Multistage Parallel Interference Cancellation: Convergence Behavior and Improved Performance Through Limit Cycle Mitigation

D. Richard Brown, III, *Member, IEEE*

Abstract—This paper investigates the convergence behavior of the hard-decision multistage parallel interference cancellation (PIC) detector in synchronous code division multiple access (CDMA) communication systems with random spreading sequences. Hard-decision multistage PIC is known to possess three desirable properties for multiuser detectors: a) low computational complexity, b) low decision latency due to parallel computation, and c) good bit error rate (BER) performance due the fact that the optimum (joint maximum likelihood) symbol estimates are a fixed point of the iteration. With respect to the third property, hard-decision multistage PIC detection is also known to sometimes demonstrate two modes of undesirable convergence behavior: convergence to suboptimum fixed points and limit cycles. The results in this paper show that limit cycles are often the dominant source of performance degradation. To improve the performance of the hard-decision multistage PIC detector, we propose a class of limit cycle mitigation algorithms that reactively correct for limit cycles and provide a tradeoff between performance gain and increased computational complexity. Computer simulations suggest that significant performance gains may be possible in some cases with only modest increases in computational complexity.

Index Terms—Code division multiple access, interference suppression, maximum likelihood detection, neural networks.

I. INTRODUCTION

THE advent of third-generation (3-G) cellular systems based primarily on code division multiple access (CDMA) technology and the improvements in signal processing hardware over the last decade have led to a renewed interest in multiuser detection [1] as a viable method to improve the throughput and quality of cellular communication systems. The parallel interference cancellation (PIC) multiuser detector is generally perceived by researchers as one of the most promising approaches [2], [3] and has been the subject of extensive research recently due to its applicability to 3-G cellular standards [4]. PIC multiuser detection was first introduced for CDMA communication systems in [5] and [6] as the "multistage detector" and was shown to have low computational complexity, good performance, and close connections to the optimum joint maximum likelihood detector. More recently, several companies

have begun to develop iterative PIC-based processors/systems for potential deployment in 3-G cellular base stations [7].

A key feature of Varanasi and Aazhang's original multistage PIC detector is that the tentative decisions at the output of each stage of the detector are *hard* decisions. Much like hard-decision and soft-decision channel decoding, various modifications to the hard-decision multistage PIC detector have since been proposed that instead use soft tentative decisions. As one example of this approach, the linear PIC detector (first described in [8]) replaces the nonlinear sgn($\cdot$) function of the hard-decision PIC detector with a linear mapping. The performance of the linear PIC detector has been extensively investigated, e.g., [9]–[12]. Other examples of the soft-decision approach are partial interference cancellation [13], weighted linear/nonlinear cancellation [14], linear clipping and deadzone nonlinearities [15], [16], and sigmoidal interference cancellation nonlinearities [17]–[19].

While soft-decision approaches have been shown to outperform the original hard-decision multistage PIC detector in some cases, the hard-decision detector remains important for several reasons. First, the hard-decision detector tends to require less computational resources, at least on a per-iteration basis, than the soft-decision approaches. This is especially true in CDMA communication systems with antipodal modulation [binary phase shift keying (BPSK) or quaternary phase shift keying (QPSK)]. Second, as discussed in this paper, it is quite simple to determine when the hard-decision multistage PIC detector has converged to a stable state. It is more difficult, in general, to determine when convergence has occurred in a soft-decision detector. Finally, the optimum (joint maximum likelihood) decisions are known to be a fixed point of the hard-decision multistage PIC detector. While convergence to the optimum fixed point is not guaranteed, it is still a desirable property and intuitively explains the near-optimum performance of the hard-decision multistage PIC detector in some cases. This property is lost when soft tentative decisions are used.

In addition to the original work in [5] and [6], various performance aspects of the hard-decision PIC detector have also been investigated in [20]–[24]. These investigations have primarily focused on the overall output performance of the detector and not on the dynamics or internal structure of the iteration. While good performance is seen in many cases, it is clear that the BER performance of the hard-decision multistage PIC detector is not equivalent to that of the optimum detector, even for an infinite number of PIC stages. What is less clear is why. The first main contribution of this paper is an investigation into this question through a study of the dynamics of the hard-decision multistage

PIC detector. We show that the suboptimum performance of the hard-decision multistage PIC detector is due to the iteration potentially possessing one or more suboptimum fixed point and/or limit cycle attractors. We also show that limit cycles are often the dominant cause of poor performance.

Based on these findings, the second main contribution of this paper is the development of a new approach toward improving the performance of the hard-decision multistage PIC detector. Research to date on improving the performance of PIC detection can generally be classified into soft-decision approaches or a modified initialization approaches (e.g. initializing the hard-decision PIC detector with decorrelator outputs as proposed in [6]). The new approach described in this paper is a class of algorithms that *reactively* mitigate limit cycle behavior. The proposed algorithms do not modify the interference cancellation nonlinearity or the initialization of the hard-decision PIC detector but rather observe the output of the hard-decision PIC iteration and reactively correct for limit cycles when they are detected. The advantages of this approach are that the correction only needs to be applied when needed, the desirable properties of the original hard-decision multistage PIC detector are retained while the undesirable properties are mitigated, and the computational complexity can be kept low.

The remainder of this paper is organized as follows. Section II describes the multiuser CDMA system model. Section III describes the hard-decision multistage PIC detector and presents new results on its convergence behavior. Section IV exploits the results of the prior section to develop a class of reactive limit cycle mitigation algorithms and evaluates the tradeoff between the performance and computational complexity of these algorithms. Section V then summarizes the conclusions of this work.

II. SYSTEM MODEL

We assume a synchronous CDMA multiuser communication scenario with binary signaling, nonorthogonal transmissions, and an additive white Gaussian noise channel. The communication system model is identical to the basic synchronous CDMA model described in [1]. The number of users in the system is denoted by K and all multiuser detectors considered in this paper operate on the K-dimensional matched filterbank output given by the expression

$$\boldsymbol{y} = \boldsymbol{R}\boldsymbol{A}\boldsymbol{b} + \sigma\boldsymbol{n} \qquad (1)$$

where $\boldsymbol{R} \in \mathbb{R}^{K \times K}$ is a symmetric matrix of normalized user signature crosscorrelations such that $\boldsymbol{R}_{kk} = 1$ for $k = 1, \ldots, K$ and $|\boldsymbol{R}_{k\ell}| \leq 1$ for all $k \neq \ell$, $\boldsymbol{A} \in \mathbb{R}^{K \times K}$ is a diagonal matrix of positive real amplitudes, $\boldsymbol{b} \in \{\pm 1\}^K$ is the vector of i.i.d. equiprobable binary user symbols, σ is the standard deviation of the additive channel noise, and $\boldsymbol{n} \in \mathbb{R}^K$ represents a matched filtered, unit variance AWGN process, where $\mathrm{E}[\boldsymbol{n}] = \boldsymbol{0}$ and $\mathrm{E}[\boldsymbol{n}\boldsymbol{n}^\top] = \boldsymbol{R}$. The channel noise and user symbols are assumed to be independent.

III. HARD-DECISION MULTISTAGE PIC DETECTION

Under the assumption that the receiver knows the amplitudes and signature crosscorrelations of all the users in the system,

the hard-decision multistage PIC detector's output after iteration $m + 1$ is given in vector form as [6]

$$\boldsymbol{d}(m+1) = \mathrm{sgn}\left(\boldsymbol{y} - (\boldsymbol{R} - \boldsymbol{I})\boldsymbol{A}\boldsymbol{d}(m)\right) \qquad (2)$$

where $\boldsymbol{d}(m)$ is the K-vector of tentative binary decisions at the output of the mth iteration, and $\mathrm{sgn}(\cdot)$ is the elementwise sign operator defined as

$$\mathrm{sgn}(x) = \begin{cases} -1, & x < 0 \\ +1, & x \geq 0 \end{cases} \qquad (3)$$

for $x \in \mathbb{R}$. Typically, the PIC iteration is initialized by setting $\boldsymbol{d}(0) = \mathrm{sgn}(\boldsymbol{y})$. The multistage PIC detector's final decisions may occur at some pre-determined final iteration M or, as is the case in this paper, the iteration may be monitored such that final decisions are generated upon convergence of the iteration.

A. Connections to Neural Networks

Despite the computational and conceptual simplicity of (2), little is actually known about the dynamics of the iteration. In this section, we describe the connections between the hard-decision multistage PIC detector and Hopfield neural networks (HNNs) in order to leverage this relatively large body of theory. It was first shown in [25] that the hard-decision multistage PIC detector is a particular case of a discrete HNN. We briefly review this result here and then describe some of the key properties from the neural network literature as they apply to the dynamics of the hard-decision PIC iteration.

HNNs were first proposed by Hopfield in 1982 [26] as a method of creating a system of "neurons" (also called "nodes") capable of performing certain "computational tasks." Despite this somewhat vague description, HNNs have since been applied to a variety of specific computational problems, the most common of which are the *content addressable memory* (as originally described in [26]) and a class of *combinatorial optimization* problems [27]. HNNs can be written with continuous or discrete time dynamics and the nodes of a HNN can be continuous or discrete valued. Denoting the number of nodes as K and the ith node as as x_i with $i \in \{1, \ldots, K\}$, each node x_i has an associated threshold value $\theta_i \in \mathbb{R}$ and each pair of nodes (x_i, x_j) has an associated connection weight $w_{ij} \in \mathbb{R}$. The ith node in a discrete valued, discrete time HNN is updated according to the rule $x_i(m+1) = \mathrm{sgn}(\sum_{j=1}^{K} w_{ij}x_j(m) - \theta_i)$, where m is the discrete time index. In *fully parallel* operation, all nodes are updated simultaneously, and the update can be written in vector form as $\boldsymbol{x}(m+1) = \mathrm{sgn}(\boldsymbol{W}\boldsymbol{x}(m) - \boldsymbol{\theta})$. An HNN is *simple* if all self connections are equal to zero (i.e., $w_{ii} = 0$ for all i) and is *symmetric* if the connection weights satisfy $w_{ij} = w_{ji}$ for all (i, j).

In this context, it is clear that the hard-decision PIC iteration in (2) is a simple, discrete time, discrete valued HNN operating in fully parallel update mode. Although the "connection matrix" $\boldsymbol{W} = -(\boldsymbol{R} - \boldsymbol{I})\boldsymbol{A}$ in (2) is not symmetric in general, it is possible to rewrite (2) in an equivalent symmetric form. Using the

fact that $\mathrm{sgn}(x) = \mathrm{sgn}(ax)$ for all $a > 0$ and that $\boldsymbol{A}$ is a diagonal matrix with strictly positive coefficients, we can rewrite (2) equivalently as

$$\underbrace{\boldsymbol{d}(m+1)}_{\boldsymbol{x}(m+1)} = \mathrm{sgn}\left(\underbrace{\boldsymbol{Ay}}_{-\boldsymbol{\theta}} - \underbrace{\boldsymbol{A}(\boldsymbol{R}-\boldsymbol{I})\boldsymbol{A}}_{-\boldsymbol{W}}\underbrace{\boldsymbol{d}(m)}_{\boldsymbol{x}(m)}\right). \quad (4)$$

Since $(\boldsymbol{R}-\boldsymbol{I})$ is symmetric with zeros on its diagonal, the hard-decision PIC iteration is both simple and symmetric.

The explicit relationship between the hard-decision multistage PIC detector and HNNs exposes an interesting connection between the hard-decision multistage PIC detector and the optimum (joint maximum likelihood) multiuser detector [28]. Since Hopfield and Tank's work on the well-known "Traveling Salesman Problem" [29], HNNs have been used as a computationally efficient (but suboptimum) approach for solving a variety of combinatorial optimization problems [27]. Iterations of a HNN are viewed as a method for finding a minimum of a Lyapunov *energy function*, usually defined as $E(\boldsymbol{x}) = 2\boldsymbol{\theta}^\top \boldsymbol{x} - \boldsymbol{x}^\top \boldsymbol{W}\boldsymbol{x}$. Using (4), we can say that hard-decision PIC iteration is a method for finding a minimum of the energy function

$$E(\boldsymbol{d}) = -2\boldsymbol{y}^\top \boldsymbol{Ad} + \boldsymbol{d}^\top \boldsymbol{A}(\boldsymbol{R}-\boldsymbol{I})\boldsymbol{Ad}.$$

While this energy function does not have a physical meaning in the context of CDMA communication systems, it can be shown that $E(\boldsymbol{d}) = c - \Lambda(\boldsymbol{d}|\boldsymbol{y})$ where c is a constant that does not depend on $\boldsymbol{d}$ and $\Lambda(\boldsymbol{d}|\boldsymbol{y}) = 2\boldsymbol{y}^\top \boldsymbol{Ad} - \boldsymbol{d}^\top \boldsymbol{ARAd}$ is the *likelihood function* [1] representing the relative posterior likelihood that the symbol vector $\boldsymbol{d}$ was transmitted conditioned on the observation $\boldsymbol{y}$. Since minimization of the energy function is equivalent to maximization of the likelihood function, the hard-decision PIC iteration can be viewed as a HNN approach to the combinatorial optimization problem of finding the optimum (joint maximum likelihood) symbol estimates.

B. Attractors of the Hard-Decision PIC Iteration

While simulations of the hard-decision multistage PIC detector show that convergence to the optimum symbol estimates may occur frequently in some scenarios, it is a suboptimum detector due to the fact that the iteration can converge to local minima or other spurious attractors of the energy function. An important result from the HNN literature (derived from an analysis of the Lyapunov energy function) explicitly states that there are only two types of attractors possible in discrete-time symmetric HNNs.

Theorem 1 (Goles [30]): Denote the maximum period of an attractor of (4) as $P(\boldsymbol{W}, \boldsymbol{\theta})$. If $\boldsymbol{W}$ is symmetric, $P(\boldsymbol{W}, \boldsymbol{\theta}) \leq 2$.

In other words, the hard-decision multistage PIC detector must converge in a finite number of iterations to either a fixed point (i.e. $\boldsymbol{x}(m + 1) = \boldsymbol{x}(m)$) or a limit cycle of period two (i.e. $\boldsymbol{x}(m+1) = \boldsymbol{x}(m-1) \neq \boldsymbol{x}(m)$). Limit cycles with period longer than two and chaotic behavior are not possible. The following three-user example demonstrates both fixed point and period-2 limit cycle convergence behavior for the hard-decision multistage PIC detector.
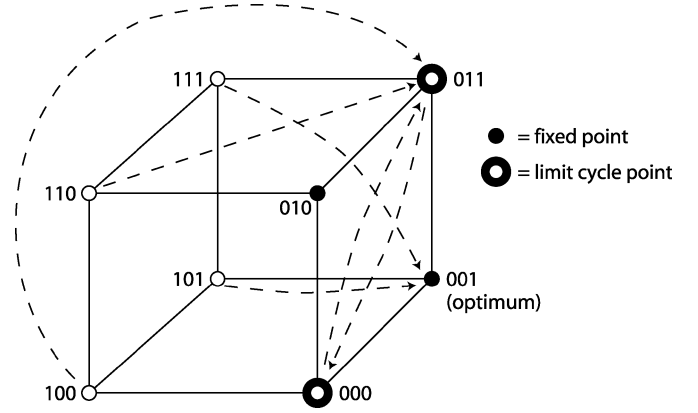

Fig. 1. Three-user hard-decision PIC iteration example. Bits shown as zero correspond to the BPSK symbol -1.

Example 1: Suppose that $\boldsymbol{A} = \boldsymbol{I}$,

$$\boldsymbol{R} = \frac{1}{3}\begin{bmatrix} 3 & 0 & -1 \\ 0 & 3 & 2 \\ -1 & 2 & 3 \end{bmatrix}$$

and that $\boldsymbol{y} = [-1.2, -0.3, 0.6]^\top$. Fig. 1 shows the eight possible states of the tentative decision vector $\boldsymbol{d}(m)$ and the flows between these states specified by (2). The iteration has two fixed points at $\boldsymbol{d} = [-1, 1, -1]^\top$ and $\boldsymbol{d} = [-1, -1, 1]^\top$ (the latter being the optimum fixed point) and two states that form a period-2 limit cycle at $\boldsymbol{d} = [-1, 1, 1]^\top$ and $\boldsymbol{d} = [-1, -1, -1]^\top$. The remaining four states are not attractors of (2).

The fact that the hard-decision multistage PIC detector converges to either a fixed point or a periodic attractor is not particularly surprising due to the deterministic nature of the update and the finite number of states. Nevertheless, the fact that periodic attractors are *always length two* is a powerful result with practical implications. Specifically, it implies that a hard-decision multistage PIC detector requires only memory of its last two states ($2K$ bits) in order to determine when an attractor has been reached. It also implies that the receiver is easily able to distinguish between fixed point convergence and convergence to a limit cycle. We will use these facts in Section IV-A in order to develop new methods for improving the performance of the hard-decision multistage PIC detector. The following subsections describe the known properties of the attractors of (2).

1) Properties of Fixed Point Attractors: Proposition 1 below describes a basic property of the fixed point attractors of (2).

Proposition 1: Given $\boldsymbol{d} \in \{\pm 1\}^K$, $\boldsymbol{d}$ is a fixed point of (2) if and only if $\boldsymbol{d}$ is a local maximum (neighborhood size of Hamming distance one) of the likelihood function $\Lambda(\boldsymbol{d}|\boldsymbol{y})$.

The proof of this proposition is given in the Appendix. This proposition implies that all fixed point attractors of the hard-decision multistage PIC detector must be separated by at least Hamming distance two (see Example 1). Consequently, the number of fixed points in (2) is upper bounded by 2^{K-1}. This bound tends to be quite loose in most cases. Proposition 1 also implies that convergence to a suboptimum fixed point results in at least two decision errors with respect to the optimum symbol estimates but that convergence to a suboptimum fixed point is
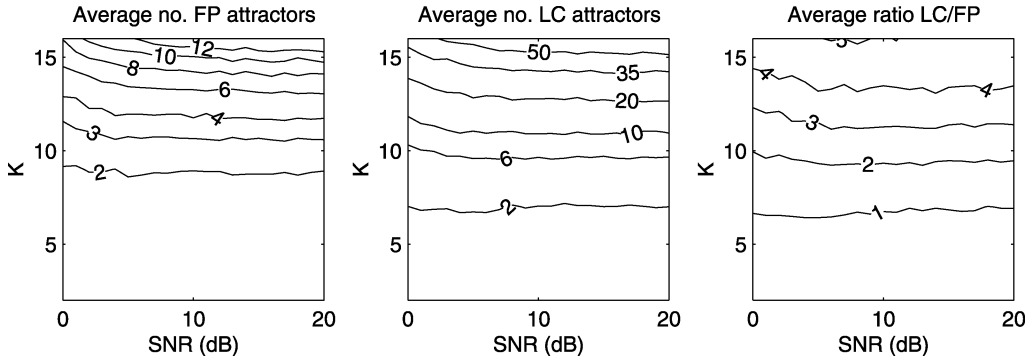
Fig. 2.   Average number of fixed point and limit cycle attractors of the hard-decision PIC iteration for a CDMA system with length $N = 16$ random spreading sequences.

better (in terms of likelihood) than convergence to any of its neighboring states.

The neural network literature also provides some insight into the properties of the fixed point attractors of (2). Due to the quadratic nature of the likelihood function, we know that there is always at least one fixed point attractor: the joint maximum likelihood symbol vector estimate corresponding to the global maximum of $\Lambda(\boldsymbol{d}|\boldsymbol{y})$. It may also be useful to know, given $\boldsymbol{W}$ and $\boldsymbol{\theta}$, if there are any additional (suboptimum) fixed points in the iteration. The neural network literature shows that this question is difficult to answer. Specifically, the question of whether there are two [31] or three [32] fixed points in a simple, symmetric HNN with fully parallel updates is $\mathcal{NP}$-complete. The problem of determining the exact number of fixed points in the same class of networks is #$\mathcal{P}$-complete [31], [32] (see [33] for a definition of the complexity class #$\mathcal{P}$).

Given that it is difficult to compute the number of fixed points in the iteration (2), it may still be useful to quantify the attraction radius of a particular attractor in the system. For instance, Example 1 shows a case where the attraction radius of both fixed points is zero. In the case of the hard-decision multistage PIC detector, a general result would give some intuition on the ability of the detector to "correct errors" in the initial decisions based on the matched filter outputs. Unfortunately, the problem of computing the attraction radius of a fixed point in a simple, symmetric HNN has been shown to be $\mathcal{NP}$-hard [34]. It was also shown in [34] that the attraction radius of a fixed point of (2) cannot even be approximated within a factor of $K^{1-\delta}$ for any fixed $\delta > 0$ in polynomial time.

In the case when the hard-decision multistage PIC detector does converge to a fixed point, it would be useful to determine whether the resulting solution is optimum or suboptimum. Not surprisingly, this is also difficult. It was shown in [35] that the question of whether there exists another state $\boldsymbol{d}'$ having lower energy (or, equivalently, greater likelihood) than the current state $\boldsymbol{d}$ turns out to be $\mathcal{NP}$-complete.

While most of the results in this section appear to be negative, they are presented here because they directly influence the strategy by which we propose to improve the performance of the hard-decision multistage PIC detector in Section IV. These results clearly state that it is difficult to predict the outcome of the iteration *a priori* and that, if a fixed point is reached, it is difficult to determine if the solution is optimum. These facts motivate the

development of the reactive limit cycle mitigation algorithms in Section IV.

2) Properties of Limit Cycle Attractors: Proposition 2 describes a basic property of the limit cycle attractors of (2).

Proposition 2: Given $\dot{\boldsymbol{d}}$ and $\ddot{\boldsymbol{d}}$ are the two states comprising a limit cycle attractor of (2), the Hamming distance between $\dot{\boldsymbol{d}}$ and $\ddot{\boldsymbol{d}}$ is at least two.

The proof of this proposition is given in the Appendix. We note that Proposition 2 does not imply that *all* of the states of the limit cycle attractors of (2) must be separated by a Hamming distance of at least two but only that the two states comprising a particular limit cycle attractor must be separated by a Hamming distance of at least two. It is possible to generate examples where a pair of limit cycle attractors (with four total states) will have neighboring states. We also note that fixed points may be neighbors of states corresponding to limit cycle attractors. This is illustrated in Example 1.

C. Convergence Behavior in CDMA Systems with Random Signature Sequences

This section presents numerical examples that illustrate the convergence behavior of the hard-decision multistage PIC detector in the case of a CDMA communication system where each user is assigned a length-N binary spreading sequence that is random, equiprobable from the set $\{\pm 1\}^N$, and independent of all other users' spreading sequences. Denoting $\boldsymbol{s}^k \in \{\pm 1\}^K$ as the kth user's random spreading sequence and $\boldsymbol{S} = [\boldsymbol{s}^1, \ldots, \boldsymbol{s}^K]$ as the $N \times K$ signature matrix, the signature cross-correlation matrix is defined as $\boldsymbol{R} = (1/N)\boldsymbol{S}^\top\boldsymbol{S}$. All of the results in this section assume that the users are all received at equal power, i.e., $\boldsymbol{A} = a\boldsymbol{I}$. The signal-to-noise ratio, which is also equal for all users, is defined as SNR $= 10\log_{10}(a^2/\sigma^2)$.

Our first result gives some intuition on the numbers of attractors present in the hard-decision PIC iteration as a function of K and SNR. Fig. 2 plots the average number of fixed point and limit cycle attractors of (2) in the case when the spreading gain $N = 16$ and also plots the average of the ratio of the number of limit cycle attractors to fixed point attractors. Since the number of attractors is computationally difficult to estimate (as discussed in Section III-B1) we performed a brute-force search over the space of 2^K states and were computationally constrained to considering only values of $K \leq 16$. Nevertheless, the results show an interesting trend. Specifically, as
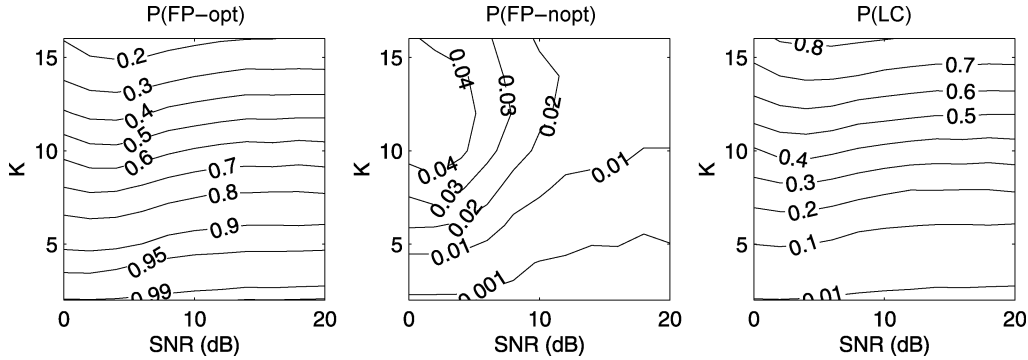
Fig. 3. Probability of the modes of convergence for hard-decision multistage PIC detection with matched filter initialization in a CDMA system with length $N = 16$ random spreading sequences. Notation: "FP-opt": optimum fixed point convergence; "FP-nopt": suboptimum fixed point convergence; "LC": period-2 limit cycle.

K approaches N, the results in Fig. 2 show that the both the number of fixed point attractors (all but one of which are suboptimum) and the number of limit cycle attractors (all of which are suboptimum) tend to increase rapidly. This implies that the multistage hard-decision PIC iteration is likely to be plagued by spurious attractors when K is close to N and that, at least intuitively, convergence to the optimum fixed point is less likely in these cases. Moreover, the results suggest that limit cycle attractors tend to be more prevalent than fixed point attractors as K approaches N and that the average ratio of limit cycle attractors to fixed point attractors can be fairly large in these cases.

Our next result considers the probability of convergence to each type of attractor under the assumption that the hard-decision multistage PIC detector is initialized with matched filter decisions. Fig. 3 plots the relative probability of each mode of convergence in the case when the spreading gain $N = 16$. Due to the fact that the results require the computation of the joint maximum likelihood bit estimates, we were again computationally constrained to considering only values of $K \leq 16$. The results in Fig. 3 confirm the intuition from Fig. 2 but also show that, when the hard-decision multistage PIC detector is initialized with matched filter decisions, the relative probability of limit cycle convegence to suboptimum fixed point convergence is even more dramatic than the results of Fig. 2 would suggest.

To confirm that these convergence trends are not an anomaly resulting from the relatively low spreading gain and/or small number of users, the next result considers the probability of limit cycle convergence for larger values of K and N. Surprisingly, the basic trends seen in Fig. 3 become even more pronounced in these cases. While it is computationally difficult to distinguish suboptimum fixed point convergence from optimum fixed point convergence for large values of K, it is easy to distinguish between limit cycles and fixed point convergence. Using this fact, Fig. 4 plots the probability of limit cycle convergence (via simulation) for the case when the all users are received at 10 dB SNR. These results suggest that limit cycle convergence is very likely for $K/N > 0.7$ when K and N are both large and that fixed point convergence is very likely for $K/N < 0.5$. Similar results are observed at other values of SNR with the main difference being that the transition occurs at slightly different values of K/N. These results also suggest that when K and N are very large, there may be a critical value β^* such that when $K/N > \beta^*$, the hard-decision multistage PIC detector

almost always converges to a period-2 limit cycle. A similar result was proved for the linear multistage PIC detector in [12], where $\beta^* \approx 0.17$. A proof of such a result for the hard-decision multistage PIC detector remains an open problem.

IV. MITIGATION OF LIMIT CYCLES

Based on the results of the prior section that suggest that limit cycles are often the dominant source of poor convergence behavior in the hard-decision multistage PIC detector, this section proposes a new approach for improving the performance of the hard-decision multistage PIC detector: reactive limit cycle mitigation. A reactive limit cycle mitigation algorithm retains the original iteration of (2) and only corrects for poor convergence behavior when it is detected. Although poor convergence behavior includes both limit cycles and suboptimum fixed points, the techniques in this section only correct for limit cycles due to a) the ease in which they are identified and b) their frequency of occurrence with respect to suboptimum fixed points, especially as $K/N \to 1$.

In the following sections we describe three reactive limit cycle mitigation algorithms with varying performance and complexity tradeoffs and then numerically compare the performance and complexity of these algorithms in CDMA systems with random signature sequences.

A. Algorithms

This section describes three reactive limit cycle algorithms with varying performance and complexity. For lack of a better naming system, we refer to these algorithms as LCM1–LCM3.

LCM1—Full Maximum Likelihood Search: When a limit cycle is detected, the LCM1 algorithm simply performs a full combinatorial optimization of the likelihood function, i.e.,

$$\hat{\boldsymbol{d}}_{\text{LCM1}} = \arg \max_{\boldsymbol{d} \in \{\pm 1\}^K} \Lambda(\boldsymbol{d}|\boldsymbol{y}). \tag{5}$$

While computing the joint maximum likelihood solution for all users' bits every time a limit cycle occurs is likely to be computationally infeasible in all but very small systems, we present it here as a benchmark because it establishes a bound on the performance that can be attained by reactive limit cycle mitigation. When the system is in an operating region where limit cycles occur frequently, e.g., when $K/N > 0.7$ in Fig. 4, the
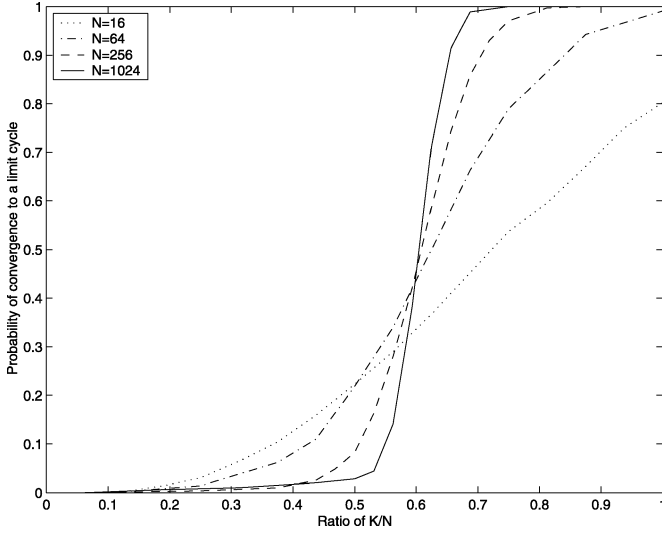
Fig. 4.   Probability of convergence to a limit cycle attractor from a matched filter initialization in a CDMA system with length $N \in \{16, 64, 256, 1024\}$ random spreading sequences and all users received at 10 dB SNR.

performance gain of this approach will be significant but the complexity will be essentially the same as optimum multiuser detection.

LCM2—Partial Maximum Likelihood Search: One of the shortcomings of the LCM1 algorithm, in addition to its high complexity, is that it does not exploit any information obtained from the results of the hard-decision PIC iteration. The LCM2 algorithm addresses this shortcoming by using the results of the hard-decision PIC iteration in order to classify users into two groups: "fixed" and "undecided." Specifically, denoting the two known states of the limit cycle as $\dot{\boldsymbol{d}}$ and $\ddot{\boldsymbol{d}}$, the fixed and undecided user index sets are defined as

$$\mathcal{F} = \left\{ k \in \{1, \ldots, K\} : \dot{d}^{(k)} = \ddot{d}^{(k)} \right\}$$
$$= \{f_1, \ldots, f_{K-\kappa}\} \text{ (fixed) and}$$
$$\mathcal{U} = \left\{ k \in \{1, \ldots, K\} : \dot{d}^{(k)} = -\ddot{d}^{(k)} \right\}$$
$$= \{u_1, \ldots, u_\kappa\} \text{ (undecided).}$$

The fixed users are the users with bit estimates that are the same in both states of the limit cycle and the undecided users are the users with bit estimates that are toggling. Note that κ denotes the number of undecided users. Denoting

$$\mathcal{X} = \Big\{ \boldsymbol{x} \in \{\pm 1\}^K : x^{(f_1)} = \dot{d}^{(f_1)}$$
$$x^{(f_2)} = \dot{d}^{(f_2)}, \ldots, x^{(f_{K-\kappa})} = \dot{d}^{(f_{K-\kappa})} \Big\}$$

as the set of all 2^κ states in $\{\pm 1\}^K$ with bits that agree with the decided users, the LCM2 algorithm then finds the maximum likelihood bit estimates for the undecided users over this set, i.e.,

$$\hat{\boldsymbol{d}}_{\text{LCM2}} = \arg \max_{\boldsymbol{d} \in \mathcal{X}} \Lambda(\boldsymbol{d}|\boldsymbol{y}). \tag{6}$$

The key difference between the LCM1 and LCM2 algorithms is that the set over which the optimization is performed in (6) is usually much smaller than the set over which the optimization is performed in (5). Intuitively, the LCM2 algorithm assumes that

the fixed bits in a period-2 limit cycle are likely to be correct and that the toggling bits are all unreliable. The LCM2 algorithm exploits this intuition to perform a joint maximum likelihood optimization only on the toggling bits, conditioned on the fixed bits, and typically at a much lower computational cost than a full, unconditional maximum likelihood search.

We note that the LCM2 algorithm is similar in spirit to the group detector proposed in [36]. In this context, the group size corresponding to the original hard-decision multistage PIC detector is one. If a limit cycle occurs, then the users are partitioned into the fixed and undecided groups of size $K - \kappa$ and κ, respectively, and a final "group detection" iteration is performed on the κ users with toggling bits. Unlike [36], however, the members, and consequently the sizes, of the groups in the LCM2 algorithm are dynamically determined by the properties of the limit cycle.

LCM3—Soft Output Combining: The final limit cycle mitigation algorithm is the least complex. When a limit cycle is detected, the LCM3 algorithm generates its bit estimates by combining the soft output statistics from the two known states of the limit cycle, $\dot{\boldsymbol{d}}$ and $\ddot{\boldsymbol{d}}$, i.e.,

$$\hat{\boldsymbol{d}}_{\text{LCM3}} = \text{sgn}(\dot{\boldsymbol{x}} + \ddot{\boldsymbol{x}}) \tag{7}$$

where $\dot{\boldsymbol{x}} = \boldsymbol{y} - (\boldsymbol{R} - \boldsymbol{I})\boldsymbol{A}\dot{\boldsymbol{d}}$, and $\ddot{\boldsymbol{x}} = \boldsymbol{y} - (\boldsymbol{R} - \boldsymbol{I})\boldsymbol{A}\ddot{\boldsymbol{d}}$. Equivalently, the LCM3 decision can be expressed as

$$\hat{\boldsymbol{d}}_{\text{LCM3}} = \text{sgn}\left(\boldsymbol{y} - (\boldsymbol{R} - \boldsymbol{I})\boldsymbol{A}\left(\frac{\dot{\boldsymbol{d}} + \ddot{\boldsymbol{d}}}{2} \right) \right). \tag{8}$$

This last expression reveals the intuition behind the LCM3 algorithm. All of the "undecided" (as defined in the LCM2 algorithm) users' decisions effectively cancel each other in (8) and all of the "fixed" users' decisions constructively combine. The LCM3 algorithm generates its final decisions by performing one hard-decision PIC iteration that attempts to cancel only the estimated interference from the fixed users and does not attempt to cancel the interference of the undecided users.

Like the LCM2 algorithm, LCM3 exploits the information contained in the limit cycle states. The complexity of this approach, however, is much lower than both the LCM1 and LCM2 algorithms.

We now have a few additional remarks.

1) The final decisions generated by the LCM1–LCM3 algorithms are not used to reinitialize the hard-decision multistage PIC detector. While this is certainly possible, it does not make sense for the LCM1 algorithm since LCM1 decisions are already a fixed point of (2). Our simulations also suggest that, for the cases tested, little or no gain is achieved by reinitializing the PIC detector with LCM outputs. We also found that it is possible to reenter *the same limit cycle* after reinitializing the PIC detector with LCM2 or LCM3 decisions.

2) It is possible to combine the LCM2 and LCM3 algorithms into a hybrid algorithm that achieves almost any point in performance/complexity space between the two approaches. Specifically, an integer parameter κ^* is specified such that $2 \le \kappa^* \le K$. When a limit cycle with κ undecided users is detected, κ is compared to this threshold.

If $\kappa \leq \kappa^*$, the more complex LCM2 algorithm is used since the number of states to search is small; otherwise, the less complex LCM3 algorithm is used. Because of the occasional occurance of large values of κ (as discussed in Section IV-C), this approach can achieve most of the performance gain of the LCM2 approach at much lower computational complexity.

B. Performance Comparison in CDMA Systems with Random Signature Sequences

This section presents performance results for the hard-decision multistage PIC detector with limit cycle mitigation and compares the performance of limit cycle mitigation to several benchmark multiuser detectors. All of the results in this section assume a CDMA communication system with random length-N spreading sequences and equal power users, as described in Section III-C.

One of the multiuser detectors considered in this performance comparison is the partial-cancellation PIC detector first proposed in [14]. The results of this section assume a three-stage partial-cancellation PIC detector, initialized with matched filter decisions, and with partial cancellation factors specified as $p_1 = 0.5$, $p_2 = 0.7$, and $p_3 = 0.9$ using the notation of [14].

The hard-decision PIC detector is also initialized with matched filter decisions. After each stage, its output is compared with the output of the prior stage to determine if the iteration converged to a fixed point. If not, the output is compared to the output of the stage twice prior to determine if the iteration entered a period-2 limit cycle. If either of these results occurs, the iteration is terminated and the hard-decision multistage PIC detector's bit estimates are set equal to the output of the final stage. If the outcome of the iteration is a fixed point, the limit cycle mitigation algorithms are not used (the unmodified PIC outputs are used as the LCM decisions). If the outcome of the iteration is a limit cycle, the appropriate LCM algorithm is run to generate their final decisions.

The first result in this section considers the BER performance of the limit cycle mitigation algorithms. Fig. 5 demonstrates the gain in BER performance obtained with limit cycle mitigation with respect to the original hard-decision multistage PIC detector and several benchmark detectors. Note that the hard-decision multistage PIC detector offers very little performance gain with respect to the matched filter detector when $K = 12$ and actually performs worse than the matched filter detector when $K = 16$. The LCM1 algorithm performs significantly better than the unmodified hard-decision multistage PIC detector and the partial-cancellation PIC detector and shows that large potential gains are possible with an effective limit cycle mitigation algorithm. The LCM2 algorithm also performs significantly better than the unmodified hard-decision multistage PIC detector as well as the partial-cancellation PIC detector and performs almost as well as the LCM1 algorithm when $K \leq 8$. The performance of the LCM2 algorithm degrades with respect to the LCM1 algorithm as $K/N \rightarrow 1$ due to the fact that the conditional optimization becomes less reliable. Finally, the LCM3 algorithm shows more modest gains with respect to the unmodified hard-decision multistage PIC detector and performs

slightly worse than the partial-cancellation PIC detector. Nevertheless, the performance gain achieved by the LCM3 algorithm with respect to the unmodified hard-decision PIC detector is achieved with almost no additional computation. Overall, the largest gains for all three limit cycle algorithms are seen when the number of users is small and the SNR is high.

The second performance result in this section looks at the BER performance of the limit cycle mitigation algorithms in a larger system and demonstrates that the hybrid LCM2/LCM3 approach described in Section IV-A can be used to achieve a desired tradeoff between performance gain and increased computational complexity. Fig. 6 shows the BER performance versus the number of users in a CDMA system with length $N = 64$ random spreading sequences. All users are received at 10 dB SNR. The partial-cancellation PIC detector and limit cycle mitigation algorithms all demonstrate BER rates significantly better than the unmodified hard-decision multistage PIC detector BER, except when the number of users is small and all algorithms are performing close to the single user bound. The two hybrid algorithms show the most performance gain (with respect to the unmodified hard-decision multistage PIC detector) when the system is approximately half-loaded, i.e., $K \approx 32$. The limit cycle mitigation algorithms tend to perform better than the partial-cancellation PIC detector when the system is lightly loaded $(K < 32)$ but, due to the fact that the hybrid LCM2/LCM3 approaches behave more like LCM3 for large K, the partial-cancellation PIC detector tends to outperform the limit cycle mitigation algorithms in the more heavily loaded cases $(K > 32)$.

This result also shows the capacity increase that can be achieved by the limit cycle mitigation algorithms for a fixed quality of service. For example, if the users require a BER of $2 \cdot 10^{-3}$ or better, the matched filter can support only two or fewer users, the unmodified hard-decision multistage PIC detector can support up to 15 users, the partial-cancellation PIC detector can support up to 21 users, and the limit cycle mitigation algorithms can support up to 22, 25, and 28 users, respectively.

C. Complexity Comparison in CDMA Systems with Random Signature Sequences

Since the essential tradeoff with any multiuser detector is complexity for performance, this section evaluates the computational complexity of the hard-decision multistage PIC detector and the additional computational complexity required by the limit cycle mitigation algorithms described in Section IV-A. The results in this section are intended to provide context for the potential performance gains of limit cycle mitigation demonstrated in Section IV-B. All of the results in this section assume a CDMA communication system with random length-N spreading sequences and equal power users received at 10 dB SNR.

For the purposes of complexity comparison, we assume that $\boldsymbol{W}$ and $\boldsymbol{\theta}$ of (4) are precomputed and available to the hard-decision PIC detector and limit cycle mitigation algorithms without any computational cost. In addition, to facilitate the comparison,
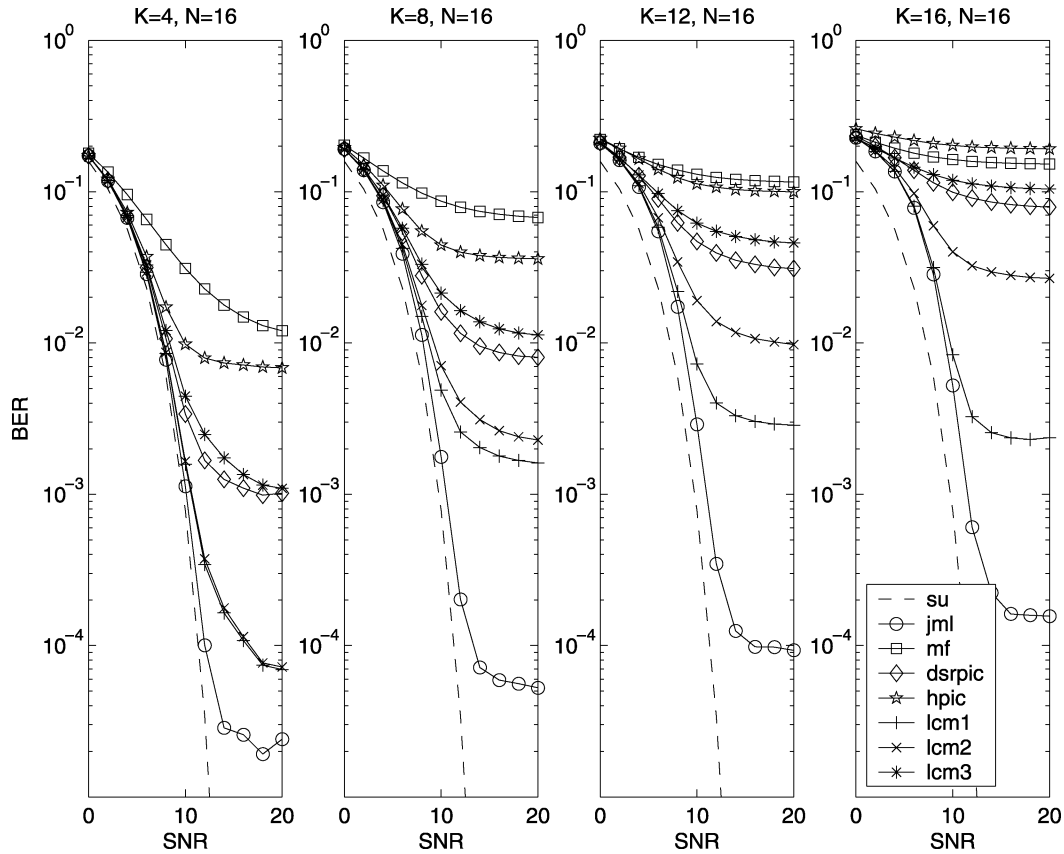
Fig. 5.   Bit error rates of hard-decision multistage PIC detection with limit cycle mitigation ($\mathsf{lcm1} - \mathsf{lcm3}$) compared with the single user bound (su), the optimum joint maximum likelihood (jml) detector, the matched filter (mf) detector, the partial-cancellation PIC detector by Divsalar *et al.* in [14] (dsrpic), and the hard-decision multistage PIC detector (hpic).
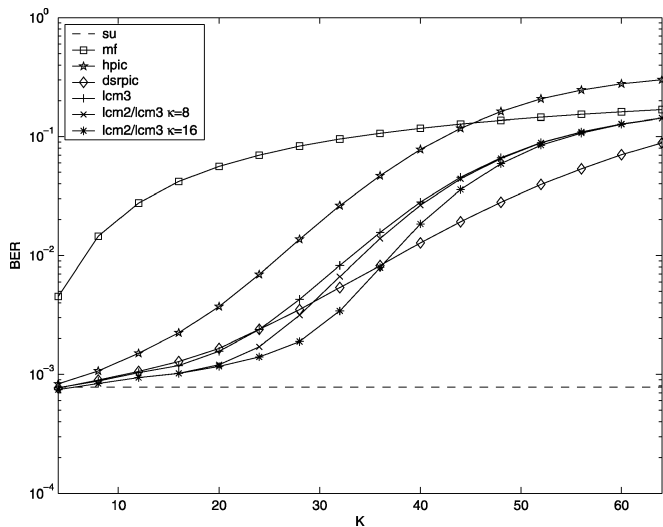

Fig. 6.   Bit error rates of hard-decision multistage PIC detection with LCM3 and hybrid LCM2/LCM3 limit cycle mitigation compared to the single user bound (su), the matched filter (mf) detector, the partial-cancellation PIC detector by Divsalar *et al.* in [14] (dsrpic), and the hard-decision multistage PIC detector (hpic).

we define a *complexity unit* as one real-binary multiplication[1] and one signed addition. When $\boldsymbol{W}$ in (4) [or, equivalently, the term $(\boldsymbol{R}-\boldsymbol{I})\boldsymbol{A}$ in (2)] is precomputed, the binary nature of $\boldsymbol{d}(m)$

[1]By "real-binary multiplication," we mean the multiplication of a real-valued number by ± 1.

implies that each iteration of the hard-decision PIC detector can be computed entirely with real-binary multiplications and signed additions. No real-real multiplications are required. This property is in contrast to linear detectors like the linear MMSE detector and the soft-decision or partial-cancellation versions of the PIC detector. While these PIC detectors may offer improved performance in some cases, they also require $\mathcal{O}(K^2)$ real-real multiplications per iteration and consequently tend to require more computational resources (at least on a per-iteration basis) than the hard-decision PIC detector. This section focuses on the computational complexity of multiuser detectors that do not require any real-real multiplications.

In a given bit interval, the total computational complexity of the hard-decision multistage PIC detector with or without limit cycle mitigation can be expressed as

$$\mathcal{C}_{\mathrm{TOTAL}} = M \cdot \mathcal{C}_{\mathrm{HPIC}} + \mathcal{C}_{\mathrm{LCMx}} \qquad (9)$$

where M is the number of PIC iterations required to reach a convergent state, $\mathcal{C}_{\mathrm{HPIC}}$ is the per-iteration complexity of the hard-decision PIC detector, and $\mathcal{C}_{\mathrm{LCMx}}$ is the computational complexity of the limit cycle mitigation algorithm (equal to zero if no limit cycle mitigation is used).

We note that unlike the majority of multiuser detectors where a deterministic amount of computation is required to compute the bit estimates in each bit interval, the hard-decision multistage PIC detector requires a nondeterministic amount of computation in each bit interval due to the random nature of M and
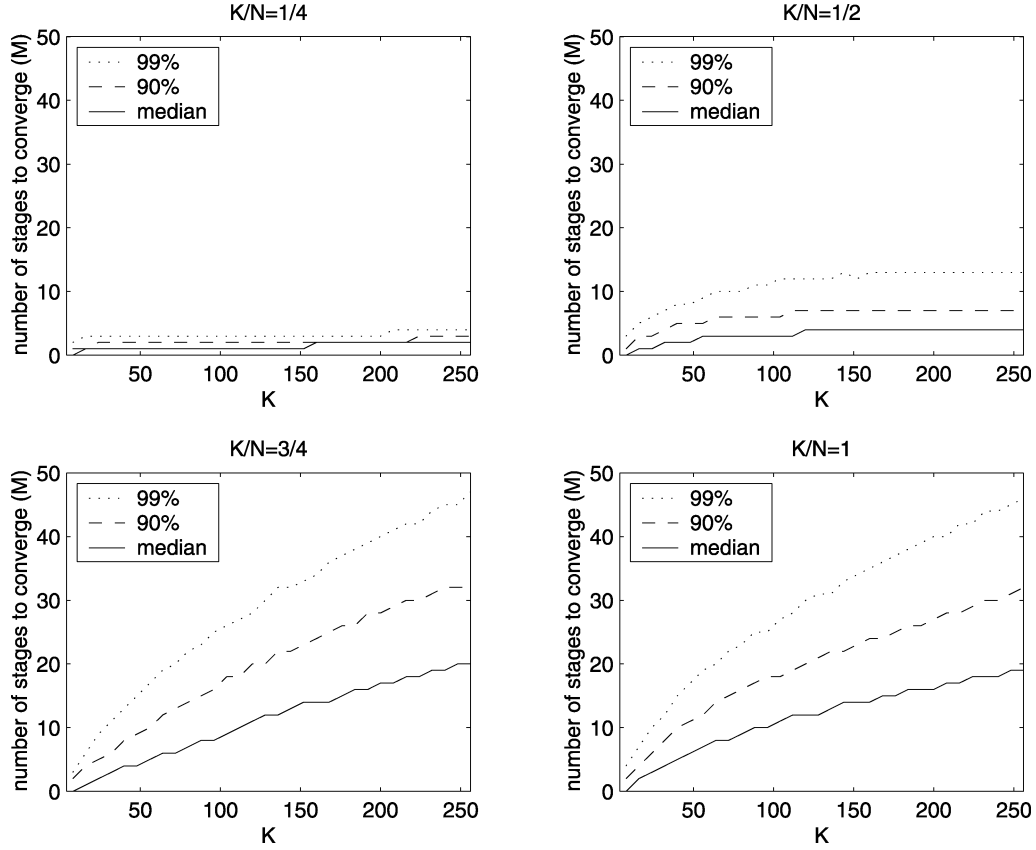
Fig. 7. Number of iterations of hard-decision PIC required to reach an attractor from a matched filter initialization for CDMA systems with length-N random spreading sequences and all users received at 10 dB SNR.

the possibly random nature of $\mathcal{C}_{\mathrm{LCMx}}$ (depending on the limit cycle mitigation algorithm employed). To address this practical implementation challenge, the following sections provide some insight into the variable computational requirements of the hard-decision multistage PIC detector with and without limit cycle mitigation.

1) Computational Complexity of Hard-Decision Multistage PIC: The per-iteration computational complexity of the hard-decision multistage PIC detector is deterministic. Computation of $\boldsymbol{W}\boldsymbol{d}(m)$ in (4) requires $K(K-1)$ real-binary multiplications and $K(K-2)$ signed additions. Subtracting $\boldsymbol{W}\boldsymbol{d}(m)$ from $\boldsymbol{\theta}$ also requires K signed additions, resulting in a total of $K(K-1)$ signed additions. Hence, the per-iteration computational complexity of the hard-decision PIC detector is

$$\mathcal{C}_{\mathrm{HPIC}} = K(K-1)$$

complexity units. We assume that the $\mathrm{sgn}(\cdot)$ operation in (4) and the binary comparisons required to determine if the iteration has converged require no additional computational complexity.

The total computational complexity of the hard-decision multistage PIC detector is nondeterministic due to the fact that the number of iterations of (2) required to reach an attractor, which is denoted M, is random. No closed-form distribution for M is currently known. The neural network literature does offer an upper bound on M in [30], but this upper bound tends to be difficult to compute and quite loose in most cases. To provide some intuition into this quantity, we instead rely on simulations

of a CDMA system with random spreading sequences. Fig. 7 shows the number of iterations required to reach an attractor when the hard-decision multistage PIC detector is initialized with the matched filter decisions. The results show the median number of iterations required to reach an attractor as well as the maximum number of iterations required to reach an attractor in 90% and 99% of the trials. These results suggest that the system loading in terms of K/N has a large impact on the number of iterations required to reach an attractor. These results also suggest that the number of iterations required to reach an attractor, even in the 99% case, is no worse than linear in K and may in fact be sublinear in K in some cases. A proof of this property is an open problem.

2) Computational Complexity of Limit Cycle Mitigation: In this section, we quantify the additional complexity required by the limit cycle mitigation algorithms described in Section IV-A.

LCM1—Full Maximum Likelihood Search: Finding the maximum of $\Lambda(\boldsymbol{d}|\boldsymbol{y})$ over $\boldsymbol{d} \in \{\pm 1\}$ requires $2^K(K^2 + K)$ real-binary multiplications and $2^K K^2$ signed additions under the assumption that each real-valued comparison involved in finding the maximum is computationally equivalent to one signed addition. For reasonably large values of K, we can apply the approximation

$$\mathcal{C}_{\mathrm{LCM1}} \approx 1_{\{\mathrm{LC}\}} 2^K K^2 \qquad (10)$$

complexity units, where $1_{\{\mathrm{LC}\}}$ is the indicator function that is equal to one when the outcome of the hard-decision PIC iteration is a limit cycle and is equal to zero otherwise.
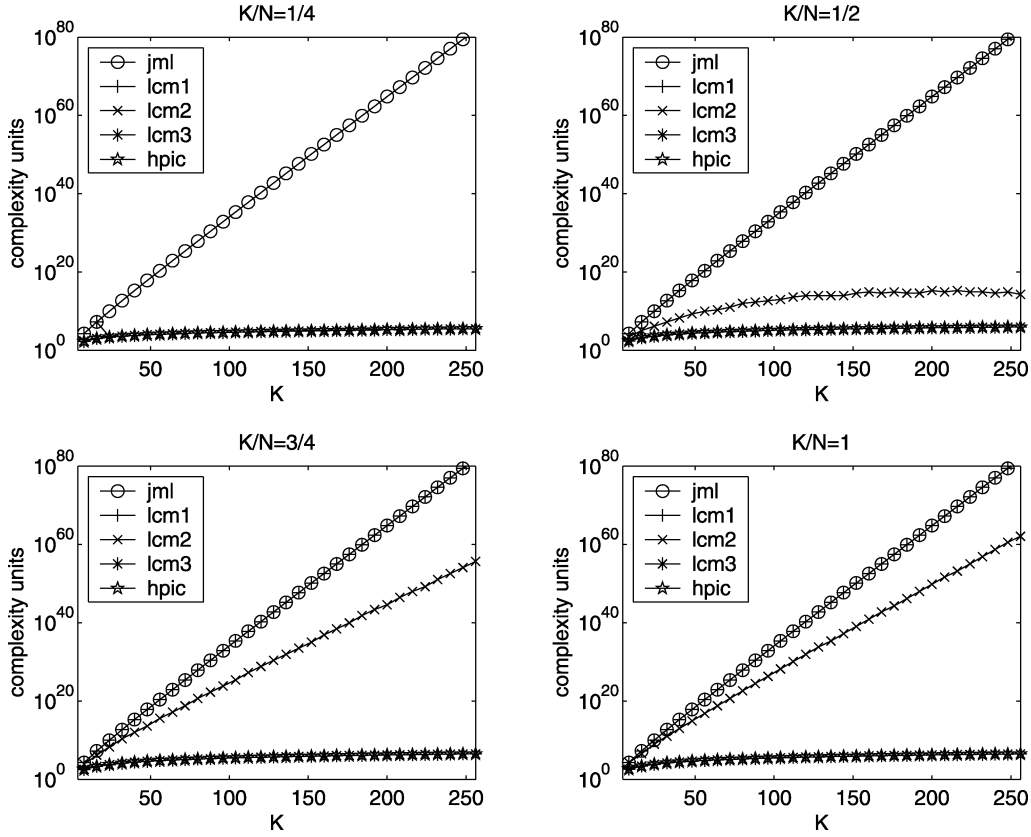
Fig. 8.   Overall complexity of hard-decision multistage PIC detection with and without limit cycle mitigation (99% thresholds) compared with optimum multiuser detection in the case of a CDMA system with random spreading sequences and all users received at 10 dB SNR.

LCM2—Partial Maximum Likelihood Search: Finding the maximum of $\Lambda(\boldsymbol{d}|\boldsymbol{y})$ over $\boldsymbol{d} \in \mathcal{X}$, where $\mathcal{X}$ is defined in (6), requires $2^\kappa(\kappa^2 + \kappa)$ real-binary multiplications and $2^\kappa \kappa^2$ signed additions assuming that the reduction of the problem dimension from K to κ requires negligable computation. When κ is reasonably large, we can apply the approximation

$$\mathcal{C}_{\mathrm{LCM2}} \approx 2^\kappa \kappa^2 \tag{11}$$

complexity units. Note that no indicator function is needed in this case since $\kappa = 0$ unless a limit cycle occurs. In addition, note that κ is a random variable for which no closed-form distribution is currently known. Empirically derived distributions for κ in a CDMA system with random spreading sequences are presented in [37].

LCM3—Soft Output Combining: Assuming that the soft outputs from the past two iterations of the hard-decision PIC iteration are available to the LCM3 algorithm, the LCM3 algorithm requires only K signed additions and no multiplications. Hence

$$\mathcal{C}_{\mathrm{LCM3}} \approx 1_{\{\mathrm{LC}\}} K \tag{12}$$

complexity units.

3) Overall Computational Complexity Comparison: This section combines the results from the prior sections to present an overall computational complexity comparison between the hard-decision multistage PIC detector with and without limit cycle mitigation. We also compare the results to the complexity of the optimum (joint maximum likelihood) multiuser detector

which has a deterministic computational complexity equal to $2^K K^2$ complexity units for reasonably large values of K.

Fig. 8 shows an overall computational complexity comparison of the optimum multiuser detector (jml), hard-decision multistage PIC with limit cycle mitigation (lcm1, lcm2, and lcm3), and hard-decision multistage PIC with no limit cycle mitigation (hpic). The results show the number of complexity units required to compute the final bit estimates in 99% of the trials.

These results show that, when $K/N = 1/4$, the computational complexity of multistage PIC detection with any of the limit cycle mitigation algorithms tends to be very similar to that of the original multistage PIC detector. This is due to the fact that limit cycles occur with very low probability in this case. When $K/N \geq 1/2$, limit cycles now occur with greater probability and the LCM1 algorithm has essentially the same computational complexity as optimum multiuser detection. The LCM2 algorithm has a computational complexity between the optimum detector and the original multistage PIC detector that is largely dependant on the distribution of κ. When $K/N = 1/2$ the computational complexity of the LCM2 algorithm appears to be approximately constant as $K \rightarrow 256$ but, when $K/N \geq 3/4$, computational complexity of the LCM2 algorithm appears to be exponential in K. Finally, the LCM3 detector is indistinguishable from the original hard-decision multistage PIC detector in all cases due to the fact that the overall complexity of this detector is dominated by the $M \cdot \mathcal{C}_{\mathrm{HPIC}}$ term.

As a final remark on the complexity of the LCM2 algorithm, we note that the 99% complexity results shown in Fig. 8 do not reflect the occasional large values of κ that tend to occur rarely (but with nonzero probability) in the cases tested. These outliers from the tail of the κ distribution require a very large number of complexity units in one bit interval and, in practical applications, motivate the use of the hybrid LCM2/LCM3 algorithm described in Section IV-A. The κ^* parameter of the hybrid LCM2/LCM3 algorithm can be specified to provide a 99% computational complexity curve almost anywhere between the LCM2 and LCM3 curves and, perhaps even more importantly, can be specified to provide a strictly upper bounded computational complexity for the limit cycle mitigation algorithm. This feature makes the hybrid LCM2/LCM3 approach attractive for practical applications with limited computational resources.

V. CONCLUSIONS

This paper presents new results on the convergence behavior of the hard-decision multistage PIC detector and a new approach toward improving the performance of this detector. Our results suggest that limit cycles are a significant source of poor performance in the hard-decision multistage PIC detector and we propose a class of limit cycle mitigation algorithms to reactively correct for limit cycle behavior. All of the proposed limit cycle mitigation algorithms retain the desirable properties of the original hard-decision multistage PIC iteration while detecting and correcting for limit cycles only when they occur. Simulation results suggest that limit cycle mitigation can significantly improve the bit error rate performance of the hard-decision multistage PIC detector in a variety of operating scenarios with the greatest improvements observed when the number of users is small with respect to the spreading gain and when the SNR is high.

The proposed limit cycle mitigation algorithms offer a tradeoff between performance gain and increased complexity. The largest performance gains are observed with limit cycle mitigation algorithms that tend to have complexity exponential or near-exponential in K. For practical applications, a hybrid algorithm is proposed that allows the specification of a design parameter to achieve a desired tradeoff in the performance/complexity space.

This paper also highlights the few analytical results that have been published on the dynamics of the hard-decision multistage PIC detector and on nonlinear iterative algorithms in general. There are several relevant open problems in this area, including the development of analytical tools to better understand of the asymptotic behavior seen in Fig. 4, the development of better bounds or distributions on the number of iterations required to reach a convergent state, and the development of a distribution on κ, which is the number of users "participating" in limit cycles.

Potential future research directions include an investigation into the dynamics of the asynchronous hard-decision multistage PIC detector [5] and the development of limit cycle mitigation algorithms for the asynchronous case. Analysis of multistage PIC detection and limit cycle mitigation for the general case of a CDMA system with arbitrary multipath channels that include the effects of intersymbol interference also remains an open problem.

APPENDIX
PROOF OF PROPOSITIONS 1 AND 2

Proof of Proposition 1: Denote $e^{(k)}$ as the kth standard basis vector and observe that the sign of the kth element of d is flipped in $d - 2d^{(k)}e^{(k)}$. By definition, d is a local maximum of the likelihood function iff

$$2d^\top Ay - d^\top ARAd > 2\left(d - 2d^{(k)}e^{(k)}\right)^\top$$
$$\times Ay - \left(d - 2d^{(k)}e^{(k)}\right)^\top ARA\left(d - 2d^{(k)}e^{(k)}\right) \quad (13)$$

holds $\forall k \in \{1, \ldots, K\}$. Expanding and canceling common terms, we can rewrite (13) as

$$0 > -4d^{(k)}e^{(k)\top}Ay + 4d^{(k)}d^\top ARAe^{(k)} - 4\left(a^{(k)}\right)^2 \quad (14)$$

where the last term is due to the fact that $(d^{(k)})^2 = 1$ and $e^{(k)\top}ARAe^{(k)} = (a^{(k)})^2$. Since $Ae^{(k)} = a^{(k)}e^{(k)}$ and $e^{(k)\top}Ay = a^{(k)}y^{(k)}$, we can divide (14) by $-4a^{(k)}$ to get the equivalent expression

$$d^{(k)}y^{(k)} - d^{(k)}d^\top ARe^{(k)} + a^{(k)} > 0. \quad (15)$$

Since $d^\top ARe^{(k)} = \sum_{\ell=1}^{K} d^{(\ell)}a^{(\ell)}\rho_{\ell k}$, we can rewrite (15) as

$$d^{(k)}y^{(k)} - d^{(k)}d^{(k)}a^{(k)}\rho_{kk} - d^{(k)}\sum_{\ell=1,\ell\neq k}^{K} d^{(\ell)}a^{(\ell)}\rho_{\ell k} + a^{(k)} > 0 \quad (16)$$

which, since $d^{(k)}d^{(k)}a^{(k)}\rho_{kk} = a^{(k)}$, simplifies to

$$d^{(k)}\left[y^{(k)} - \sum_{\ell=1,\ell\neq k}^{K} d^{(\ell)}a^{(\ell)}\rho_{\ell k}\right] > 0. \quad (17)$$

Since (17) is equivalent to the original expression (13), we conclude then that d is a local maximum of the likelihood function iff

$$d^{(k)} = \text{sgn}\left(y^{(k)} - \sum_{\ell=1,\ell\neq k}^{K} d^{(\ell)}a^{(\ell)}\rho_{\ell k}\right) \quad (18)$$

holds $\forall k \in \{1, \ldots, K\}$. This last expression can be rewritten in vector notation as

$$d = \text{sgn}\left(y - (R - I)Ad\right)$$

which, by definition, is equivalent to the statement that d is a fixed point of the hard-decision PIC iteration (2). $\qquad\square$

Proof of Proposition 2: Let $\mathcal{H}(x, y)$ denote the Hamming distance between the vectors x and y, both in $\{-1, +1\}^K$. Since $d(m)$ and $d(m-1)$ are both in $\{-1, +1\}^K$ and $d(m) \neq d(m-1)$, then $\mathcal{H}(d(m), d(m-1)) \geq 1$, and it is sufficient to show that $\mathcal{H}(d(m), d(m-1)) \neq 1$ to prove the Proposition.

We will now prove the Proposition by contradiction. Denote $e^{(k)}$ as the kth standard basis vector, and suppose that $\mathcal{H}(d(m), d(m-1)) = 1$. This is equivalent to

$$d^{(k)}(m) = -d^{(k)}(m-1)$$

for one particular $k \in \{1, \ldots, K\}$ and

$$d^{(\ell)}(m) = d^{(\ell)}(m-1)$$

$$\forall \ell \in \{1, \ldots, k-1, k+1, \ldots, K\}. \quad (19)$$

The hard-decision PIC iteration (2) states that

$$d^{(k)}(m) = \text{sgn}\left(y^{(k)} - \sum_{\ell=1,\ell\neq k}^{K} a^{(\ell)}\rho_{k\ell}d^{(\ell)}(m-1)\right) \quad (20)$$

and also that

$$d^{(k)}(m+1) = \text{sgn}\left(y^{(k)} - \sum_{\ell=1,\ell\neq k}^{K} a^{(\ell)}\rho_{k\ell}d^{(\ell)}(m)\right). \quad (21)$$

Since $d^{(\ell)}(m) = d^{(\ell)}(m-1)\forall \ell \in \{1,\ldots,k-1,k+1,\ldots,K\}$, the right-hand sides of (20) and (21) are identical, implying that $d^{(k)}(m) = d^{(k)}(m+1)$. However, by the limit cycle assumption of the Proposition, $d^{(k)}(m+1) = d^{(k)}(m-1)$; hence, $d^{(k)}(m) = d^{(k)}(m-1)$, which is a direct contradiction of (19) since $d^{(k)}(m) \neq 0$. Hence, $\mathcal{H}(\boldsymbol{d}(m), \boldsymbol{d}(m-1)) = 1$ is impossible under the assumptions of the Proposition, and $\mathcal{H}(\boldsymbol{d}(m), \boldsymbol{d}(m-1)) \geq 2$. $\qquad\square$

REFERENCES

[1] S. Verdú, *Multiuser Detection*. New York: Cambridge Univ. Press, 1998.

[2] R. Buehrer, N. Correal, and B. Woerner, "A comparison of multiuser receivers for cellular CDMA," in *Proc. Global Telecomm. Conf.*, vol. 3, London, UK, Nov. 18–22, 1996, pp. 1571–1577.

[3] Y.-F. Huang and P. Diniz, "Interference suppression in CDMA wireless communications," *IEEE Circuits Syst. Soc. Newsletter*, vol. 9, pp. 1, 8–9, 14, Sep. 1998.

[4] T. Ojanpera and R. Prasad, "An overview of air interface multiple access for IMT-2000/UMTS," *IEEE Commun. Mag.*, vol. 36, pp. 82–95, Sep. 1998.

[5] M. Varanasi and B. Aazhang, "Multistage detection in asynchronous code-division multiple-access communications," *IEEE Trans. Commun.*, vol. 38, pp. 509–519, Apr. 1990.

[6] ——, "Near-optimum detection in synchronous code-division multiple-access systems," *IEEE Trans. Commun.*, vol. 39, pp. 725–736, May 1991.

[7] P. Mannion, "Smart basestations maximize capacity," *Commun. Syst. Design*, Jun. 2002.

[8] A. Kaul and B. Woerner, "Analytic limits on the performance of adaptive multistage interference cancellation," *Electron. Lett.*, vol. 30, pp. 2093–2094, Dec. 1994.

[9] V. Ghazi-Moghadam, L. Nelson, and M. Kaveh, "Parallel interference cancellation for CDMA systems," in *Proc. 33rd Annu. Allerton Conf. Commun., Contr, Comput.*, Monticello, IL, Oct. 4–6, 1995, pp. 216–224.

[10] R. Buehrer and B. Woerner, "Analysis of adaptive multistage interference cancellation for CDMA using an improved Gaussian approximation," *IEEE Trans. Commun.*, vol. 44, pp. 1308–1321, Oct. 1996.

[11] N. Correal, R. Buehrer, and B. Woerner, "Improved CDMA performance through bias reduction for parallel interference cancellation," in *Proc. 8th Int. Symp. Personal, Indoor, Mobile Radio Commun.*, vol. 2, Helsinki, Finland, Sep. 1–4, 1997, pp. 565–569.

[12] D. Brown, M. Motani, V. Veeravalli, H. Poor, and C. Johnson Jr., "On the performance of linear parallel interference cancellation," *IEEE Trans. Inf. Theory*, vol. 47, pp. 1957–1970, Jul. 2001.

[13] B. Abrams, A. Zeger, and T. Jones, "Efficiently structured CDMA receiver with near-far immunity," *IEEE Trans. Veh. Technol.*, vol. 44, pp. 1–13, Feb. 1995.

[14] D. Divsalar, M. Simon, and D. Raphaeli, "Improved parallel interference cancellation for CDMA," *IEEE Trans. Commun.*, vol. 46, pp. 258–268, Feb. 1998.

[15] X. Zhang and D. Brady, "Soft-decision multistage detection for asynchronous AWGN channels," in *Proc. the 31st Allerton Conf. Commun., Contr., Comput.*, Monticello, IL, Sep. 1993, pp. 54–63.

[16] ——, "Asymptotic multiuser efficiencies for decision-directed multiuser detection," *IEEE Trans. Inf. Theory*, vol. 44, pp. 502–515, Mar. 1998.

[17] T. Frey and M. Reinhardt, "Signal estimation for interference cancellation and decision feedback equalization," in *Proc. 47th Veh. Technol. Conf.*, Phoenix, AZ, May 4–7, 1997, pp. 155–159.

[18] S. Gollamudi, S. Nagaraj, Y.-F. Huang, and R. Buehrer, "Optimal multistage interference cancellation for CDMA systems using the nonlinear MMSE criterion," in *Conf. Rec. Thirty-Second Asilomar Conf. Signals, Syst., Comput.*, vol. 1, Pacific Grove, CA, Nov. 1–4, 1998, pp. 665–669.

[19] S. Gollamudi and Y.-F. Huang, "Iterative nonlinear MMSE multiuser detection," in *Proc. IEEE ICASSP*, vol. 5, Mar. 15–19, 1999, pp. 2595–2598.

[20] A. Hottinen, H. Holma, and A. Toskala, "Performance of multistage multiuser detection in a fading multipath channel," in *Proc. 6th Int. Symp. Pers., Indoor, Mobile Radio Commun.*, vol. 3, Toronto, ON, Canada, Sep. 27–29, 1995, pp. 960–964.

[21] C. Hegarty and B. Vojcic, "Two-stage multiuser detection for noncoherent CDMA," in *Proc. 33rd Annu. Allerton Conf. Commun., Contr., Comput.*, Monticello, IL, Oct. 4–6, 1995, pp. 1063–1072.

[22] R. Buehrer, "On the convergence of multistage interference cancellation," in *Conf. Rec. Thirty-Third Asilomar Conf. Signals, Syst., Comput.*, vol. 1, Pacific Grove, CA, Oct. 24–27, 1999, pp. 634–638.

[23] M. Varanasi, "Decision feedback multiuser detection: a systematic approach," *IEEE Trans. Inf. Theory*, vol. 45, pp. 219–240, Jan. 1999.

[24] D. Brown and C. Johnson Jr., "SINR, power efficiency, and theoretical system capacity of parallel interference cancellation," *J. Commun. Networks*, vol. 3, pp. 228–237, Sep. 2001.

[25] G. Kechriotis and E. Manolakos, "Hopfield neural network implementation of the optimal CDMA multiuser detector," *IEEE Trans. Neural Networks*, vol. 7, pp. 131–141, Jan. 1996.

[26] J. Hopfield, "Neural networks and physical systems with emerging collective computational abilities," in *Proc. Nat. Acad. Sci.*, vol. 79, 1982, pp. 2554–2558.

[27] K. Smith, M. Palaniswami, and M. Krishnamoorthy, "Neural techniques for combinatorial optimization with applications," *IEEE Trans. Neural Networks*, vol. 9, pp. 1301–1318, Nov. 1998.

[28] S. Verdu, "Minimum probability of error for asynchronous Gaussian multiple-access channels," *IEEE Trans. Inf. Theory*, vol. IT-32, pp. 85–96, Jan. 1986.

[29] J. Hopfield and D. Tank, "Neural computation of decisions in optimization problems," *Biol. Cybern.*, vol. 52, pp. 141–152, 1985.

[30] E. Goles-Chacc, F. Fogelman-Soulie, and D. Pellegrin, "Decreasing energy functions as a tool for studying threshold networks," *Discrete Applied Math.*, vol. 12, pp. 261–277, 1985.

[31] J. Lipscomb, "On the computational complexity of finding a connectionist model's stable state vectors," M.Sc. thesis, Dept. Comput. Sci., Univ. Toronto, Toronto, ON, Canada, 1987.

[32] P. Flor'een and P. Orponen, "On the computational complexity of analyzing Hopfield nets," *Complex Syst.*, vol. 3, no. 6, pp. 577–587, 1989.

[33] M. Garey and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: W.H. Freeman, 1979.

[34] P. Flor'een and P. Orponen, "Attraction radii in binary Hopfield nets are hard to compute," *Neural Comput.*, vol. 5, pp. 812–821, 1993.

[35] F. Barahona, "On the computational complexity of Ising spin glass models," *J. Phys. A*, vol. 15, pp. 3241–3253, 1982.

[36] M. Varanasi, "Group detection for synchronous Gaussian code-division multiple-access channels," *IEEE Trans. Inf. Theory*, vol. 41, pp. 1083–1096, Jul. 1995.

[37] D. Brown, "Improved multistage parallel interference cancellation using limit cycle mitigation," in *Proc. Conf. Inform. Sci. Syst.*, Princeton, NJ, Mar. 20–22, 2002.

D. Richard Brown, III (S'97–M'00) received the B.S. and M.S. degrees in electrical engineering from the University of Connecticut, Storrs, in 1992 and 1996, respectively, and the Ph.D. degree in electrical engineering with a minor in mathematics from Cornell University, Ithaca, NY, in 2000.

From 1992 to 1997, he was with the General Electric Company, Plainville, CT, as a Development Engineer. He is currently an assistant professor with the Department of Electrical and Computer Engineering, Worcester Polytechnic Institute, Worcester, MA. His research interests include adaptive signal processing, multiuser communication systems, and interference cancellation.