

Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article was published in an Elsevier journal. The attached copy is furnished to the author for non-commercial research and education use, including for instruction at the author's institution, sharing with colleagues and providing to institution administration.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Fast communication

On the steady-state mean squared error of the fixed-point LMS algorithm

Mohamed Ghanassi*, Benoît Champagne, Peter Kabal

Department of Electrical & Computer Engineering, McGill University, 3480 University Street Montreal, Quebec, Canada H3A 2A7

Received 26 January 2007; received in revised form 1 May 2007; accepted 30 May 2007

Available online 13 June 2007

Abstract

This communication studies the quantization effects on the steady-state performance of a fixed-point implementation of the Least Mean Squares (LMS) adaptive algorithm. Based on experimental observations, we introduce a new *intermediate* mode of operation and develop a simplified theoretical approach to explain the behaviour caused by quantization effects in this mode. We also review the *stall* mode and provide a new expression that predicts the discontinuous behaviour of the steady-state mean squared error as a function of the input signal power. Combined with a previous analysis of quantization effects in *stochastic gradient* mode, this study provides analytical expressions for the steady-state mean squared error for the full range of step-size values. We present experimental results that are in a good agreement with theoretical predictions to validate our model.

© 2007 Elsevier B.V. All rights reserved.

Keywords: LMS; Fixed-point arithmetic; Quantization effects; Adaptive filtering

1. Introduction

The Least Mean Squares (LMS) algorithm [1] is widely used in adaptive filtering. This algorithm recursively updates the vector of coefficients $\mathbf{w}(k) = [w_0(k), w_1(k), \dots, w_{N-1}(k)]^t$ of an FIR filter according to the following equations

$$e(k) = y(k) - \mathbf{w}^t(k)\mathbf{x}(k), \quad (1a)$$

$$\mathbf{w}(k+1) = \mathbf{w}(k) + \mu e(k)\mathbf{x}(k), \quad (1b)$$

where k is the discrete time index, $e(k)$ is the estimation error, $\mathbf{x}(k) = [x(k), x(k-1), \dots, x(k-N+1)]^t$ is the input signal vector, $y(k)$ is the reference signal to be estimated, and μ is an

adaptation parameter. The speed of convergence of the algorithm towards the optimal Wiener solution, as well as the power of the residual error after convergence (i.e., in steady-state) depend on μ . These properties of the LMS are well established for infinite precision arithmetic (see [2] and references therein).

When LMS is implemented on a fixed-point processor, quantization errors affect its performance and the mean squared error (MSE), i.e. $E\{|e(k)|^2\}$ where $E\{\cdot\}$ denotes statistical expectation, may be significantly higher than the one expected in infinite precision. Analysis of the quantization effects on LMS performance goes back to the work of Gitlin et al. [3] who studied the variations of the MSE as a function of the step-size in a digitally implemented LMS. They reported two main

*Corresponding author. Tel.: +1 514 398 2002.

E-mail address: mohamed.ghanassi@mcgill.ca (M. Ghanassi).

observations: (a) the MSE after the algorithm converges is much higher than the one expected from a quantization of the algorithm variables and (b) due to quantization effects, the adaptation may stop and in this case, the MSE may be actually reduced by increasing the step-size. Caraiscos and Liu [4] presented an analysis of quantization errors in steady-state for fixed-point and floating-point arithmetic. In their analysis, they modelled quantization errors as white noise and obtained an analytical expression for the residual error. However, the error model they used is only valid when the adaptation is not stopped by quantization effects—in this situation, the quantization has a low impact on the steady-state MSE. Alexander [5] used the same white noise model for quantization errors to analyse the behaviour of the finite precision LMS algorithm in the transient regime.

Bermudez and Bershada [6] recognized the drawback of the error model used in [4] and [5] and its non-validity in a situation where adaptation is stopped by quantization effects. They proposed a non-linear analytical model for the quantization function. By using a conditional moment technique, for a white Gaussian input and a small adaptation step size, they derived recursive equations which can be numerically solved to give the MSE in both transient and steady-state regimes [6]. They investigated the steady-state behaviour of the quantized LMS algorithm for small step-size and showed that the stalling behaviour is indeed a “slow-down” phenomenon. Under limiting assumptions, their model predicts a steady-state MSE that is nearly independent on the number of bits [7]. In our study, we observed the steady-state behaviour of the LMS for all range of values of the adaptation step-size and noticed a clear difference between our simulation results and the MSE predicted by the model in [7], as we will show in Section 3.

The studies cited above explained and modelled the LMS behaviour in different conditions but did not provide an analytical expression of the steady-state MSE (SS-MSE) for all these conditions. In particular, the intermediate region between the *stall* mode, where adaptation is stopped, and the *stochastic gradient* mode, where the analysis of [4] is applicable, has not been previously investigated. In this work, we introduce a new *intermediate* mode to characterize the algorithm behaviour in this region and develop a simplified theoretical model that provides an analytic expression of the corresponding SS-MSE values for a white stationary

Gaussian input signal. We also review the stall mode and provide a new expression that predicts the discontinuous behaviour of the SS-MSE as a function of the input signal power. Combined with the analysis of quantization effects in stochastic gradient mode, this study provides analytical expressions for the SS-MSE for the complete range of step-size parameter values. In particular, the value of step-size corresponding to the onset of the stall mode can be predicted accurately, so that stalling can be avoided by judiciously choosing the step size value.

The outline of the paper is as follows. In Section 2 we present the theoretical analysis and developments leading to analytical expressions for the SS-MSE for different operating conditions of the algorithm. In Section 3 we present experimental results. A brief conclusion follows in Section 4.

2. Theoretical analysis

In our analysis of the finite precision LMS algorithm we assume that the input signal $x(k)$ is a white stationary Gaussian process with zero mean and variance σ_x^2 , and the reference signal $y(k)$ is written as

$$y(k) = \mathbf{w}_o^t \mathbf{x}(k) + n(k), \quad (2)$$

where \mathbf{w}_o is the optimal vector of coefficients and $n(k)$ is a white stationary Gaussian noise independent of $x(k)$, with zero mean and variance σ_n^2 . We assume that signals and filter coefficients are real-valued but this analysis can be easily generalized to the complex case.

Under the above assumption on the input signal $x(k)$, the LMS algorithm converges in the mean square for values of μ in the range [8]

$$0 < \mu < 2/(N\sigma_x^2), \quad (3)$$

where N is assumed to be large.¹ The SS-MSE, i.e. after the algorithm converges, is given by [4]

$$\xi = \lim_{k \rightarrow \infty} E\{|e(k)|^2\} = \frac{\xi_{\min}}{1 - \mu N \sigma_x^2 / 2}, \quad (4)$$

where ξ_{\min} is the MSE of the optimum filter, equal to σ_n^2 if the adaptive filter is long enough to cover the impulse response to be estimated.

¹Variations of this result can be found in the literature (see [2] and references therein). However, for the model under consideration here (i.e. white stationary Gaussian processes and large N), the values predicted by (3) are sufficiently accurate and in good agreement with experimental observations.

In fixed-point arithmetic, signals, variables, and parameters used in the algorithm as well as operation results are quantized to discrete values. We will use a fractional fixed-point representation with values lying in the interval $[-1, +1]$. We assume that addition does not lead to overflow and that only multiplication induces quantization errors. We consider a roundoff quantization using b bits (including sign) and represented by a generic function $Q[\cdot]$. We use primed symbols to represent quantities of the finite precision LMS, which can then be described by the following equations:

$$e'(k) = y'(k) - Q[\mathbf{w}'(k)^T \mathbf{x}'(k)], \quad (5a)$$

$$\delta \mathbf{w}(k) = Q[\mu e'(k)] \mathbf{x}'(k), \quad (5b)$$

$$\mathbf{w}'(k+1) = \mathbf{w}'(k) + Q[\delta \mathbf{w}(k)], \quad (5c)$$

where $Q[\cdot]$ applied to a vector acts on each of its components. All three quantizers in (5) are considered in our analysis.

In (5), the quantization operation $Q[\cdot]$ cannot be modelled in general by a smooth function and therefore an analytic expression for the SS-MSE cannot be easily derived. In the existing literature on the finite-precision LMS, we can distinguish two qualitative operating modes of the algorithm. In the first mode, the components of the updating vector $\delta \mathbf{w}(k)$ are smaller in absolute value than $\Delta/2$, where $\Delta = 2^{1-b}$ represents the least significant bit, and are quantized to zero in (5c). The adaptation is virtually stopped by quantization effects and the algorithm is in *stall* mode [3]. In the second mode, the components of $\delta \mathbf{w}(k)$ are much larger in absolute value than $\Delta/2$, and their quantization can be modelled as an additive white noise (i.e. linear noise model). With fluctuations of the stochastic gradient, the algorithm continues to adapt as in the infinite precision case [4]. We refer to this mode as *stochastic gradient* mode. In this work, we introduce a third mode as being *intermediate* between the stall and stochastic gradient modes.

2.1. Stall mode

During the adaptation, while the algorithm in (5) tends to the steady-state, the amplitude of the error $e'(k)$ decreases with the number of iterations until the components of $\delta \mathbf{w}(k)$ become smaller than $\Delta/2$ in absolute value. Hence, the adaptation of the filter coefficients stops and the algorithm enters the stall mode. To evaluate the resulting SS-MSE, we

assume that adaptation stops when the components of the quantized updating vector $Q[\delta \mathbf{w}(k)]$ in (5c), where $\delta \mathbf{w}(k)$ is defined in (5b), become equal to zero with a high probability, which is expressed as

$$\Pr[|Q[\mu e'(k)]x'(l)| < \Delta/2] \simeq 1, \quad (6)$$

where the time index $l \in \{k - N + 1, \dots, k\}$. In practice, we find that the stall mode is achieved when the absolute value of $Q[\mu e'(k)]$ becomes smaller than a certain threshold. Since the quantity $Q[\mu e'(k)]$ takes values which are integer multiples of Δ (i.e. $Q[\mu e'(k)] = \pm p\Delta$, with $p \in \{0, 1, 2, \dots\}$), this threshold corresponds to a maximum value of p and is written as $p_{\max}\Delta$. Assuming that the processes $e'(k)$ and $x(l)$ are independent, the condition (6) is satisfied whenever

$$\Pr[|\mu e'(k)| \leq (p_{\max} + 1/2)\Delta] \simeq 1, \quad \text{and} \quad (7)$$

$$\Pr[|x'(l)| < 1/(2p_{\max})] \simeq 1. \quad (8)$$

To calculate the SS-MSE, we have to determine a value for p_{\max} appropriate for use in (7) and (8). For a signal $x(l)$ with an amplitude much larger than Δ and for a small value of p_{\max} , we can replace $x'(l)$ by $x(l)$ in (8). For the Gaussian process $x(l)$, we can argue that (8) is valid for $1/(2p_{\max}) \geq l_x \sigma_x$, where l_x is a number that we can reasonably set to 3 or more (for $l_x = 3$, the probability is 0.997).² The corresponding p_{\max} is expressed as

$$p_{\max} = \left\lfloor \frac{1}{2l_x \sigma_x} \right\rfloor, \quad (9)$$

where $\lfloor x \rfloor$ is the greatest integer smaller than or equal to x .

In a similar way, (7) will be valid if

$$(p_{\max} + 1/2)\Delta = l_e \mu \sqrt{\xi'}, \quad (10)$$

where $\xi' = E\{|e'(k)|^2\}$ is the SS-MSE and l_e is a positive number we can reasonably set to 3 or more. Replacing p_{\max} from (9) into (10), the desired SS-MSE is obtained as

$$\xi' \simeq \frac{\Delta^2}{l_e^2 \mu^2} \left(\left\lfloor \frac{1}{2l_x \sigma_x} \right\rfloor + \frac{1}{2} \right)^2. \quad (11)$$

This result provides a simple expression of the SS-MSE as a function of the step-size μ , the number of bits (via $\Delta = 2^{1-b}$), and the input signal power σ_x^2 . While the general behaviour of (11) with respect to these parameters is consistent with earlier studies

²More generally, for any zero-mean random variable x , it follows from Chebyshev's inequality that if $1/(2p_{\max}) \geq l_x \sigma_x$ then $\Pr[|x| < 1/(2p_{\max})] > 1 - 1/l_x^2$.

[3], the presence of the floor function in (11) is new. It allows for a more accurate description of the SS-MSE versus σ_x relationship, as clearly evidenced by the experimental results in Section 3 (see Fig. 4).

2.2. Stochastic gradient mode

In this mode, the algorithm converges before entering the stall mode. The filter coefficients continue to adjust in a random way around the optimal Wiener solution, due to the stochastic gradient noise. The components of $\delta\mathbf{w}(k)$ in (5) have an absolute value much larger than $\Delta/2$, on average, and their quantization only introduces a small error to their values. Thus, we can describe all quantization errors in the algorithm as additive white noise. To obtain an analytic expression of the SS-MSE (needed in the analysis of the intermediate mode), we use the approach of Caraiscos and Liu [4], modified to allow for the quantization of the term $\mu e'(k)$ in (5b), which is not considered in their original paper.

The various quantization errors in the finite precision LMS are defined as follows:

$$q_y(k) = y'(k) - y(k), \quad (12a)$$

$$\mathbf{q}_x(k) = \mathbf{x}'(k) - \mathbf{x}(k), \quad (12b)$$

$$q_{wx}(k) = Q[\mathbf{w}'(k)^t \mathbf{x}'(k)] - \mathbf{w}'(k)^t \mathbf{x}'(k), \quad (12c)$$

$$q_{\mu e}(k) = Q[\mu e'(k)] - \mu e'(k), \quad (12d)$$

$$\mathbf{q}_{\delta w}(k) = Q[\delta\mathbf{w}(k)] - \delta\mathbf{w}(k). \quad (12e)$$

The following assumptions are made (see also [2,9] for a further discussion): the errors $q_y(k)$, $q_{wx}(k)$, and $q_{\mu e}(k)$ have zero mean and respective variances Σ_y^2 , Σ_{wx}^2 , and $\Sigma_{\mu e}^2$; the error vectors $\mathbf{q}_x(k)$ and $\mathbf{q}_{\delta w}(k)$ have independent components of zero mean and respective variances Σ_x^2 and $\Sigma_{\delta w}^2$; and finally, the quantization error of a variable is independent of that variable, of other variables and of their respective quantization errors.

Using these expressions in (5) and neglecting the quantization errors terms of order greater than one, a development similar to the one presented in [4] (see this reference for details) leads to the following expression for the SS-MSE for the finite precision LMS algorithm:

$$\xi' \simeq \frac{1}{1 - \mu N \sigma_x^2 / 2} \left[\xi_{\min} + \|\mathbf{w}_o\|^2 \Sigma_x^2 + \Sigma_y^2 + \Sigma_{wx}^2 + \frac{N}{2\mu} (\sigma_x^2 \Sigma_{\mu e}^2 + \Sigma_{\delta w}^2) \right]. \quad (13)$$

In expression (7) of [4], replacing the trace of the correlation matrix ($\text{tr}\mathbf{R}$) with $N\sigma_x^2$, which corresponds to our assumption on the input signal $x(k)$, leads to a result similar to (13). However, in [4], authors assumed that, as for a particular LMS implementation in which the step size value is a power-of-two, the elements of the term $Q[\delta\mathbf{w}(k)]$ in (5c) are calculated by first computing the product $e'(k)x'(k)$, right shifting the result and then quantizing. In this case, there is no quantization error for the product $\mu e'(k)$ and thus there is no term corresponding to $\Sigma_{\mu e}^2$.

2.3. Intermediate mode

In this analysis of the quantization effects we emphasize the important role of the step-size μ in the operation of the quantized LMS algorithm. When all other parameters used in the algorithm are fixed we can explain all the modes of operation of the algorithm after convergence in terms of the choice of the value of μ .

For small values of μ , the algorithm is in stall mode and the SS-MSE decreases according to (11) (i.e., as $1/\mu^2$) as μ increases. However, this effect must be limited since the SS-MSE cannot decrease below a limiting value ξ_{\lim} for which the vector of estimated coefficients becomes equal to a quantized value of the optimal Wiener solution \mathbf{w}_o .

A simple expression for ξ_{\lim} can be obtained by substituting $\mathbf{w}_o = Q[\mathbf{w}_o] - \mathbf{q}_w$ in (2), where \mathbf{q}_w is a vector whose components are independent quantization errors of zero mean and variance Σ_w^2 . This results in

$$y(k) = Q[\mathbf{w}_o]^t \mathbf{x}(k) + n'(k), \quad (14)$$

where the uncorrelated modelling error is now given by

$$n'(k) = n(k) - \mathbf{q}_w^t \mathbf{x}(k). \quad (15)$$

The limiting SS-MSE, defined as $\xi_{\lim} = E\{|n'(k)|^2\}$, is

$$\xi_{\lim} = \xi_{\min} + N\sigma_x^2 \Sigma_w^2. \quad (16)$$

In a fixed-point context, the best result that can be aimed for is that the LMS converges to the quantized value of the optimal vector of coefficients $Q[\mathbf{w}_o]$. Accordingly, it is reasonable to replace the reference signal model (2) by (14) and thus, ξ_{\lim}

should be used in place of ξ_{\min} in the SS-MSE expressions.³

When the step-size μ increases and the SS-MSE gets closer to the limiting value in (16), the algorithm leaves the stall mode to enter in the stochastic gradient mode, where the quantization errors contribute in the algorithm as additive white noise. Simulations results of the SS-MSE as a function of μ showed that for a limited range of values for μ , corresponding to a region between the stall and stochastic gradient modes, SS-MSE values cannot be explained by expressions (11) or (13), corresponding to these modes. Based on this, we introduce an *intermediate* mode, corresponding to the transitional region between the stall and the stochastic gradient modes.

We consider this mode as a modified stochastic gradient mode where the components of the updating vector $\delta\mathbf{w}(k)$ (5b) are still small in absolute value, but may exceed with non-zero probability the first quantizer decision level $\Delta/2$. Specifically, to obtain a simplified expression of the steady-state MSE in the intermediate mode, we make the following additional assumptions:

- A1. The quantizer in (5c) is operated in the small input regime. That is, the variance of the components of $\delta\mathbf{w}(k)$ is of the order of $\Delta^2/12$.
- A2. The uniform quantizer $Q[\mu e'(k)]$ in (5b) approximately satisfies the centroid condition for minimum squared error distortion [9].

All other quantization errors are treated as in the analysis of the stochastic gradient mode (i.e., independent additive noise model).

For a uniform quantizer in the small input regime, the variance of the quantization error can be approximated by that of its input.⁴ The l th component of the updating vector $\delta\mathbf{w}(k)$ (5b) is given by

$$\delta w_l(k) = Q[\mu e'(k)]x'(k-l). \quad (17)$$

Thus, under assumption A1, the variance of the quantization error in (5c) can be approximated as

$$\Sigma_{\delta w}^2 \simeq \text{var}\{Q[\mu e'(k)]x'(k-l)\} = \sigma_x^2 E\{Q[\mu e'(k)]^2\}, \quad (18)$$

³For the stochastic gradient mode, this modification does not result in appreciable changes of the SS-MSE values over the range of interest of μ .

⁴See [10] for a more precise analytical characterization of the input–output variance relationship.

where the second equality follows from the zero-mean and independence assumptions.

For a quantizer satisfying the centroid condition, the quantizer output is uncorrelated with the quantization error (see [9, Lemma 6.2.2]). Thus, invoking (12d) and assumption A2, we obtain

$$\mu^2 \xi' = \Sigma_{\mu e}^2 + E\{Q[\mu e'(k)]^2\}. \quad (19)$$

Combining (18) and (19), we obtain⁵

$$\sigma_x^2 \Sigma_{\mu e}^2 + \Sigma_{\delta w}^2 = \mu^2 \sigma_x^2 \xi'. \quad (20)$$

Substituting this result in (13), replacing ξ_{\min} by ξ_{\lim} (16) and solving for ξ' , we finally obtain the desired SS-MSE expression:

$$\xi' \simeq \frac{1}{1 - \mu N \sigma_x^2} [\xi_{\lim} + \|\mathbf{w}_o\|^2 \Sigma_x^2 + \Sigma_y^2 + \Sigma_{wx}^2]. \quad (21)$$

In practice, the quantity $\mu e'(k)$ will not be uniformly distributed between adjacent quantization levels, so the uniform quantizer in (5b) will not exactly satisfy the centroid condition (i.e. the quantization levels are not optimal) and the resulting SS-MSE may exceed that prescribed by (21).

2.4. Summary of modes

Table 1 summarizes the different operating modes of the fixed-point LMS algorithm with the corresponding expressions for the steady-state MSE. The step-size values μ_L and μ_U delimiting the three regions can be obtained by equating the appropriate expressions for ξ' and solving for μ .

3. Experimental study

To validate our analysis, we consider an application of the fixed-point LMS adaptive filter to the identification of a linear system. The input signal $x(k)$ is a stationary white Gaussian process with zero mean and variance $\sigma_x^2 = 0.1$. The reference signal $y(k)$ is obtained via convolution of $x(k)$ with a vector \mathbf{w}_o of length $N = 100$. We used this signal without additional noise to ensure that the SS-MSE is only due to quantization effects. In the fixed-point LMS, the signal $x'(k) = Q[x(k)]$ is obtained by a uniform quantization of $x(k)$ and the samples with a value outside the interval $[-1, +1)$ are saturated to

⁵While this derivation is based on the general LMS implementation (not a power-of-two step size case), the results can also be applied to the power-or-two case by setting the variable $\Sigma_{\mu e}$ to zero in (20).

Table 1
Summary of fixed-point LMS modes and corresponding steady-state MSE values

Mode	μ	Steady-state MSE, ζ'
Stall	$\mu \leq \mu_L$	$\frac{A^2}{l_e^2 \mu^2} \left(\left\lfloor \frac{1}{2l_x \sigma_x} \right\rfloor + \frac{1}{2} \right)^2$
Intermediate	$\mu_L \leq \mu \leq \mu_U$	$\frac{1}{1 - \mu N \sigma_x^2} [\zeta_{\min} + \ \mathbf{w}_o\ ^2 \Sigma_x^2 + \Sigma_y^2 + \Sigma_{wx}^2 + N \sigma_x^2 \Sigma_w^2]$
Stochastic gradient	$\mu_U \leq \mu$	$\frac{1}{1 - \mu N \sigma_x^2 / 2} [\zeta_{\min} + \ \mathbf{w}_o\ ^2 \Sigma_x^2 + \Sigma_y^2 + \Sigma_{wx}^2 + \frac{N}{2\mu} (\sigma_x^2 \Sigma_{\mu e}^2 + \Sigma_{\delta w}^2)]$

± 1 ; the reference signal $y'(k) = Q[y(k)]$ is similarly obtained. The optimal vector \mathbf{w}_o was randomly generated and scaled so that $\|\mathbf{w}_o\|^2 = 0.4$.

We developed a set of C routines that emulate arithmetic operations of a fixed-point DSP [11] with a number of bits $b = 16$ and implemented the fixed-point LMS on a personal computer using these routines. We also implemented the fixed-point LMS on a fixed-point DSP emulator software provided by Texas Instruments and validated the accuracy of our routines. We performed proper scaling of the various signals relative to the available dynamic range so that saturation effects are avoided at initial stage of operation of the algorithm.

To simulate different modes of operation of the algorithm, we varied the step-size μ and measured the SS-MSE. Two slightly different implementations of the error calculation in the fixed-point LMS were considered:

- *Without accumulator*: to emphasize the quantization effects, the elements of the inner product $\mathbf{w}^t(k)\mathbf{x}(k)$ were individually quantized and then added together.
- *With accumulator*: the individual products were computed in double-precision, summed and the final result was quantized.

For all values of μ used in our experiment, we ran simulations for up to several millions samples. Fig. 1 shows experimental results versus the number of iterations for the fixed-point LMS with an intermediate accumulator and for three different μ values of 0.002, 0.075 and 0.16 corresponding, respectively, to stall, intermediate and stochastic gradient modes. For all three μ values the steady-state behaviour is clearly reached. As we show in the inset of Fig. 1 for $\mu = 0.002$, there is a difference between our simulations results and the SS-MSE




Fig. 1. MSE versus number of iterations for fixed-point LMS with an intermediate accumulator ($b = 16$, $N = 100$). The three curves are for three different μ values corresponding to the three modes (0.002 for stall, 0.075 for intermediate, and 0.16 for stochastic gradient). The inset shows MSE for $\mu = 0.002$ and for a larger number of iterations (up to 30 millions) as well as values obtained using the recursive equations (18) and (26) from [6] (the smooth curve).

predicted by Bermudez's model using Eqs. (18) and (26) from [6].

Figs. 2 (without accumulator) and 3 (with accumulator) show the theoretical SS-MSE versus μ for the stall, intermediate, and stochastic gradient modes, as well as the limiting SS-MSE, defined, respectively, in (11), (21), (13), and (16). In the evaluation of these SS-MSE expressions, we use $\Sigma_x^2 = \Sigma_y^2 = \Sigma_w^2 = \Sigma_{\mu e}^2 = \Sigma_{\delta w}^2 = A^2/12$, and $\Sigma_{wx}^2 = c A^2/12$ where $c = N$ in the implementation without accumulator and $c = 1$ in the one with accumulator.

The results in Figs. 2 and 3 show the different modes of operation of the algorithm. For small values of μ , the algorithm operates in the stall mode and the SS-MSE decreases with increasing μ . As shown in the figures, the decrease in the SS-MSE in




Fig. 2. Steady-state MSE versus μ for fixed-point LMS without an intermediate accumulator ($b = 16$, $N = 100$). The three solid curves are for the three modes (stall, intermediate, and stochastic gradient). The circles show the experimental results. The horizontal line is the limiting SS-MSE ξ_{lim} .




Fig. 3. Steady-state MSE versus μ for fixed-point LMS with an intermediate accumulator ($b = 16$, $N = 100$). The three solid curves are for the three modes (stall, intermediate, and stochastic gradient). The circles show the experimental results. The horizontal line is the limiting SS-MSE ξ_{lim} .

the stall mode is well described by a $1/\mu^2$ function, such as that given by (11) with $l_x = 3.2$ and $l_e = 4.5$.

As the step size μ increases, the SS-MSE decreases to reach a minimum value corresponding to the boundary point μ_L between the stall and intermediate modes; as μ is increased beyond this limit, the SS-MSE starts increasing and the algorithm leaves the stall mode to enter the intermediate mode. Increasing μ further leads to the stochastic gradient mode.

Fig. 4 shows experimental results of the SS-MSE in stall mode as a function of the variance σ_x^2 of the

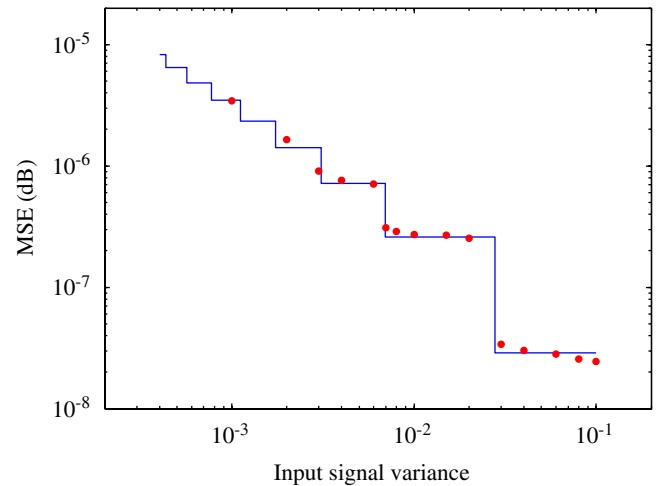


Fig. 4. Steady-state MSE in stall mode versus the input signal variance, for $b = 16$, $N = 100$, $\mu = 0.02$. The solid line is from (11). The circles are experimental points.

input signal $x(k)$. We have a good agreement with the expression of (11) with parameters $l_x = 3.2$ and $l_e = 4.5$. In particular, the presence of the floor function in (11), not present in earlier analysis of the stall mode, enables a good prediction of the stair case behaviour of the SS-MSE as a function of the input signal variance.

In the intermediate and stochastic gradient modes, the SS-MSE increases with μ . Experimental results of the SS-MSE for these modes are in a very good agreement with theoretical results of (13) and (21).

4. Conclusion

In this paper, we studied the impact of quantization effects on the steady-state behaviour of the LMS algorithm. We introduced a new *intermediate* mode to characterize the algorithm behaviour for a specific range of step-size parameter values and developed a simplified theoretical model that provides an analytic expression of the corresponding steady-state MSE values for a white stationary Gaussian input signal. We also reviewed the *stall* mode and provided a new expression that predicts the discontinuous behaviour of the steady-state MSE as a function of the input signal power. Combined to the analysis of quantization effects in stochastic gradient mode, our study provides unifying set of analytic expressions for the steady-state MSE for the complete range of step-size parameter values. Experimental results show good agreement with our theoretical analysis.

References

- [1] B. Widrow, S.D. Stearns, Adaptive Signal Processing, Prentice-Hall, Englewood Cliffs, NJ, 1985.
- [2] S. Haykin, Adaptive Filter Theory, fourth ed., Prentice-Hall, Englewood Cliffs, NJ, 2002.
- [3] R.D. Gitlin, J.E. Mazo, M.G. Taylor, On the design of gradient algorithms for digitally implemented adaptive filters, *IEEE Trans. Circuit Theory* 20 (2) (March 1973) 125–136.
- [4] C. Caraiscos, B. Liu, A roundoff error analysis of the LMS adaptive algorithm, *IEEE Trans. Acoust. Speech Signal Process.* 32 (1) (February 1984) 34–41.
- [5] S.T. Alexander, Transient weight misadjustment properties for the finite precision LMS algorithm, *IEEE Trans. Acoust. Speech Signal Process.* 35 (9) (September 1987) 1250–1258.
- [6] J.C.M. Bermudez, N.J. Bershad, A nonlinear analytical model for the quantized LMS algorithm—the arbitrary step size case, *IEEE Trans. Signal Process.* 44 (5) (May 1996) 1175–1183.
- [7] N.J. Bershad, J.C.M. Bermudez, New insights on the transient and steady-state behavior of the quantized LMS algorithm, *IEEE Trans. Signal Process.* 44 (10) (October 1996) 2623–2625.
- [8] T.C. Hsia, Convergence analysis of LMS and NLMS adaptive algorithms, *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 2, Boston, MA, April 1983, pp. 667–670.
- [9] A. Gersho, R.M. Gray, *Vector Quantization and Signal Compression*, Kluwer, Dordrecht, 1992.
- [10] A.B. Sripad, D.L. Snyder, A necessary and sufficient condition for quantization errors to be uniform and white, *IEEE Trans. Acoust. Speech Signal Process.* 25 (5) (October 1977) 442–448.
- [11] TMS320C54x DSP Reference Set, Volume 4: Application Guide, Texas Instruments, 1996.