Quiver Diagrams and Signed Adaptive Filters

CRAIG R. ELEVITCH, WILLIAM A. SETHARES, MEMBER, IEEE, GONZALO J. REY, AND C. RICHARD JOHNSON, JR., SENIOR MEMBER, IEEE

Abstract—**Little is known about the convergence properties of the sign-sign variant of the well-known least mean square (LMS) adaptive algorithm. Extrapolation from the behavior of the unsigned LMS adaptive algorithms to the behavior of the signed versions is substantially misleading. While the parameter estimates of the FIR adaptive LMS filter are always bounded, the parameter estimates of the sign-sign variant may diverge. In cases where the correct parameterization of a properly chosen IIR adaptive LMS filter is locally stable, the correct parameterization of the IIR sign-sign variant may be locally *unstable*. Unlike unsigned IIR LMS, the sign-sign variant may have multiple stable limit cycles.**

This paper examines a graphical technique with which to explore the behavior of deterministic discrete time adaptive algorithms excited by periodic inputs. The constraint to periodic inputs is convenient; the underlying concepts used also apply to stationary inputs drawn from finite alphabets. The resulting graph is based on plotting single period parameter trajectories in the parameter error space and is essentially a discrete version of flow diagrams associated with continuous time systems. For the LMS adaptive FIR filter with the algorithm step size sufficiently small, the motion of adapted parameters from a particular initial value is given by the total update over the input period with the parameterization frozen at its "initial value." This results in the familiar elliptical hyperparaboloid surface steepest-descent interpretation. This method can be extended to the LMS adaptive IIR filter with suitable approximations based on a small step size. A less familiar interpretation, based on the parameter estimates following a possibly nonsteepest-descent, pseudogradient trajectory, can be applied to the resulting picture. For variants of LMS with signed error, the novel feature is that single period update trajectories indicating the directions of these averaged movements can be bundled into "quiver diagrams" which accurately predict parameter trajectories, the presence of stable limit cycles, as well as stable and unstable stationary loci, and offer insights into the influence of the character of the input on signed error algorithm stability.

"I shot an arrow into the air.
It fell to earth, I knew not where."

—H. W. Longfellow

I. Introduction

In a number of high speed communications applications where computational resources are at a premium, the computational requirements of even the simple, discrete-time LMS adaptive filter [1] may be excessive. Replacing regressor and prediction error components of the update term by their signs reduces computing time and dynamic range requirements by turning multiplications into bit shifts. Such signed versions of LMS adaptive FIR and IIR filters are called "sign-sign" algorithms. An early suggestion of the applicability of the sign-sign LMS algorithm appears in [2] for use in channel equalization. More recently, an international standard for 32 kbit/s adaptive differential pulse code modulation (ADPCM) in telephony [3] has been adopted that employs the sign-sign variant in an adaptive IIR filtering task [4].

The stability and convergence properties of the sign-sign FIR and IIR adaptive filters remain an open issue, and the relatively well-developed understanding of unsigned adaptive filter algorithms gives meager intuition as to the behavior of their signed cousins. For instance, the persistency of excitation (PE) condition for unsigned LMS, which is a condition requiring sufficient spectral richness in the input, guarantees convergence of the parameter estimates to the unique setting that zeros the parameter error in ideal use [5]. As shown in [6], this PE condition does not guarantee parameter convergence with the use of sign-sign LMS. In fact, some inputs that are PE for unsigned LMS can lead to parameter divergence for the sign-sign version of the same LMS algorithm. A generic description of all inputs which guarantee parameter convergence at least to within a small ball about the correct answer in ideal use of sign-sign LMS (which is one feature of a PE condition) is as yet unknown.

This paper presents "arrow diagrams," which are essentially a discrete time version of continuous time flow diagrams. As we will see, arrow diagrams complement the frequently exploited gradient descent interpretation of the parameter space behavior of adaptive filters, such as LMS. The geometry of these diagrams depends on the adaptive filter input and, thus, is strongly coupled to the PE condition for LMS. A certain class of arrow diagrams called "quiver diagrams" arises when the prediction error part of the usual LMS update is replaced by its sign. Quiver diagrams help clarify complicated parameter estimate behavior by graphically displaying the collection of short-term average trajectory motions for the sign-sign algorithms. Even though quiver diagrams will be derived quantitatively, their real use is qualitative. When an input generates a quiver diagram for which any starting point (within a particular region) "flows" to the vicinity of a setting that zeros the parameter error in ideal use, this input is considered a candidate "PE" signal for the sign-sign FIR algorithm. Sufficient diagram properties for such parameter convergence can be translated into an algebraic definition of classes of associated persistently exciting in-

Manuscript received August 8, 1987; revised May 25, 1988. This work was supported by NSF under Grant MIP—8608787.

C. R. Elevitch, G. J. Rey, and C. R. Johnson, Jr., are with the School of Electrical Engineering, Cornell University, Ithaca, NY 14853.

W. A. Sethares was with the School of Electrical Engineering, Cornell University, Ithaca, NY. He is now with the Department of Electrical and Computer Engineering, University of Wisconsin, Madison, WI 53706.

IEEE Log Number 8825132.

puts. This graphical analysis of the sign-sign LMS adaptive FIR filter is extended to the IIR case, but the transfer to a definition of persistently exciting signals associated with the latter is substantially more complicated, since the parameter estimate update is nonlinear and the "PE" condition may only be a local condition.

Another distinction between the sign-sign algorithms and their unsigned counterparts concerns the relation between the FIR and IIR cases. Inputs which are persistently exciting for unsigned adaptive FIR filters are also persistently exciting for unsigned adaptive IIR filters, if a particular transfer function is strictly positive real (SPR) [7]. Simulations show that inputs which are persistently exciting for the sign-sign adaptive FIR filter (in terms of the quiver diagrams and confirmed by simulations) are not necessarily persistently exciting for the sign-sign IIR filter, even if the same SPR condition is fulfilled.

The following section introduces quiver diagrams by first examining single period update trajectories in the parameter space for a given periodic input to unsigned LMS, and then extends the concept to three signed variants of LMS: signed regressor, signed error, and sign-sign. Section III will construct a quiver diagram for the 2-parameter sign-sign LMS adaptive filter, and present simulations that display how the diagram corresponds to "usual" algorithm behavior. Brief comments will be provided as to how the simple geometrical insights given by quiver diagrams help define algebraic conditions on the inputs which guarantee stability of sign-sign LMS in the FIR case. Section IV extends the quiver diagram concept to the sign-sign adaptive IIR filter, demonstrating the rather surprising (and unwelcome) potential for local instability about the desired parameterization and the existence of multiple stable limit cycles.

II. Introduction of the Quiver Diagram Concept

Consider the mth-order single-input single-output FIR plant in the familiar form

$$y_k = \sum_{j=1}^{m} b_j u_{k-j} = X_{k-1}^T \theta \tag{2.1}$$

where the b_j's are plant parameters, u_k's are inputs to the plant, and the regressor vector X_{k-1} and plant parameter vector θ are defined as

$$X_{k-1} = [u_{k-1}, u_{k-2}, \cdots, u_{k-m}]^T,$$
$$\theta = [b_1, b_2, \cdots, b_m]^T. \tag{2.2}$$

The plant output estimate is then

$$\hat{y}_k = \sum_{j=1}^{m} \hat{b}_{j_{k-1}} u_{k-j} = X_{k-1}^T \hat{\theta}_{k-1}, \tag{2.3}$$

where

$$\hat{\theta}_k = [\hat{b}_{1_k}, \hat{b}_{2_k}, \cdots, \hat{b}_{m_k}]^T. \tag{2.4}$$

The standard least mean squares (LMS) parameter update [1] is

$$\hat{\theta}_k = \hat{\theta}_{k-1} + \mu X_{k-1} e_k, \tag{LMS}$$

where μ is a small positive step size and $e_k = y_k - \hat{y}_k$ is the prediction error. The prediction error can be rewritten as

$$e_k = y_k - \hat{y}_k = X_{k-1}^T \theta - X_{k-1}^T \hat{\theta}_{k-1}$$
$$= X_{k-1}^T (\theta - \hat{\theta}_{k-1}) = X_{k-1}^T \tilde{\theta}_{k-1}, \tag{2.5}$$

where $\tilde{\theta}_k \equiv \theta - \hat{\theta}_k$ is the parameter error vector. Implicit in this expression for e_k is exact plant and model matching, noise free measurements, etc.; in short, this paper considers LMS only in the ideal case. Subtracting both sides of (LMS) from θ, the parameter error system for (LMS) can be written

$$\tilde{\theta}_k = \tilde{\theta}_{k-1} - \mu X_{k-1} (X_{k-1}^T \tilde{\theta}_{k-1}). \tag{2.6}$$

It will often be convenient to refer to this latter form of (LMS).

A simplistic interpretation of (LMS) shows that the single-step update of the $\hat{\theta}$ vector is parallel to X_{k-1}, with direction (plus or minus) and scaling given by the prediction error e_k. Since X contains only delayed inputs, each X provides a fixed (i.e., independent of the parameter estimate) one-dimensional space in which to move, and the error provides a direction and magnitude to that motion. Equation (2.5) shows that e_k changes sign along the hyperplane defined by the nullspace of X_{k-1}, that is, about any $\tilde{\theta}_{k-1}$ for which $X_{k-1}^T \tilde{\theta}_{k-1} = 0$. Notice that this hyperplane passes through $\tilde{\theta} = 0$ (i.e., the correct setting) for all X_{k-1}, under ideal conditions. It is easy to see that for each regressor X_{k-1}, motion of the parameter estimates is *always toward the nullspace of* X_{k-1}.

Parameter motion for the single step update $\Delta\tilde{\theta}_k = \tilde{\theta}_k - \tilde{\theta}_{k-1}$ corresponding to a 2-dimensional X_{k-1} is depicted in Fig. 1(a). It is evident from Fig. 1(a) (and this can be shown formally as in [8]) that under ideal conditions parameter motion given by (LMS) will never increase the summed squared parameter error, assuming a small enough step size. In fact, one can visualize from Fig. 1(a) that if the regressors regularly span the parameter space $\mathbb{R}^2$, the algorithm will "push" the parameter estimates toward the nullspace common to all regressors, eventually settling on the correct setting.

Before turning to single-step motion in the sign-sign LMS case [Fig. 1(b)], consider the τ-term trajectory of (LMS) from an initial parameter estimate setting $\tilde{\theta}_i$ calculated as

$$\Delta_\tau = \tilde{\theta}_{i+\tau} - \tilde{\theta}_i = -\mu \sum_{k=i}^{i+\tau-1} X_{k-1} e_k. \tag{2.7}$$

If τ is the period of the input, the average update can be calculated simply by using a finite window (of length τ). The update represented by (2.7) over a period τ will be called the "periodic update." Note that to define the average over a single-step scaling of Δ_τ by $1/\tau$ would be necessary, but has been omitted for convenience. For each initial $\tilde{\theta}_i$ in (2.7), a unique final $\tilde{\theta}_{i+\tau}$ can be calculated. Connecting the initial and final points yields a vector (Δ_τ) which represents the motion of the parameter estimate

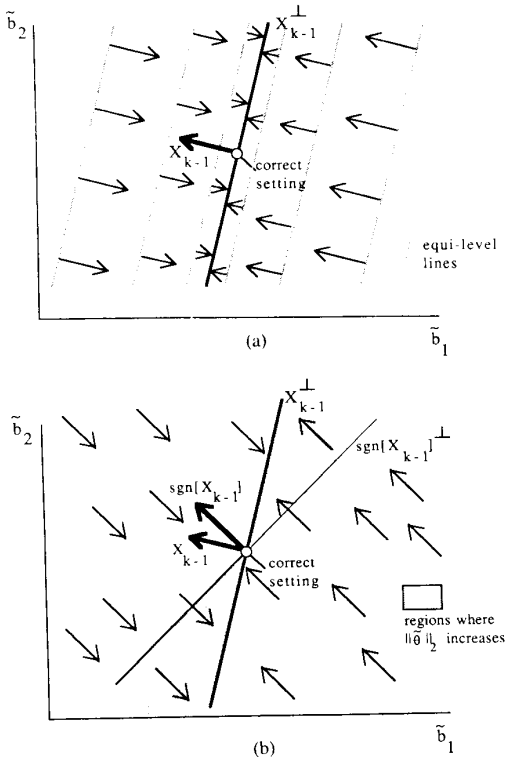

Fig. 1. (a) Parameter trajectories for the one step update of (LMS). All trajectories are parallel to X_{k-1}. (b) Parameter trajectories for sign-sign LMS. Updating occurs parallel to sgn $[X_{k-1}]$, leading to regions where $\|\tilde{\theta}\|_2$ increases.

over the τ samples of the input. Plotted in the parameter error space, these "trajectory snapshots" or "arrows" give rise to an "arrow diagram." The arrow diagram of Fig. 2 for a grid of $\tilde{\theta}_i$ candidates in the parameter error space is based on the 2-dimensional plant and model

$$y_k = b_1 u_{k-1} + b_2 u_{k-2}$$
$$\hat{y}_k = \hat{b}_{1_{k-1}} u_{k-1} + \hat{b}_{2_{k-1}} u_{k-2} \tag{2.8}$$

with 3-periodic input $u_k = \{\cdots, -1.5, -1, 3, \cdots\}$, and plant parameters $(b_1, b_2) = (0.6, 1.0)$. The arrow diagram of the kind shown in Fig. 2 may be thought of as a 2-dimensional projection onto the parameter error space of (closely approximate) instantaneous downward gradient lines of the 3-dimensional error surface $J(e) = \text{avg}[e_k^2]$. The error surface given by J can be described as an elliptical hyperparabolic "bowl" with a unique minimum at the correct setting. For any 2-dimensional FIR plant in the form of (2.8) and any τ-periodic input sequence ($\tau \geq 2$), the equilevel curves of J are ellipses with axes oriented along the 45° lines. Fig. 2 reflects this elliptical bowl shape: connecting the starting points of average trajectories of similar length will sketch in the level curve ellipses, and the arrows become shorter as the bowl "flattens out" near $\tilde{\theta} = 0$. The above corresponds to the standard steepest-descent interpretation of the unsigned LMS algorithm.

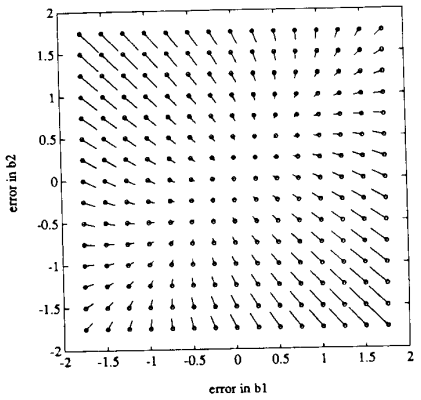

Fig. 2. Arrow diagram for unsigned LMS. Each arrow represents the trajectory for one period of the input $u_k = \{\cdots, -1.5, -1, 3, \cdots\}$ with initial condition at the base of the arrow, which is marked by an ○.

In developing some insight into the sign-sign LMS adaptive filter, it is useful to compare (LMS) to three signed versions of (LMS): signed regressor LMS, in which the regressor alone is replaced by its sign; signed error LMS, in which the prediction error alone is replaced by its sign; and the sign-sign LMS, in which both the regressor and prediction error are replaced by their signs. These three algorithms replace the (LMS) update of $\hat{\theta}_{k-1}$ with

$$\hat{\theta}_k = \hat{\theta}_{k-1} + \mu \text{ sgn}[X_{k-1}] e_k, \tag{SR}$$
$$\hat{\theta}_k = \hat{\theta}_{k-1} + \mu X_{k-1} \text{ sgn}[e_k], \tag{SE}$$

and

$$\hat{\theta}_k = \hat{\theta}_{k-1} + \mu \text{ sgn}[X_{k-1}] \text{ sgn}[e_k], \tag{SS}$$

where "sgn $[\cdot]$" is the usual signum function to be interpreted as an element by element operator when applied to a vector. Although the geometrical descriptions of (LMS) of Figs. 1(a) and 2 are not new, they lend a familiar backdrop for similar graphical analysis of signed variants of LMS. The first question one might ask is: how do sign operators alter the single-step update of LMS? For both the signed regressor and sign-sign LMS variants, the parameter update is parallel to sgn $[X_{k-1}]$, which may be significantly out of alignment with X_{k-1}. The prediction error for both of these algorithms serves the same purpose as in (LMS); it directs the update to move with or against sgn $[X_{k-1}]$. That is, the nullspace of X_{k-1} given by $e_k = X_{k-1}^T \tilde{\theta}_{k-1} = 0$ divides the space into $\pm$[signed regressor direction] motion. Fig. 1(b) depicts motion for a single update in the parameter estimate space for (SS). Due to the sign operators in (SS), *all updates for a single regressor are of equal magnitude, colinear and opposite in direction in the two half-spaces defined by the boundary.*

As noted above, each update for unsigned LMS moves in the direction of X_{k-1} toward the boundary given by the nullspace of X_{k-1}. Since sgn $[X]$ always has a positive projection of X (sgn $[X]^T X = \|X\|_1 \geq 0$, $X \in \mathbb{R}^m$), the component of the update direction for (SS) perpendicular

to the boundary is pointed toward the boundary. However, unlike (LMS) or (SE), motion is not necessarily orthogonal to the nullspace of X (since, in general, $X \neq \alpha$ sgn $[X]$, where $\alpha \in \mathbb{R}$). As with (SR) relative to unsigned LMS [9], frequent misalignment of sgn $[X]$ and X can conceivably cause poor behavior in sign-sign LMS, versus the behavior of (SE) as analyzed in [10]. For both (SR) and (SS), there exist regions in the parameter estimate space for which parameter motion *increases* the summed squared parameter error [see region shaded in Fig. 1(b)]. An expression for the change in the squared parameter error shows this:

$$\tilde{\theta}_{i+1}^T \tilde{\theta}_{i+1} - \tilde{\theta}_i^T \tilde{\theta}_i = \left(\tilde{\theta}_i - \mu \operatorname{sgn}[X_i] X_i^T \tilde{\theta}_i\right)^2 - \tilde{\theta}_i^T \tilde{\theta}_i$$
$$= -2\mu \tilde{\theta}_i^T \operatorname{sgn}[X_i] X_i^T \tilde{\theta}_i + o(\mu^2). \tag{2.9}$$

When sgn $[\tilde{\theta}^T \operatorname{sgn}[X_i]] = -\operatorname{sgn}[X_i^T \tilde{\theta}_i]$, i.e., when $\tilde{\theta}_i$ lies in the shaded region of Fig. 1(b) and μ is small, (2.9) shows that the change in the squared parameter error is positive.

The signed error algorithm update moves parallel to X_{k-1}, but is scaled only by the sign of the prediction error, i.e., ± 1. Here again the decision to move with or against X_{k-1} is delineated by the nullspace of the regressor. As seen in Fig. 1(b), all update arrows have the same magnitude. In common to all of these LMS variants is that the nullspace of each X_{k-1} sets up a boundary along $X_{k-1}^T \tilde{\theta}_{k-1} = 0$ across which the direction of parameter motion changes sign. These boundaries prove to be significant in dissecting the behavior of signed algorithms. Arrow diagrams for each of the three signed LMS algorithms are presented in Fig. 3, for the plant given by (2.8) and the same input as used for Fig. 2. The nullspace of each regressor has been superimposed on these diagrams.

Notice that within certain regions for both (SE) and (SS), all of the arrows are parallel and of the same magnitude; these conic regions delineated by regressor nullspaces will be termed "sectors." The periodic update term from an initial setting $\tilde{\theta}_i$ explains this behavior as we now show. An expression for the periodic update for (SS) which is analogous to that of (2.7) is

$$\Delta_\tau = \tilde{\theta}_{i+\tau} - \tilde{\theta}_i$$
$$= -\mu \sum_{k=i+1}^{i+\tau} \operatorname{sgn}[X_{k-1}] \operatorname{sgn}[X_{k-1}^T \tilde{\theta}_{k-1}]. \tag{2.10}$$

When no boundaries are crossed in τ steps, sgn $[X_{k-1}^T \tilde{\theta}_i]$ = sgn $[X_{k-1}^T \tilde{\theta}_j]$ where $\tilde{\theta}_i$ and $\tilde{\theta}_j$ are any two points in the same sector. Thus, (2.10) can be written as

$$\Delta_\tau = -\mu \sum_{k=i+1}^{i+\tau} \operatorname{sgn}[X_{k-1}] \operatorname{sgn}[X_{k-1}^T \tilde{\theta}_i], \tag{2.11}$$

as if each were calculated with *fixed* $\tilde{\theta}_i$ for motion not crossing boundaries. The reasoning is identical for (SE). Without the sign operator on the error, rewriting (2.10) as (2.11) would not be possible. In other words, it is the sign operator on the error which makes each update sensitive only to the side of the boundary on which $\tilde{\theta}_{k-1}$ lies, and insensitive to the distance to the boundary. With X periodic, the sum of individual updates is the same over each period of X when no boundary is crossed in that period, explaining the bundling of arrows in the arrow diagrams for (SE) and (SS). By (2.11), a single periodic update arrow of a sector *exactly* describes the motion of the parameter estimates from any initial condition in the sec-

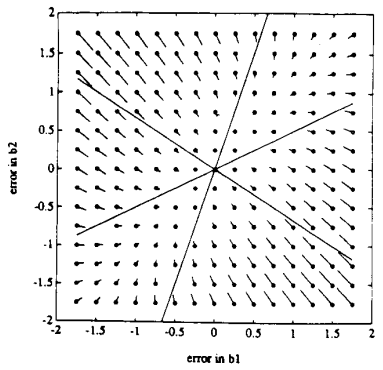

(SR)

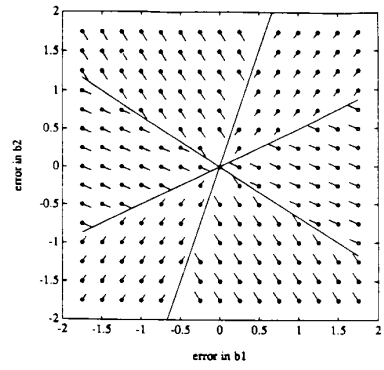

(SE)

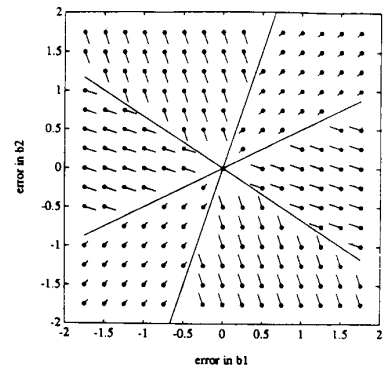

(SS)

Fig. 3. Arrow diagrams for (SR), (SE), and (SS) which are companions to the arrow diagram for (LMS) of Fig. 2. The input is again $u_k = \{\cdots, -1.5, -1, 3, \cdots\}$. For both (SS) and (SE), the periodic trajectory arrows are identical in large open conic regions of the space. Since arrows in each region can be bundled into a single representative arrow for the region, these diagrams are termed "quiver diagrams."

tor except when a boundary is crossed. The arrow can be said to reflect periodic behavior within a sector excluding a "β-region" about the boundary, where the Euclidean distance β is the maximum possible magnitude of the periodic update trajectory ($\beta = \mu\tau m^{1/2}$). Since all of the arrows in each sector are identical, they may be bundled into "quivers," creating the quiver diagram. From a computational standpoint, only one arrow need be calculated per sector.

It should be pointed out that assuming a deterministic τ-periodic input in (2.11) serves only as an instrument in forming average trajectories within sectors. Any input whose elements are restricted to a finite alphabet yields a finite number of regressors and thus boundaries. Within sectors delineated by these boundaries, an average trajectory (also implicit in periodic trajectories) somehow must be determined. A variety of averaging techniques are presented in [11] and [12]. For an algorithm using the correlation update form of LMS, the average trajectories will be the same throughout each sector if a function is applied to the prediction error which is constant throughout each sector, such as the sign operator. Both the signed error and sign-sign LMS algorithms give rise to quiver diagrams. When the sign operator is applied to the error term in LMS, the boundaries are the nullspaces of the regressors, and in each sector the average trajectory (or the sum of the individual regressor updates as with periodic inputs) is the same from all initial conditions, excluding the β-regions about the boundaries. A similarly bundled geometry also arises with piecewise constant functions of the prediction error, such as multilevel quantization. These and other potential extensions/applications of quiver diagrams are mentioned in Section V of this paper. The use of deterministic τ-periodic inputs in this paper simplifies the discussion, but does not restrict the qualitative value of quiver diagrams. To be sure, quiver diagram insights are considerably diminished if too many boundaries are present. In the following section, the building of a quiver diagram for the sign-sign LMS adaptive FIR filter using τ-periodic inputs will be detailed.

III. Building a Quiver Diagram for Sign-Sign LMS

The periodic trajectory when no boundary crossings are involved as given by (2.11) implies that the ordering of the regressors has no bearing on the periodic updates represented by quiver arrows (although along the boundaries regressor order plays a nontrivial role). For the purposes of constructing a quiver diagram then, instead of considering the ordered periodic input stream $u_k = \{\cdots, u_1, u_2, u_3, \cdots\}$, we may consider the input to be the unordered *regressor set* $\mathfrak{X} = \{X_1, X_2, X_3, \cdots\}$.

For the 2-parameter FIR plant of (2.8) with the 3-periodic input used in the example $u_k = \{\cdots, -1.5, -1, 3, \cdots\}$, the set of regressors is

$$\mathfrak{X} = \left\{\begin{bmatrix} -1 \\ -1.5 \end{bmatrix}, \begin{bmatrix} -1.5 \\ 3 \end{bmatrix}, \begin{bmatrix} 3 \\ -1 \end{bmatrix}\right\}. \tag{3.1}$$

Each of these defines a boundary in the parameter error space, and via the prediction error, the sign of the update direction. One way to draw a quiver diagram would be to construct the boundaries for all regressors, choose a point within each sector, and calculate the periodic update arrow according to (2.11) using a fixed $\tilde{\theta}_i$. Equivalently, one may construct the boundaries and updates for each regressor, and then "superimpose" these over one another, summing the arrows graphically in each sector. This summing of the unordered regressor effects to obtain sector arrows is seen implicitly in (2.11): from a fixed $\tilde{\theta}_i$ (within a sector but excluding the β-region), terms in the last sum may be arranged in any order. The signed regressors corresponding to $\mathfrak{X}$ are

$$\mathfrak{X}_{sgn} = \left\{\begin{bmatrix} -1 \\ -1 \end{bmatrix}, \begin{bmatrix} -1 \\ -1 \end{bmatrix}, \begin{bmatrix} 1 \\ -1 \end{bmatrix}\right\}. \tag{3.2}$$

For each regressor in $\mathfrak{X}$, its nullspace and the updates of $\mathfrak{X}_{sgn}$ are plotted in the parameter estimate space in Fig. 4(a)–(c). The "sum" of these three figures is depicted in Fig. 4(d). Fig. 4(d) is the quiver diagram.

The arrows in the quiver diagram indicate the direction and magnitude of motion over one period, and it is straightforward to "predict" the behavior from any initial condition. From points well interior to a sector, the periodic trajectory will have a direction and relative magnitude as shown by the arrows. Note that the arrows have been scaled to make them clearly visible. In general, it seems logical to suggest that upon reaching a boundary, the estimate trajectory will either cross the boundary or "chatter" along the boundary. Chattering seems likely if the arrows on either side of these boundaries share a radial component toward the desired setting, and have opposing projection on the perpendicular to the boundary. Simulation evidence suggests that behavior along the boundaries is as expected based on the direction of sector arrows: parameter estimates chatter along the boundary in the direction common to the arrows when both sector arrows point toward it; estimates pass through boundaries where arrows have the same directional component orthogonal to the boundary. It can be proved that boundary behavior is as expected [13] assuming only one boundary is crossed in the input period, which is always the case in two dimensions except for within a region about the origin. Thus, by the quiver diagram of Fig. 4, the algorithm will evidently take all initial conditions to a small ball about the correct setting, since the radial component of motion in each sector points towards this setting. We conclude that $u_k = \{\cdots, -1.5, -1, 3, \cdots\}$ may be called persistently exciting for 2-dimensional (SS).

The concept of quiver diagram sector arrows may be used to define a class of PE inputs for (SS) in the FIR case. Imagine such a condition based on testing that all sector arrows from initial condition on a ball about the correct setting have a radial component towards the correct setting. The arrows of Fig. 4, for example, satisfy such a condition on a sufficiently large ball about the or-

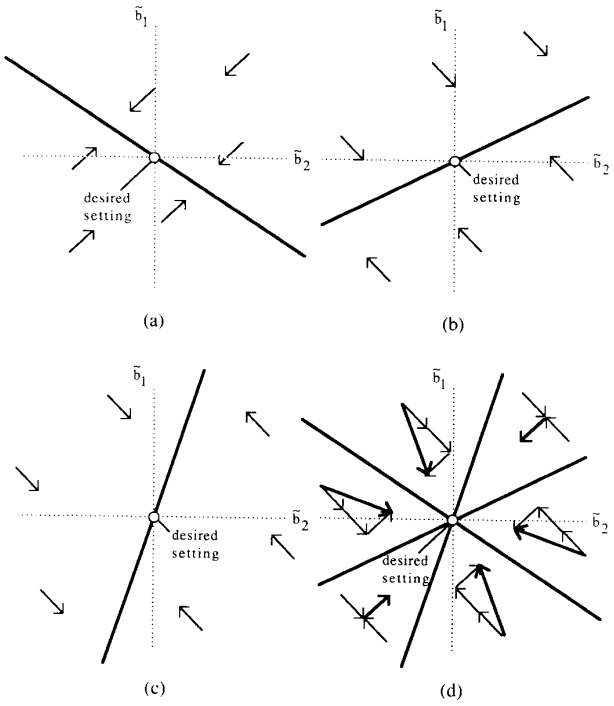

Fig. 4. (a)-(c) Boundaries corresponding to the regressors in $\mathfrak{X}$ and corresponding update directions given by $\mathfrak{X}_{sgn}$. (d) Graphical sum of (a)-(c) with update arrows laid end to end, and periodic update trajectories marked in bold for each sector. This is the quiver diagram for the input $u_k = \{\cdots, -1.5, -1, 3, \cdots\}$ is the FIR case.

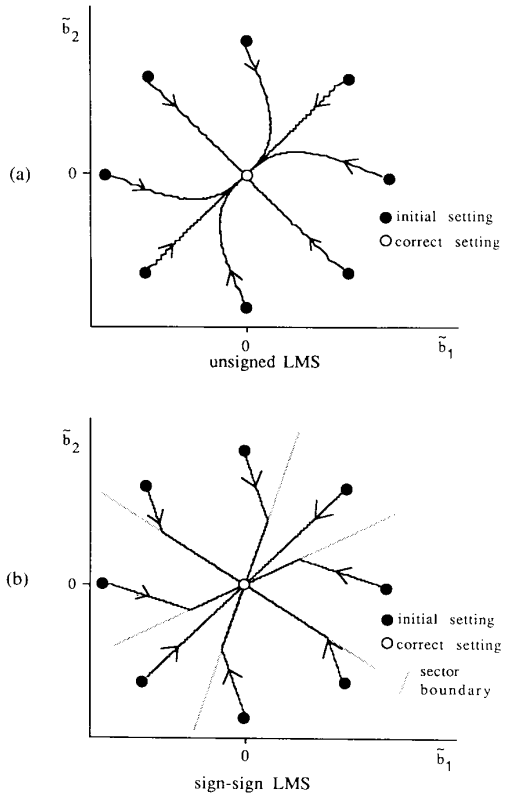

Fig. 5. Parameter error space trajectories for the unsigned and sign-sign LMS algorithm for several different initial conditions. In (b) the shaded lines are the quiver diagram boundaries. The input is $u_k = \{\cdots, -1.5, -1, 3, \cdots\}$.

igin. By the convexity of sectors (since sectors are the intersection of half-spaces), if the sector arrow points outward from the origin anywhere in the sector, it will point outward at some sector boundary. Therefore, a contraction condition need be checked only for initial conditions lying on the sector boundaries. Such a test is simply carried out by determining whether the projections of the periodic trajectory onto the sector boundaries are inward toward the origin. One can see from Fig. 4 that the projections of each arrow onto its sector's boundaries point radially inward. Note that this condition is a sufficient condition, since a trajectory which has a radially outward component within one sector may point into a sector whose trajectory points radially inward and assures convergence. In higher dimensions, it can be shown that sectors can be defined by a convex combination of their "edges," or one-dimensional boundaries [13]. Consider an algebraic formulation of a test that all sector arrows have a sufficient radial component that is based on projecting sector arrows onto these one-dimensional boundaries. One is presented in [13], which results in a sufficient condition guaranteeing stability, in the sense that all possible periodic trajectories from a large enough ball about the correct setting have an inward radial component.

Simulations confirm behavior based on the arrow diagram (Fig. 2) for unsigned LMS, and on the quiver diagram [Fig. 4(d)] for sign-sign LMS. Plotting parameter trajectories in the parameter space or phase plane, rather than as time functions, offers the opportunity to plot several trajectories from different initial conditions in a single diagram; several such trajectories are depicted in Fig. 5. The parameter trajectories in Fig. 5 for unsigned LMS with a small step size are almost smooth curves approaching the correct parameter values, $\hat{\theta} = \theta$. In the sign-sign LMS case, the trajectories are not smooth, but appear to be roughly piecewise linear. The average motion of the parameters is linear for a time, then "turns sharply" to a trajectory which "slides" along the boundary, as expected from the direction of quiver diagram arrows. The smooth legs of the "spider" diagram of Fig. 4(a) reveal the average error surface descended by (LMS), while the angular legs of Fig. 4(b) are the parameter motions for (SS).

IV. Extension of Quiver Diagrams to the Sign-Sign LMS Adaptive IIR Filter

Consider the IIR plant model having the form

$$y_k = \sum_{i=1}^{n} a_i y_{k-i} + \sum_{j=1}^{m} b_j u_{k-j}. \tag{4.1}$$

Replacing the a_i's and b_j's by their filter estimates at time $k-1$, $\hat{a}_{i_{k-1}}$ and $\hat{b}_{j_{k-1}}$, the output estimate $\hat{y}_k$ is formed as

$$\hat{y}_k = \sum_{i=1}^{n} \hat{a}_{i_{k-1}} \hat{y}_{k-i} + \sum_{j=1}^{m} \hat{b}_{j_{k-1}} u_{k-j}. \tag{4.2}$$

Since delayed output estimates appear on the right-hand side of (4.2), this model has an infinite response (IIR), and the associated regressor contains past inputs as well as past output estimates

$$X_{k-1} = [\hat{y}_{k-1}, \cdots, \hat{y}_{k-n}, u_{k-1}, \cdots, u_{k-m}]^T. \tag{4.3}$$

The appropriate definition of the parameter estimate vector is

$$\hat{\theta}_k = [\hat{a}_{1k}, \cdots, \hat{a}_{nk}, \hat{b}_{1k}, \cdots, \hat{b}_{mk}]^T, \tag{4.4}$$

which gives the output estimate $\hat{y}_k = X_{k-1}^T \hat{\theta}_{k-1}$. The prediction error can be shown to be a function of the inner product of the regressor and parameter error vector filtered through the unknown autoregression of the plant [8], as

$$e_k = y_k - \hat{y}_k = \frac{1}{1 - \sum_{i=1}^{n} a_i q^{-i}} [X_{k-1}^T \tilde{\theta}_{k-1}], \tag{4.5}$$

where $\tilde{\theta}_k = \theta - \hat{\theta}_k$, and q^{-1} indicates the time-delay operator. The recursive LMS adaptive IIR filter [14] takes the form of (LMS) with $\hat{\theta}_k$ and X_k defined as in (4.3) and (4.4). The sign-sign LMS adaptive IIR filter is of the same form as in (SS).

For unsigned LMS in the IIR case, convergence to the correct setting has been shown [7] in ideal use of inputs which are PE for LMS in the FIR case, if, in addition, the transfer function $\{1/1 - \Sigma_{i=1}^{n} a_i q^{-i}\}$ in (4.5) is SPR. An SPR discrete-time system transfer function is strictly minimum phase and has a positive real part when evaluated at any point on the unit circle. By rough analogy with the unsigned case, one might hypothesize that inputs which are PE for the sign-sign FIR filter (for which all trajectories ultimately settle in a small region about $\hat{\theta} = \theta$), plus satisfaction of the standard SPR condition on the transfer function in (4.5), would imply convergence to the correct setting for the sign-sign IIR filter. The apparently innocent input used above in the FIR case provides a counterexample to this hypothesis as will be demonstrated with quiver diagrams, and confirmed by simulations. It is, of course, possible to find inputs which do cause convergence in the IIR case considered here; one such input is $u_k = \{\cdots, -1, 1, 0.1, \cdots\}$.

Much of the analysis of the sign-sign FIR adaptive filter can be extrapolated to the sign-sign IIR form. The parameter-estimate space is again divided into regions by hypersurfaces. However, in this IIR case they are nonlinear functions of the inputs *and* the parameter estimates, and partition the space into regions other than conic sectors emanating from the desired setting as in the FIR case, making for a substantially more intricate quiver diagram.

Consider the first-order system and model ($m = n = 1$)

$$y_k = a y_{k-1} + b u_{k-1}$$
$$\hat{y}_k = \hat{a}_{k-1} \hat{y}_{k-1} + \hat{b}_{k-1} u_{k-1} \tag{4.6}$$

where a and b are to be estimated. Since the transfer function formed by dividing one by the autoregressive delay polynomial of a stable first-order plant is always SPR (i.e., Re $\{1/1 - ae^{-j\omega}\} > 0$ for all $|a| < 1$), an adaptive IIR filter of the form

$$\begin{bmatrix} \hat{a}_k \\ \hat{b}_k \end{bmatrix} = \begin{bmatrix} \hat{a}_{k-1} \\ \hat{b}_{k-1} \end{bmatrix} + \mu \operatorname{sgn} \begin{bmatrix} \hat{y}_{k-1} \\ u_{k-1} \end{bmatrix} \operatorname{sgn} [e_k] \tag{4.7}$$

seems appropriate. Recall, by the above hypothesis, that this algorithm should converge to the correct setting given the input $u_k = \{\cdots, -1.5, -1, 3, \cdots\}$, which leads to parameter convergence in the 2-dimensional FIR case.

Updates for the two estimates $\hat{a}$ and $\hat{b}$ can be considered separately:

$$\Delta \hat{a}_k = \mu \operatorname{sgn} [\hat{y}_{k-1}] \operatorname{sgn} [e_k] \tag{4.8a}$$

$$\Delta \hat{b}_k = \mu \operatorname{sgn} [u_{k-1}] \operatorname{sgn} [e_k]. \tag{4.8b}$$

Update (4.8a) will be zeroed when $\hat{y}_{k-1} = 0$ *or* $e_k = 0$. Both of these involve past values of the model output, which is a nonlinear function of the estimated parameters. For a small step size, the steady-state values for $\hat{y}_k$ can be accurately approximated by the periodic output resulting from a given periodic input and a fixed $\hat{\theta}_i$. The validity of this *steady-state* approximation is crucial to the analysis of the IIR algorithm; it does not hold, for example, for parameter estimate points outside the stable region or for excessively large step sizes.

Notice that except for the trivial case where $u_k = 0$, only the two cases where

$$\hat{y}_{k-i} = 0 \quad \text{and } e_{k-i+1} = 0, \quad (i = 1, 2, \cdots, \tau) \tag{4.9}$$

need be considered with u_k and, thus, y_k, $\hat{y}_k^q$ (from a fixed $\hat{\theta}_i$), and e_k^q all τ-periodic where the superscript q denotes steady-state values calculated for the purpose of drawing the quiver diagram in the IIR case. The solutions to (4.9) describe curves across which the sign of the update term changes for at least one regressor in the period, and divide the space into regions. As in the FIR case, all interior points of a region are updated identically for each particular regressor.

The quiver diagrams are substantially more complicated for the IIR filter, since the functions of the parameter estimates which zero the update term are nonlinear. These curves are obtained by using (4.6) to solve for steady-state, periodic sequences $\{y_k^q\}$ and $\{\hat{y}_k^q\}$ at a point in the parameter space by solving the set of simultaneous

equations

$$\left.\begin{aligned} y_1^q &= ay_3^q + bu_3 \\ y_2^q &= ay_1^q + bu_1 \\ y_3^q &= ay_2^q + bu_2 \end{aligned}\right\} \Rightarrow y_i^q$$

$$= \frac{a^2u_{i-3} + au_{i-2} + u_{i-1}}{1 - a^3}\, b, \qquad (i = 1, 2, 3). \tag{4.10}$$

The $\hat{y}_i^q$'s are obtained similarly, and their steady-state solution is the same as in (4.10) with a replaced by $\hat{a}^q$. Points for which the update is zeroed are found where

$$e_{k-i}^q = y_{k-i}^q - \hat{y}_{k-i}^q$$

$$= \frac{a^2u_{i-3} + au_{i-2} + u_{i-1}}{1 - a^3}\, b - \frac{(\hat{a}^q)^2u_{i-3} + \hat{a}^qu_{i-2} + u_{i-1}}{1 - (\hat{a}^q)^3}\, \hat{b}^q = 0 \tag{4.11}$$

and where

$$\hat{y}_{k-i}^q = \frac{(\hat{a}^q)^2u_{i-3} + \hat{a}^qu_{i-2} + u_{i-1}}{1 - (\hat{a}^q)^3}\, \hat{b}^q = 0, \qquad (i = 1, 2, 3). \tag{4.12}$$

Note that $\hat{b}^q = 0$ is always a solution of (4.12), yielding a horizontal line in the quiver diagram across which the update of $\hat{a}$ changes sign. Similarly, two more vertical lines are solutions to (4.12), also across which the sign of the $\hat{a}$ update changes for the associated regressor (these may not fall within the stable region). These lines cut the space into regions other than conical sections originating at the desired setting.

There are three sets of solutions to (4.11) and (4.12) for a given 3-periodic input sequence, one for each update in the period. These solutions, with input $u_k = \{\cdots, -1.5, -1, 3, \cdots\}$ and $(a, b) = (0.6, 1.0)$, are depicted in the quiver diagrams of Fig. 6 as the solid, dashed, and dot-dashed curves. In each region, a periodic update is shown as an arrow. These diagrams may appear at first sight to be a shocking miasma of arrows, due to the multitude of regions, each of which has an arrow.

This quiver diagram in Fig. 6 illustrates a remarkable behavior for the input $u_k = \{\cdots, -1.5, -1, 3, \cdots\}$, namely, that the algorithm may converge to any point on a locus in the parameter space; for this example, such a locus is marked in bold. This locus is along a boundary defined by the update-zeroing curve in the parameter space where the estimate chatters back and forth across the boundary in a local limit cycle. Such chattering corresponds to the periodic update trajectory in the neighboring regions being opposite in direction along the entire boundary between them. Notice that the correct parameterization in this example is an unstable equilibrium, since there are arrows pointing away from it, as seen in the closeup of the region surrounding the correct setting in Fig. 6(b). Simulations confirm this behavior. This example shows that sign-sign adaptive IIR filters can have multiple stable and unstable stationary points, and that the correct setting can be an unstable stationary point.

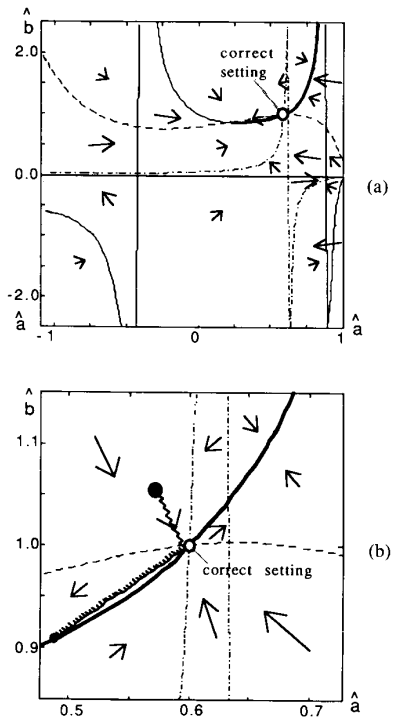

Fig. 6. Quiver diagram for two-dimensional sign-sign adaptive IIR algorithm with the input $\{\cdots, -1.5, -1, 3, \cdots\}$. The three boundaries are drawn as solid, dashed, and dot-dashed curves in the figures. (a) The true setting $(a, b) = (0.6, 1.0)$ is one point on a locus of stationary points as marked in bold. (b) A closeup of the correct setting. Arrows with radial component pointing away from (a, b) indicate this setting is unstable. A simulated trajectory depicted as a jagged line illustrates motion along sector arrows away from the correct setting to a setting on the locus of stationary points.

V. Conclusion

Quiver diagrams offer new insight into signed algorithm behavior. Beyond behavior in the ideal 2-dimensional case, quiver diagrams lend insight into higher dimensional problems in finding, e.g., a stability-inducing persistent excitation condition for sign-sign FIR LMS, as outlined in Section III. Quiver diagrams can also be extended [15] to Quantized Error (QE) versions of LMS, which use a finer quantizer than the one bit quantization of the signum function. In a quiver diagram for QE algorithms, a regressor yields as many boundaries in the parameter estimate space as there are quantization levels. An example of concepts related to quiver diagrams for

stochastic inputs is found in [16] and [17], where quiver-diagram-type insights are used to locate stationary points for a decision-directed equalizer with binary transmission. A related future application might be Decision Feedback Equalizers [18] when a signed-error-type LMS algorithm is used.

Behavior along the quiver diagram boundaries when multiple boundaries may be crossed remains an open theoretical issue. In higher dimensional cases ($m > 2$), multiple boundaries may be crossed where they intersect transversally far from the origin. Higher dimensional boundary behavior is the subject of present research.

Unfortunately, rough extrapolation of results from the unsigned LMS adaptive FIR and IIR filters offers incorrect expectations of the behavior of the sign-sign versions. Quiver diagrams show that the sign-sign adaptive IIR algorithm can converge to stable loci far from the correct setting, which cannot occur for the corresponding unsigned algorithm. These findings may have some bearing on the applications of sign-sign algorithms as in the CCITT 32 kbit/s standard ADPCM algorithm for use in telephony, which is a form of sign-sign adaptive IIR algorithm with leakage. Periodic update arrows as presented here may incorporate leakage, which is a small radial component pointing toward the origin of the parameter estimate space. Sign-sign adaptive IIR filters with regressor or prediction error filtering, the unsigned versions of which are discussed in [19], will have an effect on the appearance of the quiver diagram. Prediction error filtering which "straightens" the update zeroing curves can reduce the IIR case to something similar to the FIR case, which is conceptually simpler.

Acknowledgment

The encouragement from the Research Laboratory of Tellabs for examination of sign-sign adaptive IIR filters is gratefully acknowledged.

References

[1] B. Widrow and S. Stearns, *Adaptive Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1985.

[2] R. W. Lucky, "Techniques for adaptive equalization of digital communication systems," *Bell. Syst. Tech. J.*, vol. 45, pp. 255–286, Feb. 1966.

[3] N. S. Jayant and P. Noll, *Digital Coding of Waveforms: Principles and Applications to Speech and Video*. Englewood Cliffs, NJ: Prentice-Hall, 1984, section 6.5.3.

[4] C. R. Johnson, Jr., "Adaptive IIR filter(s) in the CCITT 32 kbps ADPCM standard," in *Proc. 19th Asilomar Conf. Circuits, Syst., Comput.*, Pacific Grove, CA, Nov. 1985, pp. 631–635.

[5] R. R. Bitmead, "Persistence of excitation conditions and the convergence of adaptive schemes," *IEEE Trans. Inform. Theory*, vol. IT-30, pp. 183–191, Mar. 1984.

[6] S. Dasgupta and C. R. Johnson, Jr., "Some comments on the behavior of sign-sign adaptive identifiers," *Syst. Contr. Lett.*, vol. 7, pp. 75–82, Apr. 1986.

[7] C. R. Johnson, Jr., M. G. Larimore, J. R. Treichler, and B. D. O. Anderson, "SHARF convergence properties," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-29, pp. 659–670, June 1981.

[8] C. R. Johnson, Jr., *Lectures on Adaptive Parameter Estimation*. Englewood Cliffs, NJ: Prentice-Hall, 1988.

[9] W. A. Sethares, I. M. Y. Mareels, B. D. O. Anderson, C. R. Johnson, Jr., and R. R. Bitmead, "Excitation conditions for signed regressor least mean square adaptation," *IEEE Trans. Circuits Syst.*, vol. 35, pp. 613–624, June 1988.

[10] A. Gersho, "Adaptive filtering with binary reinforcement," *IEEE Trans. Inform. Theory.*, vol. IT-30, pp. 191–199, Mar. 1984.

[11] L. Ljung and T. Söderström, *Theory and Practice of Recursive Identification*. Cambridge, MA: M.I.T. Press, 1983.

[12] B. D. O. Anderson, R. R. Bitmead, C. R. Johnson, Jr., P. V. Kokotovic, R. L. Kosut, I. M. Y. Mareels, L. Praly, and B. D. Riedle, *Stability of Adaptive Systems: Passivity and Averaging Analysis*. Cambridge, MA: M.I.T. Press, 1986.

[13] C. R. Elevitch, "Sign-sign adaptive LMS," Masters thesis, Cornell Univ., Jan. 1988.

[14] P. L. Feintuch, "An adaptive recursive LMS filter," *Proc. IEEE*, vol. 64, pp. 1622–1624, Nov. 1976.

[15] W. A. Sethares, "Quantized state adaptive algorithms," Ph.D. dissertation, Cornell Univ., Aug. 1987.

[16] J. E. Mazo, "Analysis of decision-directed equalizer convergence," *B.S.T.J.*, vol. 59, pp. 1857–1876, Dec. 1980.

[17] O. Macchi and E. Eweda, "Convergence analysis of self-adaptive equalizers," *IEEE Trans. Inform. Theory*, vol. IT-30, pp. 161–176, Mar. 1984.

[18] B. D. O. Anderson, R. A. Kennedy, and R. R. Bitmead, "Decision feedback equalizers: Concepts towards design guidelines," in *Proc. ISSPA87*, Brisbane, Australia, Aug. 1987, pp. 1–7.

[19] C. R. Johnson, Jr., "Adaptive IIR filtering: Current results and open issues," *IEEE Trans. Inform. Theory*, vol. IT-30, pp. 237–250, Mar. 1984.

Craig R. Elevitch was born in San Francisco, CA, on November 7, 1960. In 1983 he received the B.Sc. degree in systems science and the B.A. degree in philosophy from the University of California, San Diego. In the course of his undergraduate studies he spent $2\frac{1}{2}$ years as an exchange student. During the academic year 1980–1981 he participated in a one year exchange at Leeds University, England, studying philosophy. During the period 1982–1983 he spent a $1\frac{1}{2}$ year exchange at Lund Institute of Technology, Sweden, studying control systems. Immediately upon graduating from UCSD, he continued his studies for two additional years in Lund. In 1986 he began his studies at Cornell University and received the M.S. degree in January 1988.

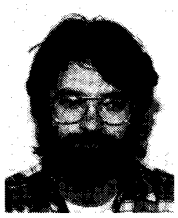

William A. Sethares (S'84–M'87) received the B.A. degree in mathematics from Brandeis University, Waltham, MA, in 1978 and the M.S. and Ph.D. degrees in electrical engineering from Cornell University, Ithaca, NY, in 1982 and 1987.

He has worked at the Raytheon Company as a Systems Engineer and is currently on the Faculty of the Department of Electrical and Computer Engineering at the University of Wisconsin in Madison. His research interests include adaptive systems, neural networks, and electronic music.

Gonzalo J. Rey was born in Buenos Aires, Argentina, in 1962. He received the B.S. degree in electrical engineering in 1986 from the Instituto Tecnologico Buenos Aires.

He is currently a graduate student at Cornell University, Ithaca, NY, and his research interest is in the dynamics of adaptive systems.

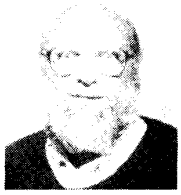

C. Richard Johnson, Jr. (S'74-M'77-SM'82) was born in Macon, GA, in 1950. He received the Ph.D. degree in electrical engineering with minors in engineering-economic systems and art history from Stanford University in 1977.

He is currently Professor and Associate Director of the School of Electrical Engineering at Cornell University, Ithaca, NY. His research interests are in adaptive parameter estimation theory as applied to digital filtering, identification, and control. He is the author of *Lectures on Adaptive Parameter Estimation* (Englewood Cliffs, NJ: Prentice-Hall, 1988).

Dr. Johnson was selected by Eta Kappa Nu as the Outstanding Young Electrical Engineer in 1982 and as the C. Holmes MacDonald Outstanding Teacher of 1983. He was a recipient of the 1982 Senior (Best Paper) Award of the Acoustics, Speech, and Signal Processing Society of the IEEE. He has served as an Editor of the *International Journal of Adaptive Control and Signal Processing* and as an Associate Editor for the IEEE TRANSACTIONS ON ACOUSTICS, SPEECH, AND SIGNAL PROCESSING, the IEEE TRANSACTIONS ON AUTOMATIC CONTROL, *Automatica*, and *Systems and Control Letters*.