 15–25 dB. The challenge is that a very quick handoff decision is needed right after the corner is turned, but the system cannot know if the user is going to turn the corner until after that happens. So the pattern needs to be located *after* the corner. But pattern stretches need to be long enough to keep the probability of misclassification low. In the system simulated, pattern stretches needed to be on the order of half a block length (about 50 m) to perform well. So if we only place pattern stretches after the corner, the handoff decisions would be made only half a block length away from the corner, when the call may have already been dropped due to poor signal quality. Therefore, problems are encountered if the regular patterns are placed only before the corner or only after the corner.

Our solution is to locate a regular pattern before the corner and special patterns after it and to use a two-state decision process (two-state decision machine in the trigger box in Fig. 2). The two states are *regular* and *alert*. The pattern action of the pattern located before the corner is not to perform any particular handoff but to go to the alert state. In the alert state, only three short patterns are looked for, corresponding to the user's going straight, turning left and turning right. The pattern action of *these* patterns would be the appropriate handoff. The system leaves the alert state and returns to the regular state when either a short pattern is matched or when the alert timer runs out. The purpose of the alert timer is to avoid having the system stuck in the alert state in the rare event that none of

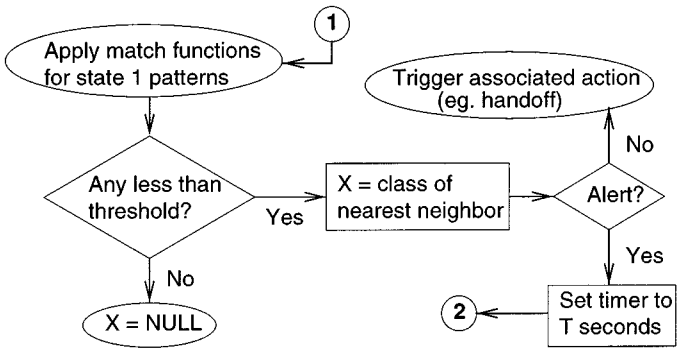

Fig. 5. Flow diagram of state 1 (regular state) of the two-state handoff algorithm. X represents the pattern class match.

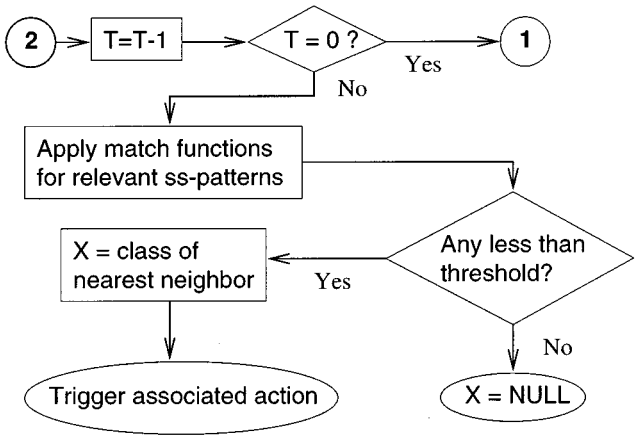

Fig. 6. Flow diagram of state 2 (alert state) of the two-state handoff algorithm. X represents the pattern class match. ss_patterns refers to second-state patterns, i.e., the alert patterns.

the short patterns is matched. This design also incorporates two regular patterns after the corner, to correct any erroneous decision made from mismatching of the short patterns. These are called the backup patterns. The pattern placement of the solution is illustrated in Fig. 7. The behavior of each of the two decision states is summarized in the flow diagrams in Figs. 5 and 6.

The following should be noted.

1) Handoff signaling and execution takes time. With a two-state decision process, handoff signaling and execution can begin earlier, once the system moves into the alert state. The result would be less overall handoff execution delay.
2) Subscriber speed information is not explicitly used, except to perform spatial sampling.

V. Performance of PRHO

Simulations were run to investigate the performance of PRHO against THO. The results demonstrate that PRHO performs very well in comparison with THO. Table I lists the simulation parameters. The base-station arrangement and the path of motion of the subscriber are shown in Fig. 4. In the simulations, handoffs are assumed to be instantaneous (no network delays). The performance measures are the number of handoffs and handoff delay. In interpreting the results, it

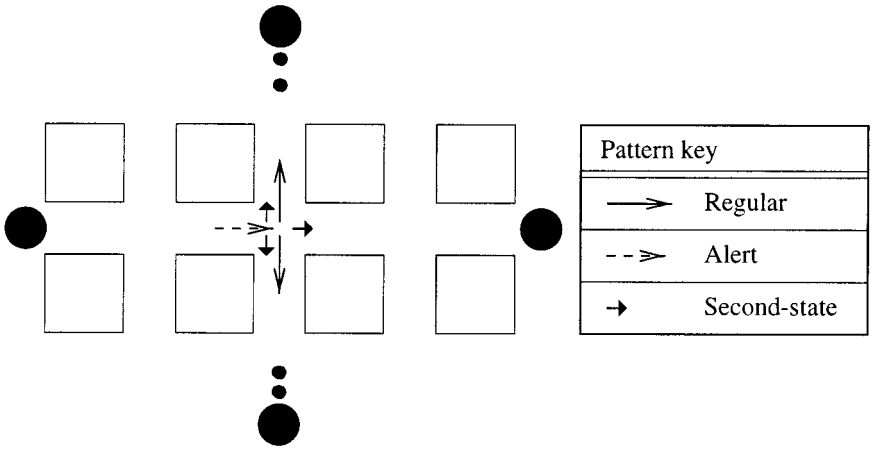

Fig. 7. Pattern layout of our two-state solution to the turning-around-the-corner handoff problem. The large solid circles are base stations and the squares are city blocks.

TABLE I
SELECTED SIMULATION PARAMETERS FOR PRHO SIMULATIONS

Parameter	Value
Spatial sampling interval	1/6 meters
(City) Block length	100.0 meters
Sliding window size, N	50 samples
No. of patts. per base station, Ω	4
Match threshold, T	variable
No. of base station components, $\mathcal{M}$	8
Std. dev. (σ) of lognormal fading	3.0
No. of independent samples input	6

should be noted that given the base-station arrangement and the path of motion of the subscriber, one handoff is needed (from BS0 to BS1). This is because the simulated microcellular propagation environment was chosen such that the BS0 can definitely provide coverage up until the corner, but not all the way until the subscriber reaches BS1. As for handoff delay, it is recognized that the handoff should occur close to the corner, not too long after it is turned. Hence, in comparing the handoff algorithms, good performance is represented by an algorithm that typically initiates one handoff close to the corner.

Selected simulation parameters are listed in Table I. Note that the value of σ (standard deviation of the lognormal fading) is somewhat on the low side. σ in practice ranges from 3 to 10 dB, depending on the environment. The low value is chosen because that is the worst case. PRHO performs better with larger σ, because there is more difference between different pattern classes. As for the sliding window size, we chose 50 samples (spaced a meter apart, i.e., a 50-m sliding window) as the size of the regular patterns. In our investigations, this was found to be a good tradeoff between two extremes. On the one hand, a sliding window that is too short (e.g., 10 m) results in a significant percentage of missed or erroneous matches. On the other hand, the slight pattern matching performance improvement associated with larger sliding windows (e.g., 100 m) does not justify the increased storage and computation costs. A reasonable exception to the rule is for the short patterns used in the alert state. As explained in Section IV, decision speed is more critical in the alert state, so we use these short patterns (10 m long in the simulations), coupled with the backup patterns.

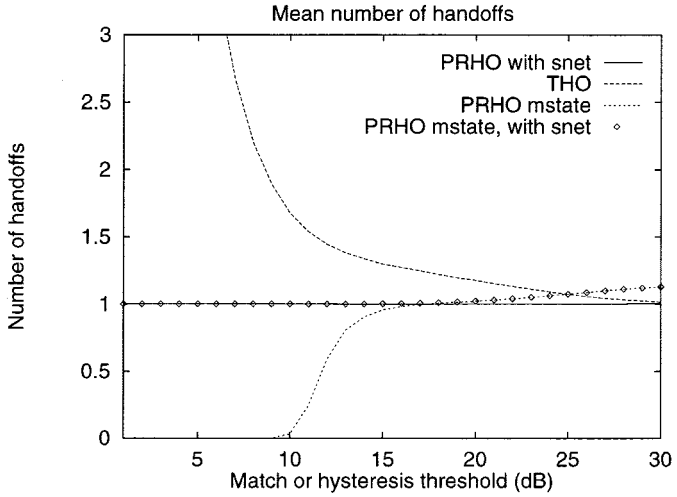

Fig. 8. Comparisons of mean number of handoffs of various algorithms.

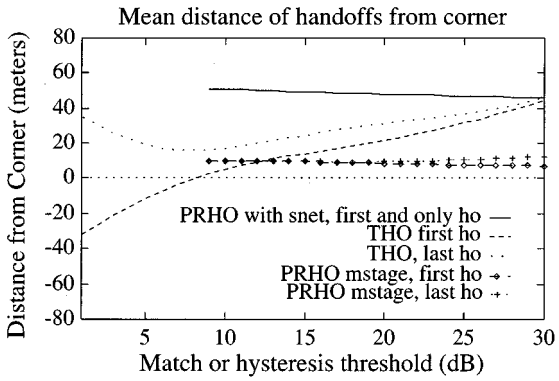

Fig. 9. Delay performance of the various algorithms.

The performance of THO and some variations of PRHO are shown in terms of the above performance measures, i.e., number of handoffs and handoff delay, in Figs. 8 and 9, respectively. In both figures, the abscissa for THO is the hysteresis threshold, while the abscissa for the PRHO variants is the match threshold.

This is not meant to imply that the hysteresis threshold and match threshold are related. In Fig. 8, the ordinate shows the mean number of handoffs for the various algorithms. For the situation simulated, it is desirable that the mean number of handoffs be approximately one. However, just looking at the mean number of handoffs is not enough, since a handoff that is too late may be useless. Hence, delay also needs to be considered. In Fig. 9, the ordinate shows the mean delay. We define delay (of a handoff) to be the number of samples after the corner that it takes for the handoff to occur. The mean delay is averaged over those runs in which at least one handoff occurred. Thus, for example, some of the curves in Fig. 8 only have meaningful values for a range of thresholds and are only plotted over that range. *First-handoff delay* is the number of samples after the corner at which the first decision is made to handoff, and *final-handoff delay* is the number of samples after the corner at which the last decision is made to handoff. The first and final handoffs are different if multiple handoffs occur. If handoff occurs before the corner is reached, it is assigned a negative delay value.

It can be seen from Fig. 8 that for small values of hysteresis threshold, THO has many unnecessary handoffs. These waste network resources. Also, handoff back to BS0 some distance after the corner might result in a dropped call. For larger values of hysteresis threshold—for example, 10 dB and greater—THO performs fewer handoffs. However, looking at Fig. 9, it can be seen that delay becomes a problem for these larger values of hysteresis threshold. On the average, for a hysteresis threshold of 10 dB, the final handoff to the correct base station occurs 20 m after the corner. This location only gets further away from the corner as the hysteresis threshold increases. The two figures together therefore illustrate the dilemma of THO in dealing with the corner-turning situation. For smaller values of hysteresis threshold, many unnecessary handoffs occur, whereas for larger values, delay becomes too high. In either case, the performance is poor.

On the other hand, the two-state algorithm we described earlier, *PRHO mstate*, has excellent delay performance, completing handoffs to BS1 within 10 m of the corner. The number of handoffs, though, depends on the match threshold. It has no handoffs for match thresholds that are too small. So these small values should not be used. As the match threshold increases, the algorithm begins to perform better. In particular, for a match threshold of 15, over 95% of the time, a timely handoff is made to BS1 as desired. For a match threshold of 17, it goes up to 99%. For larger match thresholds, the mean number of handoffs is slightly greater than one. This is because the balance shifts from erring on the side of possibly missing a pattern match to avoid unnecessary handoffs by using a smaller match threshold to allowing a small percentage of unnecessary handoffs to avoid missing a pattern match by using a larger match threshold. The exact choice used in a system will depend on the balance that the system operator desires. In any case, this algorithm can achieve accurate and timely handoffs over 95% of the time or more. In comparison, THO and vanilla PRHO (represented by *PRHO with snet*, a slight variant on vanilla PRHO that will be explained in the next paragraph) do not come close to this level of accuracy and timeliness as consistently.

The other two curves, *PRHO with snet* and *PRHO mstate, with snet*, incorporate a fail-safe mechanism, a safety net we call "snet" for short. With the fail-safe mechanism, handoffs are triggered as usual by pattern class matches. However, handoffs are also triggered when the relative signal strength of another base station exceeds that of the serving base station by a large amount, 35 dB in this case. Because this hysteresis threshold is so large, it serves mainly to catch the few cases in which the user ends up with the wrong serving base station because of pattern misclassification. The fact that *PRHO mstate, with snet* is so close to *PRHO mstate* for match thresholds over 15 dB confirms that in those cases, the two-state algorithm by itself makes the right decisions most of the time, without needing the fail-safe mechanism. The delay performance of *PRHO mstate, with snet* is virtually identical to that of *PRHO mstate* and is therefore not plotted in Fig. 9 to reduce clutter. From comparing *PRHO mstate* with *PRHO mstate, with snet*, it can be concluded that *PRHO mstate* works so well that the fail-safe mechanism becomes redundant.

PRHO with snet is the case in which PRHO is used with regular-length patterns placed after the corner, together with a fail-safe mechanism. The reason for the inclusion of *PRHO with snet* is so that the performance of this regular one-state PRHO can be compared with that of *PRHO mstate* to demonstrate why the two-state PRHO works so much better than the one-state PRHO in this case. The mean number of handoffs for *PRHO with snet* is one because either the regular-length patterns trigger the right handoff or the fail-safe mechanism is triggered, if the match threshold is too low. Does this mean that the *PRHO with snet* has the best performance? No, because for *PRHO with snet*, the delay performance is unacceptable. The handoffs tend to be triggered too far after the corner (50 m, as seen in Fig. 9) when the call may have already been dropped. Therefore, *PRHO, mstate* is the only algorithm with acceptable delay performance for a wide range of match thresholds, with the handoffs occurring around 10 m or less from the corner.

A. Discussion of Simulation Results

The simulation results justify the two-state PRHO design. The following points can be observed.

1) The shortness of the second-state patterns means they can be matched very close to the corner. This mitigates the decision delay problem. THO, on the other hand, would only be able to do this with Δ very small and so would have to suffer many unnecessary handoffs. But the two-state PRHO also improves on the case where PRHO is used with just regular patterns after the corner (vanilla PRHO), since delay for that case is too high, as seen in the simulation results in the next section. On the other hand, vanilla PRHO would be able to achieve the same low delay if three short regular patterns are used after the corner. However, short patterns are more prone to mismatch. The two-state design tightly controls the period of time in which these short patterns are looked for, which is only while in the second state. This eliminates unnecessary spurious matches to the short patterns elsewhere and provides the option of using a different match threshold in the second state.

2) Since the second-state patterns are short, there is a small probability of misclassification in the second state. This probability is not insignificant when large values of match threshold are used, as seen in the simulation results. The backup patterns are used to ensure that accuracy is maintained and the user ends up with the right serving base station. The only cost in these cases is an additional handoff (to BS7) and some delay before the user hands off. For example, through certain techniques, e.g., making the "forward" pointing short pattern slightly longer than the ones representing right and left turns, it can be ensured that mismatch errors will only be between the left-turn and right-turn short patterns. So in the unlikely event that a right turn is made but the left-turn short pattern is erroneously matched, a handoff to the wrong base station will occur, but only temporarily—as soon as the associated backup pattern is matched, a second handoff occurs, this time to the correct base station. Most of the time, this extra handoff is not necessary, but when it is, the delay before handing off to BS1 is not critical if the extra handoff is to BS7, since the signal strength of BS7 is likely to be usable further along the street toward BS1 than is that of BS0.

The simulation results also provide guidance for setting the thresholds. In using PRHO, unnecessary handoffs (in addition to the necessary handoff between BS0 and BS1) occur when the match thresholds are too *high*. Handoffs would then be triggered too often, sometimes in false alarms. However, a genuine pattern match can be missed when the thresholds are too *low*. Handoffs would rarely be triggered, and sometimes not triggered when they should be. Therefore, careful setting of the match thresholds in PRHO is important. However, the setting is not as critical as the settings of the hysteresis thresholds in THO.

VI. Concluding Remarks

THOs trade off unnecessary handoffs with delay in the decision process. PRHOs have been introduced and compared with THOs and found to be able to operate with fewer unnecessary handoffs and less delay in execution in a variety of handoff scenarios. In this paper, we have described a particularly challenging handoff scenario. We also describe how a multistate PRHO works well in meeting the challenge where THO and vanilla PRHO perform poorly.

In practice, traditional handoff algorithms are adequate for many situations. However, in certain situations—when users are turning around certain corners, for example—they may perform very poorly, resulting in frequent dropped calls. Special measures can be taken for these special situations, and the two-state PRHO is one example of something that can be applied. For each base station, there may be a few such situations for which such special care is desirable, so each base station would not have to perform pattern recognition for more than just a couple of situations at most. Hence, the memory storage requirements of the algorithm should not be a problem. Furthermore, the computational complexity need not be excessive, since a base station would at any given time be attempting matches with a subset of the patterns in its memory banks. In particular, it would typically only look for matches with state 1 patterns most of the time, and for a limited selection of state 2 patterns when in the alert state. However, the kind of measurements required (spatial sampling of signal strength from a set of base stations) may not be available in existing cellular systems. If this type of algorithm were to be adopted in future systems, changes would need to be made in the measurement procedures as well.

Future work may involve relaxing the assumption that spatial sampling is performed and investigating the possible performance degradation resulting from not having spatial sampling. It may also be of interest to add the effects of cochannel interference into the modeling, using insights from recent research like [21].

Acknowledgment

The authors are grateful to the anonymous reviewers for insightful suggestions that helped improve the clarity and presentation of the material in this paper.

References

[1] R. Vijayan and J. M. Holtzman, "A model for analyzing handoff algorithms," *IEEE Trans. Veh. Technol.*, vol. 43, pp. 351–356, Aug. 1993.

[2] D. Wong and D. Cox, "A pattern recognition system for handoff algorithms," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 1301–1312, July 2000.

[3] O. Kennemann, "Pattern recognition by hidden markov models for supporting hand over decisions in the GSM system," in *Proc. 6th Nordic Seminar Digital Mobile Radio Communications*, June 1994, pp. 195–202.

[4] H. Maturino-Lozoya, D. Munoz-Rodríguez, and H. Tawfik, "Pattern recognition techniques in handoff and service area determination," in *Proc. IEEE Vehicular Technology Conf.*, Stockholm, Sweden, June 1994, pp. 96–100.

[5] R. Narasimhan and D. Cox, "A handoff algorithm for wireless systems using pattern recognition," in *Proc. IEEE Int. Symp. Personal, Indoor, and Mobile Radio Communications*, Boston, MA, September 1998.

[6] R. Steele, J. Whitehead, and W. C. Wong, "System aspects of cellular radio," *IEEE Commun. Mag.*, vol. 33, Jan. 1995.

[7] Research and Development Center for Radio Systems (RCR), "Personal handy phone systems,", Dec. 1993.

[8] L. J. Greenstein, N. Amitay, T.-S. Chu, L. J. Cimini, G. J. Foschini, M. J. Gans, I. Chih-Lin, A. J. Rustako, R. A. Valenzuela, and G. Vannucci, "Microcells in personal communication systems," *IEEE Commun. Mag.*, pp. 76–88, Dec. 1992.

[9] N. Amitay, "Modeling and computer simulation of wave propagation in lineal line-of-sight microcells," *IEEE Trans. Veh. Technol.*, pp. 337–342, Nov. 1992.

[10] S. Ramo, J. R. Whinnery, and T. Van Duzer, *Fields and Waves in Communication Electronics*. New York: Wiley, 1965.

[11] F. Lotse and A. Wejke, "Propagation measurements for microcells in central Stockholm," in *Proc. IEEE Vehicular Technology Conf.*, Orlando, FL, May 1990, pp. 539–541.

[12] J. Whitteker, "Measurements of path loss at 910 MHz for proposed microcell urban mobile systems," *IEEE Trans. Veh. Technol.*, vol. 37, Aug. 1988.

[13] M. D. Austin and G. L. Stüber, "Velocity adaptive handoff algorithms for microcellular systems," *IEEE Trans. Veh. Technol.*, vol. 43, pp. 549–561, August 1994.

[14] A. Murase, I. C. Symington, and E. Green, "Handover criterion for macro and microcellular systems," in *Proc. IEEE Vehicular Technology Conf.*, St. Louis, MO, May 1991, pp. 524–529.

[15] R. Verdone and A. Zanella, "Handover algorithm to counteract corner effects in microcellular mobile networks," *Electron. Lett.*, vol. 34, pp. 950–951, May 1998.

[16] W. C. JakesA. Zanella, Ed., *Microwave Mobile Communications*. New York: Wiley, 1974.

[17] D. Wong and D. Cox, "Estimating local mean signal power level in a Rayleigh fading environment," *IEEE Trans. Veh. Technol.*, vol. 48, pp. 956–959, May 1999.

[18] M. Nadler and E. Smith, *Pattern Recognition Engineering*. New York: Wiley, 1993.
[19] T. M. Cover and P. E. Hart, "Nearest neighbor pattern classification," *IEEE Trans. Inform. Theory*, vol. IT-13, pp. 21–27, Jan. 1967.
[20] R. Prakash and V. Veeravalli, "Accurate performance analysis of hard handoff algorithms," in *Proc. IEEE Int. Symp. Personal, Indoor, and Mobile Radio Communications*, Boston, MA, Sept. 1998.
[21] M. Ruggieri, F. Santucci, M. Pratesi, and F. Graziosi, "A general analysis of signal strength handover algorithms with cochannel interference," *IEEE Trans. Commun.*, vol. 48, pp. 231–241, Feb. 2000.

K. Daniel Wong (S'90–M'98) was born in Ipoh, Malaysia, in 1971. He received the B.S.E. degree in electrical engineering (with highest honors) from Princeton University, Princeton, NJ, in 1992 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1994 and 1998, respectively.

He joined Bellcore (now Telcordia Technologies) in 1998 as a Research Scientist and is currently working on wireless access to next-generation networks. His research interests include handoff algorithms in cellular systems, wireless broadband technologies and network protocols, multicarrier systems, MAC and routing protocols for wireless IP networks, and third-generation mobile systems.

Dr. Wong is a member of Tau Beta Pi, Sigma Xi, and Phi Beta Kappa. He is also an executive committee member of the IEEE New Jersey Coast Section.

Donald C. Cox (S'58–M'61–SM'72–F'79) received the B.S. and M.S. degrees in electrical engineering from the University of Nebraska in 1983.

From 1960 to 1963, he did communications system design at Wright-Patterson AFB, OH. From 1963 to 1968, he was at Stanford University doing tunnel diode amplifier design and research on microwave propagation in the troposphere. From 1968 to 1973, his research at Bell Laboratories, Holmdel, NJ, in mobile radio propagation and on high-capacity mobile radio systems provided important input to early cellular mobile radio system development, and is continuing to contribute to the evolution of digital cellular radio, personal communications systems, and cordless telephones. From 1973 to 1983, he was Supervisor of a group at Bell Laboratories that did innovative propagation and system research for millimeter-wave satellite communications. In 1978, he pioneered radio system and propagation research for low-power wireless personal communications systems. At Bell Laboratories in 1983, he organized and became the Head of the Radio and Satellite Systems Research Department that became a division in Bell Communications Research (Bellcore) with the breakup of the Bell System in 1984. He was Division Manager of the Radio Research Division until it again became a department in 1991. He continued as Executive Director of the Radio Research Department, where he championed, led, and contributed to research on all aspects of low-power wireless personal communications, entitled Universal Digital Portable Communications (UDPC). He was instrumental in evolving the extensive research results into specifications that became the U.S. standard for the wireless or personal access communications system (WACS or PACS). In 1993, he became a Professor of electrical engineering and Director of the Center for Telecommunications at Stanford University, where he continues to pursue research and teaching of wireless mobile and personal communications. He was appointed Harald Tap Friis Professor of Engineering in 1994. He is a member of Commissions B, C, and F of USNC/URSI. He is a member of the URSI Intercommission Group on Time Domain Waveform Measurements (1982–1984). He is an author or coauthor of many papers and conference presentations, including many invited and several keynote addresses, and books. He has received more than a dozen patents.

Dr. Cox is a member of Sigma Xi, Sigma Tau, Eta Kappa Nu, Phi Mu Epsilon, and the National Academy of Engineering. He is a Fellow of AAS and the Radio Club of America. He received the IEEE 1993 Alexander Graham Bell Medal "for pioneering and leadership in personal portable communications" and was a Corecipient of the 1983 International Marconi Prize in Electromagnetic Wave Propagation (Italy). He received the Bellcore Fellow award in 1991, the IEEE 1985 Morris E. Leeds Award, the IEEE Communications Society 1992 L.G. Abraham Prize Paper Award, the 1990 Communications Magazine Prize Paper Award, and the 1983 IEEE Vehicular Technology Society Paper of the Year award. He was an Associate Editor of the IEEE TRANSACTIONS ON ANTENNAS AND PROPAGATION (1983–1986) and a member of the Administrative Committee of the IEEE Antennas and Propagation Society (1986–1988). He is a Registered Professional Engineer in Ohio and Nebraska.