

A NEW ROBUST FREQUENCY DOMAIN ECHO CANCELLER WITH CLOSED-LOOP LEARNING RATE ADAPTATION

Jean-Marc Valin and Iain B. Collings

CSIRO ICT Centre, Sydney, Australia

{jean-marc.valin, iain.collings}@csiro.au

ABSTRACT

One of the main difficulties in echo cancellation is the fact that the learning rate needs to vary according to conditions such as double-talk and echo path change. Several methods have been proposed to vary the learning. In this paper we propose a new closed-loop method where the learning rate is proportional to a misalignment parameter, which is in turn estimated based on a gradient adaptive approach. The method is presented in the context of a multidelay block frequency domain (MDF) echo canceller. We demonstrate that the proposed algorithm outperforms current popular double-talk detection techniques by up to 6 dB.

Index Terms—Acoustic echo cancellation, adaptive filters, learning rate, double-talk

1. INTRODUCTION

In any echo cancellation system, the presence of near end speech (double-talk) tends to make the adaptive filter diverge. To counter the effect, robust echo cancellers require adjustment of the learning rate to account for the presence of double-talk in the signal.

Most echo cancellation algorithms attempt to explicitly detect double-talk [1] conditions and then react by freezing the coefficients of the adaptive filter (setting the learning rate to zero). Reliable double-talk detection is a difficult problem and sometimes it is not clear what should be considered as double-talk, especially in an acoustic echo cancellation context with stationary background noise.

In previous work [2], we proposed a new approach to make echo cancellation more robust to double-talk. Instead of attempting to explicitly detect double-talk conditions, a continuous learning rate was used. The learning rate depends on a misalignment estimate, which is obtained through a linear regression. While the technique gives good results, the estimation of the misalignment remains a difficult problem.

In this paper, we propose a new approach where the misalignment is estimated in closed-loop based on a gradient adaptive approach. This closed-loop technique is applied to the block frequency domain (MDF) adaptive filter [3] and shows a significant improvement over previous approaches.

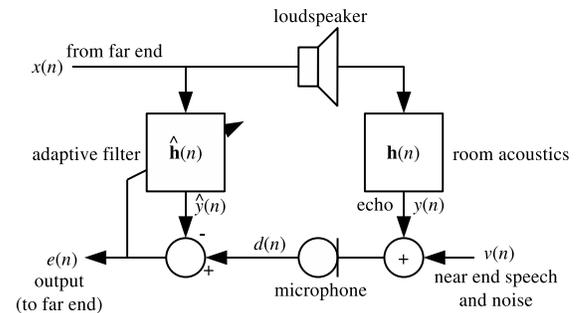


Fig. 1. Block diagram of echo cancellation system.

In Section 2, we discuss the optimal learning rate in presence of noise. Section 3 introduces the multidelay frequency domain (MDF) adaptive filter. In Section 4, we propose a gradient adaptive technique for adjusting the learning rate of the MDF algorithm. Experimental results and a discussion are presented in Section 5 and Section 6 concludes this paper.

2. OPTIMAL LEARNING RATE IN THE PRESENCE OF DOUBLE-TALK

In the acoustic echo cancellation context illustrated in Fig. 1, the speech signal $x(n)$ received from the far end is played into a loudspeaker. The microphone signal $d(n)$ captures both the convoluted far end speech $y(n)$ and the near end noisy speech $v(n)$. The adaptive filter attempts to estimate the impulse response $\hat{h}(n)$ to be as close as possible to the real impulse response $h(n)$. The estimated echo $\hat{y}(n)$ is subtracted from $d(n)$, so the output signal $e(n)$ contains both double-talk $v(n)$ and residual echo $r(n) = y(n) - \hat{y}(n)$.

The conventional approach to double-talk robustness consists of setting the learning rate to zero when double-talk is detected. Double-talk detectors [1] are thus an important aspect of the approach. Unfortunately, they are sometimes unreliable, especially in acoustic echo cancellation context when background noise is present. In this paper, we investigate continuous learning rates that do not depend on a binary double-talk decision.

Whenever an adaptive filter is not perfectly adjusted, its residual signal $r(n)$ can be used to gain better information about the exact (time-varying) filter weights $\mathbf{h}(n)$. However, the amount of information about $\mathbf{h}(n)$ present in $e(n)$ decreases with the amount of noise and near end speech $v(n)$. In the case of the normalised least mean square (NLMS) filter, it means that the stochastic gradient becomes less reliable when the noise increases or when the filter misalignment decreases (as the filter converges). The theoretical optimal learning rate is approximately equal to the residual-to-error ratio [2]:

$$\mu_{opt}(n) \approx \frac{E\{r^2(n)\}}{E\{e^2(n)\}} \quad (1)$$

where $r(n) = y(n) - \hat{y}(n)$ is the (unknown) residual echo and $e(n)$ is the error signal.

One possible method to vary the learning rate $\mu(n)$ would be to use the generalized normalized gradient descent (GNGD) algorithm [4], which includes the NLMS learning rate:

$$\mu(n) = \frac{\mu_0 E\{x^2(n)\}}{E\{x^2(n)\} + \epsilon(n)} \quad (2)$$

where $\epsilon(n)$ is adapted based on the NLMS stochastic gradient behaviour. To examine $\epsilon(n)$ more closely, it is reasonable to surmise that (2) eventually converges to the optimal learning rate defined by (1). Assuming steady state behaviour ($\epsilon(n)$ is stable) and $\mu_0 = 1$, we find that:

$$\frac{E\{r^2(n)\}}{E\{r^2(n)\} + \gamma(n)\epsilon(n)} \approx \frac{E\{r^2(n)\}}{E\{e^2(n)\}} \quad (3)$$

where $\gamma = E\{r^2(n)\}/E\{x^2(n)\}$. Knowing that we have $E\{e^2(n)\} = E\{r^2(n)\} + E\{v^2(n)\}$, we find the relation $\epsilon(n) \approx E\{v^2(n)\}/\gamma(n)$. In other words, the gradient-adaptive parameter $\epsilon(n)$ is approximately proportional to the variance of the near-field signal and independent of the far-field signal. Because $\epsilon(n)$ can only be adapted slowly over time, it is clear that (2) implicitly assumes that $E\{v^2(n)\}$ also varies slowly. While this is a reasonable assumption in applications where the background noise is constant or slowly varying, the assumption does not hold for acoustic echo cancellation, where double-talk can start or stop at any time.

In previous work [2], we proposed to use (1) directly to adapt the learning rate. While $E\{e^2(n)\}$ can easily be estimated, the estimation of the residual echo $E\{r^2(n)\}$ is difficult because one does not have access to the real filter coefficients. One reasonable assumption we make is that:

$$E\{r^2(n)\} = \eta(n) E\{y^2(n)\} \approx \eta(n) E\{\hat{y}^2(n)\} \quad (4)$$

where $\eta(n)$ is the normalised filter misalignment (or the inverse of the echo return loss enhancement) and is easier to estimate because it is assumed to vary slowly as a function of time. Although direct estimation of $\eta(n)$ through linear regression can lead to good results, estimating $\eta(n)$ remains a difficult problem. In this paper we propose to apply a gradient adaptive approach to the problem of estimating $\eta(n)$.

3. THE MDF ALGORITHM

In this paper, we consider the special case of the multidelay block frequency domain (MDF) adaptive filter [3]. The MDF algorithm in matrix form is detailed here for the sake of clarity. Let N be the MDF block size, K be the number of blocks and \mathbf{F} denote the $2N \times 2N$ Fourier transform matrix, we denote the frequency-domain signals for frame ℓ as:

$$\underline{\mathbf{e}}(\ell) = \mathbf{F} [\mathbf{0}_{1 \times N}, e(\ell N), \dots, e(\ell N + N - 1)]^T \quad (5)$$

$$\underline{\mathbf{x}}_k(\ell) = \text{diag} \left\{ \mathbf{F} [x((\ell-k-1)N), \dots, x((\ell-k+1)N-1)]^T \right\} \quad (6)$$

$$\underline{\mathbf{X}}(\ell) = [\underline{\mathbf{x}}_0, \underline{\mathbf{x}}_1, \dots, \underline{\mathbf{x}}_{K-1}] \quad (7)$$

$$\underline{\mathbf{d}}(\ell) = \mathbf{F} [\mathbf{0}_{1 \times N}, d(\ell N), \dots, d(\ell N + N - 1)]^T \quad (8)$$

The MDF algorithm is then expressed in matrix form as:

$$\underline{\mathbf{e}}(\ell) = \underline{\mathbf{d}}(\ell) - \underline{\mathbf{y}}(\ell) \quad (9)$$

$$\underline{\hat{\mathbf{y}}}(\ell) = \mathbf{G}_1 \underline{\mathbf{X}}(\ell) \underline{\hat{\mathbf{h}}}(\ell) \quad (10)$$

$$\underline{\hat{\mathbf{h}}}(\ell + 1) = \underline{\hat{\mathbf{h}}}(\ell) + \mathbf{G}_2 \underline{\boldsymbol{\mu}}(\ell) \nabla \underline{\hat{\mathbf{h}}}(\ell) \quad (11)$$

$$\nabla \underline{\hat{\mathbf{h}}}(\ell) = \Phi_{\mathbf{xx}}^{-1}(\ell) \underline{\mathbf{X}}^H(\ell) \underline{\mathbf{e}}(\ell) \quad (12)$$

where $\Phi_{\mathbf{xx}}(\ell)$ is the diagonal normalisation matrix as computed in [3], \mathbf{G}_1 and \mathbf{G}_2 are the constraint matrices:

$$\mathbf{G}_1 = \mathbf{F} \begin{bmatrix} \mathbf{0}_{N \times N} & \mathbf{0}_{N \times N} \\ \mathbf{0}_{N \times N} & \mathbf{I}_{N \times N} \end{bmatrix} \mathbf{F}^{-1} \quad (13)$$

$$\tilde{\mathbf{G}}_2 = \mathbf{F} \begin{bmatrix} \mathbf{I}_{N \times N} & \mathbf{0}_{N \times N} \\ \mathbf{0}_{N \times N} & \mathbf{0}_{N \times N} \end{bmatrix} \mathbf{F}^{-1} \quad (14)$$

$$\mathbf{G}_2 = \text{diag} \left\{ \tilde{\mathbf{G}}_2, \tilde{\mathbf{G}}_2, \dots, \tilde{\mathbf{G}}_2 \right\} \quad (15)$$

and $\underline{\boldsymbol{\mu}}(\ell)$ is the $2KN \times 2KN$ diagonal learning rate matrix:

$$\underline{\boldsymbol{\mu}}(\ell) = \text{diag} \left\{ [\boldsymbol{\mu}^T(\ell), \boldsymbol{\mu}^T(\ell), \dots, \boldsymbol{\mu}^T(\ell)]^T \right\} \quad (16)$$

$$\boldsymbol{\mu}(\ell) = [\mu_0(\ell), \mu_1(\ell), \dots, \mu_{2N-1}(\ell)]^T \quad (17)$$

If $\mu_k(\ell) = \mu_o$, we have the standard MDF algorithm.

4. GRADIENT-ADAPTIVE LEARNING RATE

In [2], we proposed to use the frequency-dependent learning rate:

$$\mu_k(\ell) = \min \left(\eta(\ell) \frac{|\underline{\hat{\mathbf{y}}}_k(\ell)|^2}{|\underline{\mathbf{e}}_k(\ell)|^2}, \mu_0 \right) \quad (18)$$

where $\eta(\ell)$ is a direct estimation of the normalised misalignment exploiting the non-stationarity of the signals and using linear regression between the power spectra of the estimated echo and the output signal. The motivation behind this formulation is that it factors the residual echo estimation into a

slowly-evolving (but unfortunately difficult to estimate) normalised misalignment $\eta(\ell)$ and a rapidly-evolving (but easy to estimate) term far-end term $|\underline{\mathbf{y}}_k(\ell)|^2$. The learning rate can thus react quickly to double-talk even if the estimation of the residual echo (through the misalignment estimate) requires a longer time period. A remaining problem with that approach is that $\eta(\ell)$ is difficult to estimate and the algorithm does not know whether its estimate of $\eta(\ell)$ is too low or too high. In this sense, the $\eta(\ell)$ update in [2] is an open-loop estimate.

4.1. Adaptation algorithm

In this paper we bypass the difficulty of estimating $\eta(\ell)$ directly and instead propose a closed-loop gradient adaptive estimation of $\eta(\ell)$. The parameter $\eta(\ell)$ is no longer an estimate of the normalised misalignment, but is instead adapted in closed-loop in such a way as to maximise convergence of the adaptive filter. As with other gradient-adaptive methods [5, 4] we compute the derivative of the mean square error $E(\ell) = \underline{\mathbf{e}}^H(\ell)\underline{\mathbf{e}}(\ell)$, this time with respect to $\eta(\ell-1)$, using the chain derivation rule:

$$\frac{\partial E(\ell)}{\partial \eta(\ell-1)} = \frac{\partial E(\ell)}{\partial \underline{\mathbf{e}}(\ell)} \frac{\partial \underline{\mathbf{e}}(\ell)}{\partial \underline{\hat{\mathbf{y}}}_k(\ell)} \frac{\partial \underline{\hat{\mathbf{y}}}_k(\ell)}{\partial \underline{\hat{\mathbf{h}}}(\ell)} \frac{\partial \underline{\hat{\mathbf{h}}}(\ell)}{\partial \mu(\ell-1)} \frac{\partial \mu(\ell-1)}{\partial \eta(\ell-1)} \quad (19)$$

$$= -\underline{\mathbf{e}}^H(\ell) \mathbf{G}_1 \underline{\mathbf{X}}(\ell) \mathbf{G}_2 \Phi_{\mathbf{xx}}^{-1}(\ell) \underline{\mathbf{X}}^H(\ell) \underline{\mathbf{e}}(\ell) \frac{|\underline{\hat{\mathbf{y}}}_k(\ell)|^2}{|\underline{\mathbf{e}}_k(\ell)|^2} \quad (20)$$

$$\approx -\frac{|\underline{\hat{\mathbf{y}}}_k(\ell)|^2}{|\underline{\mathbf{e}}_k(\ell)|^2} \nabla \underline{\hat{\mathbf{h}}}^H(\ell) \Phi_{\mathbf{xx}} \nabla \underline{\hat{\mathbf{h}}}(\ell) \quad (21)$$

We propose to use a filtered version of the gradient with a multiplicative update, similar to the general approach in [6], and to drop the normalisation term $\Phi_{\mathbf{xx}}$. This results in the new gradient-based learning rate adaptation rule:

$$\eta(\ell+1) = \eta(\ell) \exp \left[\rho \frac{|\underline{\hat{\mathbf{y}}}_k(\ell)|^2 \Re \{ \underline{\psi}^H(\ell) \nabla \underline{\hat{\mathbf{h}}}(\ell) \}}{|\underline{\mathbf{e}}_k(\ell)|^2 \sum_k \Re \{ \psi_k^*(\ell) \nabla \hat{h}_k(\ell) \}} \right] \quad (22)$$

$$\underline{\psi}(\ell+1) = \alpha \underline{\psi}(\ell) + \nabla \underline{\hat{\mathbf{h}}}(\ell) \quad (23)$$

where ρ is the step size and α controls the gradient smoothing (typically $\rho = 1$, $\alpha = 0.9$).

An intuitive interpretation for the update equation in (22) is that when the learning rate is too high (because $\eta(\ell)$ is too high), the gradient oscillates, causing $\eta(\ell)$ to decrease. On the other hand, when the learning rate is too low (perhaps because the echo path has changed) the gradient keeps pointing in the same direction and $\eta(\ell)$ increases. Also, because of the $\frac{|\underline{\hat{\mathbf{y}}}_k(\ell)|^2}{|\underline{\mathbf{e}}_k(\ell)|^2}$ factor, we can apply (22) at each frame regardless

of double-talk or even when no far-end speech is present (in which case $|\underline{\hat{\mathbf{y}}}_k(\ell)|^2 = 0$ anyway).

The last aspect that needs to be addressed is the initial condition. When the filter is initialised, all the weights are set to zero, which means that $\underline{\hat{\mathbf{y}}}_k(\ell) = \mathbf{0}$ and no adaptation can take place in (18) and (22). In order to bootstrap the adaptation process, the learning rate $\mu_k(\ell)$ is set to a fixed constant (we use $\mu_k(\ell) = 0.25$) for a short time equal to twice the filter length (only non-zero portions of signal $x(n)$ are taken into account). This procedure is only necessary when the filter is initialised and is not required in case of echo path change.

4.2. Algorithm behaviour

Here we show that the adaptive learning rate described above is able to deal with both double-talk and echo path change without explicit modelling. From (18) we can see that when double-talk occurs, the denominator $|\underline{\mathbf{e}}_k(\ell)|^2$ rapidly increases, causing an instantaneous decrease in the learning rate that lasts only as long as the double-talk period lasts. In the case of background noise, the learning rate depends on both the presence of an echo signal as well as the misalignment estimate. As the filter misalignment becomes smaller, the learning rate also becomes smaller.

One major difficulty involved in double-talk detection is the need to distinguish between double-talk and echo path change since both cause a sudden increase in the filter error signal. This distinction is made implicitly in the gradient-based adaptation of $\eta(\ell)$. When the echo path changes, $\nabla \underline{\hat{\mathbf{h}}}(\ell)$ starts pointing steadily in the same direction, thus significantly increasing $\eta(\ell)$, which is a clear sign that the filter is no longer properly adapted.

In gradient adaptive methods [5, 6], the implicit assumption is that both the near-end and the far-end signals are nearly stationary. We have shown that the GNGD algorithm [4] only requires the near-end signal to be nearly stationary. In our new proposed method, both signals can be non-stationary, which is a requirement for double-talk robustness.

5. RESULTS AND DISCUSSION

The proposed system is evaluated in an acoustic echo cancellation context with background noise, double-talk and a change in the echo path (16 seconds into a 32-second recording). The two impulse responses used are 1024-sample long and measured from real recordings in a small office room with both the microphone and the loudspeaker resting on a desk.

The proposed algorithm is compared to our previous work [2], to the normalised cross-correlation (NCC) method [1] and to a baseline with no double-talk detection (no DTD). The optimal threshold found for the NCC algorithm was 0.35 with a learning rate $\mu = 0.25$. It was found that choosing $\mu_{max} = 0.75$ as the upper bound on the learning rate gave good results for the proposed algorithm. In practise, μ_{max} has little impact

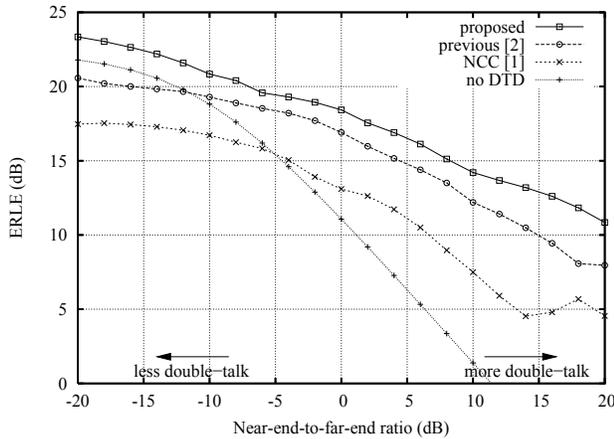


Fig. 2. Steady-state ERLE (first two seconds of adaptation not considered) as a function of the near-end to far-end ratio.

on the algorithm because the gradient-based adaptation would compensate for a higher value of μ_{max} by reducing $\eta(\ell)$. The GNGD algorithm [4] is not included in the evaluation because it is not intended for applications where the near-field signal is highly non-stationary and it was not possible to obtain better results than the baseline.

Fig. 2 shows the average steady-state (the first 2 seconds of adaptation are not considered) ERLE for the test data with different ratios of near-end signal and echo. Clearly, the proposed algorithm performs better than both our previous work (2 dB average improvement) and the NCC algorithm (6 dB average improvement). The perceptual quality of the output speech signal is also evaluated by comparing it to the near field signal $v(n)$ using the Perceptual Evaluation of Speech Quality (PESQ) ITU-T recommendations P.862 [7]. The perceptual quality of the speech shown in Fig. 3 is evaluated based on the entire file, including the adaptation time. It is again clear that the proposed algorithm performs better than all other algorithms. It is worth noting that the reason why the results in Fig. 2 improve with double-talk (unlike in Fig. 2) is that the signal of interest is the double-talk $v(n)$, so the higher the double-talk the less (relative) echo in the input signal.

6. CONCLUSION

We have demonstrated a novel method for adjusting the learning rate of frequency-domain adaptive filters based on a gradient adaptive estimation of the current misalignment. The proposed method performs better than a double-talk detector and previous work using direct estimation of the misalignment. In addition, the closed-loop gradient adaptive estimation of $\eta(\ell)$ makes the algorithm conceptually simple and means that there are very few important parameters to be tuned. Although the proposed algorithm is presented in the context of the MDF al-

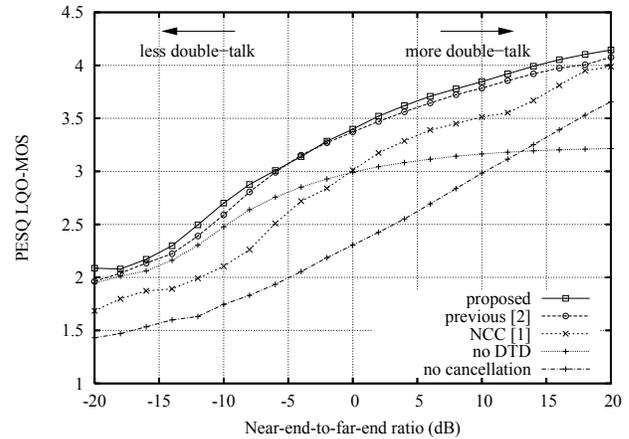


Fig. 3. PESQ objective listening quality measure (LQO-MOS) as a function of the near-end to far-end ratio.

gorithm, we believe future work could apply it to other adaptive filtering algorithms, including the NLMS algorithm.

7. REFERENCES

- [1] J. Benesty, D.R. Morgan, and J.H. Cho, "A new class of doubletalk detectors based on cross-correlation," *IEEE Trans. on Speech and Audio Processing*, vol. 8, no. 2, pp. 168–172, 2000.
- [2] J.-M. Valin, "On adjusting the learning rate in frequency domain echo cancellation with double-talk," *To appear in IEEE Trans. on Audio, Speech and Language Processing*, http://people.xiph.org/~jmv/papers/valin_tasl2006.pdf.
- [3] J.-S. Soo and K.K. Pang, "Multidelay block frequency domain adaptive filter," *IEEE Trans. ASSP*, vol. 38, no. 2, pp. 373–376, 1990.
- [4] D.P. Mandic, "A generalized normalized gradient descent algorithm," *IEEE Signal Processing Letters*, vol. 11, no. 2, pp. 115–118, 2004.
- [5] V.J. Mathews and Z. Xie, "A stochastic gradient adaptive filter with gradient adaptive step size," *IEEE Trans. on Signal Processing*, vol. 41, pp. 2075–2087, 1993.
- [6] W.-P. Ang and B. Farhang-Boroujeny, "A new class of gradient adaptive step-size LMS algorithms," *IEEE Trans. on Signal Processing*, vol. 49, no. 4, pp. 805–810, 2001.
- [7] ITU-T, *P.862: Perceptual evaluation of speech quality (PESQ): An objective method for end-to-end speech quality assessment of narrow-band telephone networks and speech codecs*, 2001.