Katholieke Universiteit Leuven

Departement Elektrotechniek

ESAT-SISTA/TR 05-258

Robust and Efficient Implementation of the PEM-AFROW Algorithm for Acoustic Feedback Cancellation[1]

Geert Rombouts[2] [3], Toon van Waterschoot[2], and Marc Moonen[2]

December 2005
Published in *Journal of the Audio Engineering Society*, vol. 55, no. 11, Nov. 2007, pp. 955–966

[1]This report is available by anonymous ftp from *ftp.esat.kuleuven.be* in the directory *pub/sista/vanwaterschoot/reports/05-258.pdf*

[2]K.U.Leuven, Dept. of Electrical Engineering (ESAT), Research group SCD(SISTA), Kasteelpark Arenberg 10, B-3001 Leuven, Belgium, Tel. +32 16 321927, Fax +32 16 321970, WWW: *http://www.esat.kuleuven.be/sista-cosic-docarch*. E-mail: *toon.vanwaterschoot@esat.kuleuven.be*.

[3]This research work was carried out at the ESAT laboratory of the Katholieke Universiteit Leuven, in the frame of the Belgian Programme on Interuniversity Attraction Poles, initiated by the Belgian Federal Science Policy Office IUAP P5/22 (“Dynamical Systems and Control: Computation, Identification and Modelling”) and P5/11 (“Mobile Multimedia Communication Systems and Networks”), the Concerted Research Action GOA-AMBioRICS and IWT project 040803: “SMS4PA-II: Sound Management System for Public Address Systems”. The scientific responsibility is assumed by its authors.

Robust and Efficient Implementation of the PEM–AFROW Algorithm for Acoustic Feedback Cancellation*

GEERT ROMBOUTS, TOON VAN WATERSCHOOT, *AES Student Member,* **and MARC MOONEN**
(geert.rombouts@esat.kuleuven.be) (toon.vanwaterschoot@esat.kuleuven.be) (marc.moonen@esat.kuleuven.be)

Katholieke Universiteit Leuven ESAT/SCD, 3001 Heverlee, Belgium

The PEM–AFROW algorithm has been developed to model a room impulse response under closed-loop conditions, such as in microphone–amplifier–loudspeaker systems, and is particularly suited for a speech signal input. If a controller is inserted in the loop that uses this model to remove the loudspeaker signal component from the microphone signal, an efficient feedback cancellation can be performed. Crucial implementation-oriented issues that have not been studied so far are addressed. In particular, adaptation control mechanisms are proposed for feedback cancellation based on PEM–AFROW, and efficient subband and frequency-domain implementations are provided for this algorithm.

0 INTRODUCTION

In closed-loop electroacoustic systems such as public-address (PA) systems or hearing aids, system instability due to acoustic feedback (the so-called Larsen effect [1], or howling) is a well-known phenomenon. It occurs when the system loop gain exceeds unity on frequencies ω_i where the loop phase equals $2n\pi$ rad, with n integer. An acoustic feedback canceler is a controller for this system, which attempts to maintain stability, resulting in "added gain." This performance measure is the difference between the maximum gain for a stable controlled system and that for a stable uncontrolled system.

The most important disadvantage of traditional acoustic feedback solutions [2], [3], mainly based on the insertion of notch filters centered on ω_i, is that they are reactive. (The instability first has to occur and be detected before the notch filter can be designed, and typically this detection takes up to 0.5 s.) Notch filters in the signal path also result in signal distortion. They can not remove the reverberation-like sound that occurs in marginally stable systems. On the other hand, a major advantage is that the notch-filtering technique can be made quite robust against fast changes in the acoustic environment. Typically a small displacement of a microphone will affect the center frequencies on which instability occurs only modestly. This means that by using wider notch filters, robustness can be achieved.

*Manuscript received 2006 November 16; revised 2007 August 16.

In [4], [5] an alternative acoustic feedback-canceling method is derived, which is based on the so-called PEM–AFROW algorithm (which is a prediction error method adaptive filter using matrix-row operations as a preprocessing step). This method belongs to a class of adaptive feedback cancellation algorithms, which are based on adaptive filtering and are similar to standard acoustic echo canceling [6], [10]. Fig. 1 provides a diagram where g is the gain and q^{-1} is the delay operator. Hence q^{-D} is a delay of D samples. The speech source signal $v(k)$ is assumed to be a time-variant autoregressive (TVAR) process, and is modeled as the output of an IIR filter

$$H(q, k) = \frac{1}{1 + a_1(k)q^{-1} + \cdots + a_P(k)q^{-P}} \tag{1}$$

driven by white noise $w(k)$. Hence $v(k) = H(q, k)w(k)$. Because of the closed signal loop, standard adaptive filters such as used in acoustic echo cancellation are found to provide biased solutions [10]. However, the PEM–

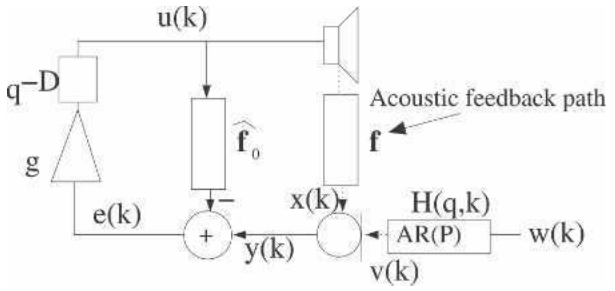

Fig. 1. Acoustic feedback cancellation.

AFROW algorithm allows for an unbiased joint identification of a linear model $\hat{\boldsymbol{f}}$ of the room impulse response (RIR) $\boldsymbol{f}$, and of the TVAR coefficients $\hat{\boldsymbol{a}}(k)$ of the denominator of $H(q, k)$, under closed-loop conditions (see also Fig. 1). Here $\hat{\boldsymbol{f}}$ is a vector that contains estimates of the coefficients of the finite impulse response (FIR) filter $\boldsymbol{f}$, modeling the room impulse response (acoustic feedback path)

$$\hat{\boldsymbol{f}}(k) = \begin{pmatrix} \hat{f}_0(k) \\ \hat{f}_1(k) \\ \vdots \\ \hat{f}_{L-1}(k) \end{pmatrix}$$

$\hat{\boldsymbol{a}}(k)$ is a vector,

$$\hat{\boldsymbol{a}}(k) = \begin{pmatrix} 1 \\ \hat{a}_1(k) \\ \vdots \\ \hat{a}_P(k) \end{pmatrix} \tag{2}$$

and q^{-D} is a delay of D taps. In [11] it is shown that for $D > P$, both $\boldsymbol{f}$ and $\boldsymbol{a}$ can be uniquely estimated. For a real speech input signal, apart from a short-term predictor with coefficients $\hat{\boldsymbol{a}}(k)$, in addition a long-term predictor is used, but for simplicity we will ignore this here (details are found in [5]). The estimated room impulse response $\hat{\boldsymbol{f}}$ can then be used in a so-called controller $\hat{\boldsymbol{f}}_0$, where the loudspeaker signal (far-end signal) $u(k)$, filtered by the model of the room impulse response (the controller), is subtracted from the microphone signal $y(k)$. The resulting signal $e(k)$ ideally has no remaining component stemming from the loudspeaker, but only contains the so-called near-end signal $v(k)$. The echo canceler terminology near-end and far-end will also be used in this paper to denote the corresponding signals (far-end for the loudspeaker signal and near-end for the speech signal), although the term "far end" does not strictly fit into the acoustic feedback cancellation context.

The PEM–AFROW algorithm [4], [5] operates in two steps (see Fig. 2). In a first step the estimate $\hat{\boldsymbol{f}}$ is assumed to be correct, and it is used to find a linear prediction error filter $\hat{\boldsymbol{a}}$ that whitens the speech source signal. In a second step this linear prediction error filter is in turn assumed to be correct, and it is used to perform a whitening of the speech source component in the input and desired signal for an adaptive filter, which then provides an update of the estimate for $\hat{\boldsymbol{f}}$. The linear prediction filter $\hat{\boldsymbol{a}}$ is—after convergence—equal to the inverse of the speech source model $H(q, k)$, and hence the adaptive filter $\hat{\boldsymbol{f}}$ then indeed operates on whitened versions of the signals $u(k)$ and $y(k)$ as its input and desired signal, respectively. It is easily shown that from the point of view of the adaptive filter $\hat{\boldsymbol{f}}$, this scheme, with a TVAR input and whitening filters as described, corresponds to a scheme without whitening filters, but driven by a white-noise signal. For a white-noise input, decorrelation between $u(k)$ and $v(k)$ is achieved when $D \geq 1$. Hence due to the decorrelation provided by the delay, the adaptive filter $\hat{\boldsymbol{f}}$ will converge to an unbiased solution.

The advantages of using such a scheme are obvious. It is proactive (no instability has to occur before feedback path identification can take place), and the room impulse response is continuously tracked. Second no (notch) filtering occurs in the signal path, which means that in case perfect identification is achieved, no signal distortion occurs. Finally this scheme also removes the "reverberating" sound, which occurs near the edge of stability.

In this paper we address crucial implementation-oriented issues that have not been studied so far. We will focus on some robustness issues (Section 1) that exist in the context of PEM–AFROW-based acoustic feedback cancellation. In particular, we will study the topics adaptation control (Section 1.1), use of prior knowledge (1.2), use of a foreground/background filter pair (Section 1.3), proactive notch filtering (Section 1.4), using nonlinearities (Section 1.5), and finally the inclusion of a high-pass filter

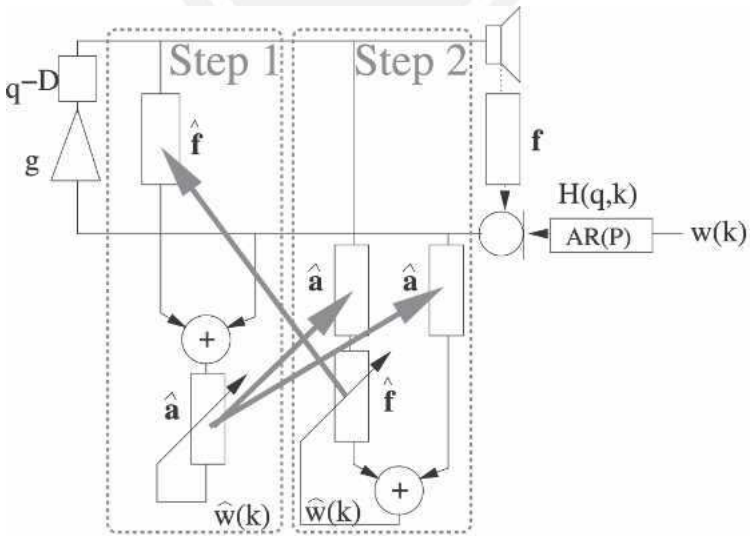

Fig. 2. PEM–AFROW identification algorithm.

(Section 1.6). Section 2 deals with subband and frequency-domain implementations of PEM–AFROW; Section 3 provides simulation results; and Section 4 gives the overall conclusions.

1 ROBUSTNESS

The controller in a PEM–AFROW-based acoustic feedback cancellation scheme stabilizes an unstable system (loudspeaker–room–microphone–amplifier). The controller $\hat{\boldsymbol{f}}_0$ should be matched to the acoustic feedback path $\boldsymbol{f}$ at all times to avoid instabilities. Fast changes in the environment can often not be tracked fast enough by the adaptive filter, and hence several robustness issues arise. These will be addressed in the following sections.

1.1 Adaptation Control

1.1.1 Problem Statement

In the context of standard acoustic echo cancellation, a near-end signal forms a disturbance (measurement noise) for the identification process performed by the adaptive filter. If the near-end signal $v(k)$ is correlated with the far-end signal $u(k)$, this will result in a bias in the estimate of the filter coefficients $\hat{\boldsymbol{f}}(k)$. If the signals are uncorrelated, variance on $\hat{\boldsymbol{f}}(k)$ will still result because of the measurement noise. When normalized least-mean-square (NLMS)–based algorithms are used (as is often the case because of their low complexity), an excess mismatch on $\hat{\boldsymbol{f}}(k)$ is introduced in the presence of near-end signals. Hence a double-talk detector is used to switch off adaptation during periods where near-end speech is present in order to avoid variance on the estimate of the RIR, and a far-end activity detector is used to switch off adaptation when the far-end signal energy is below the noise level.

In the acoustic feedback cancellation context (see Fig. 1) this far-end activity detector is also required (adaptation may only occur when a signal is present), but a double-talk detector is obviously irrelevant. The loudspeaker signal is correlated with the near-end signal due to the presence of the forward path gq^{-D}. Hence the feedback cancellation system operates continuously in double-talk mode. In order to be able to perform continuous adaptation when the far-end (and hence also near-end) signals are present, the PEM–AFROW algorithm decorrelates $u(k)$ and $v(k)$ by the whitening procedure combined with a delay in the forward path. This effectively avoids a bias in the estimate. Due to the continuous adaptation during signal presence, the convergence speed in a real-world scenario is fast enough to turn the audio volume level from $-\infty$ dB to 5–10 dB added gain in 5 s without the occurrence of howling. Note that in this case 0 dB is defined as the level at which in the system without a feedback controller howling would start to occur. The PEM–AFROW algorithm avoids the bias, but still a large variance may result. The amount of variance depends on the far-end to near-end ratio (FNR) between the two components in the microphone signal $y(k)$—the far-end component $x(k)$ and the near-end component $v(k)$. We define

$$\mathrm{FNR} = \frac{E\{x^2(k)\}}{E\{v^2(k)\}}$$

where $E\{\}$ is the expectation operator. The near-end component can be seen as measurement noise for the identification process, while the far-end component in the microphone signal is indeed the loudspeaker signal filtered by the room impulse response $\boldsymbol{f}$, and so represents the signal part, which is useful for the identification process. Hence for a large value of the FNR the variance will be low. For high gains in the forward path of an acoustic feedback application, the identification works better, which is of course useful since the Larsen effect occurs at high gains. On the other hand, in several scenarios the FNR may be low (see hereafter), and this will lead to a large variance on the estimate of the room impulse response, which may then result in instability of the system (Larśen effect).

1.1.2 Onset Detection

Instabilities primarily occur at speech onsets, or in general at sudden level increases of the near-end signal. At a speech onset [a sudden level increase of $v(k)$] the FNR is temporarily very small because the corresponding level increase in $x(k)$ is delayed by the forward path q^{-D} and the delay due to the room impulse response $\boldsymbol{f}$ itself (see Fig. 3). Hence if no control algorithm is incorporated in the system, onsets of speech may cause the estimate of the room impulse response to drift away from the correct so-

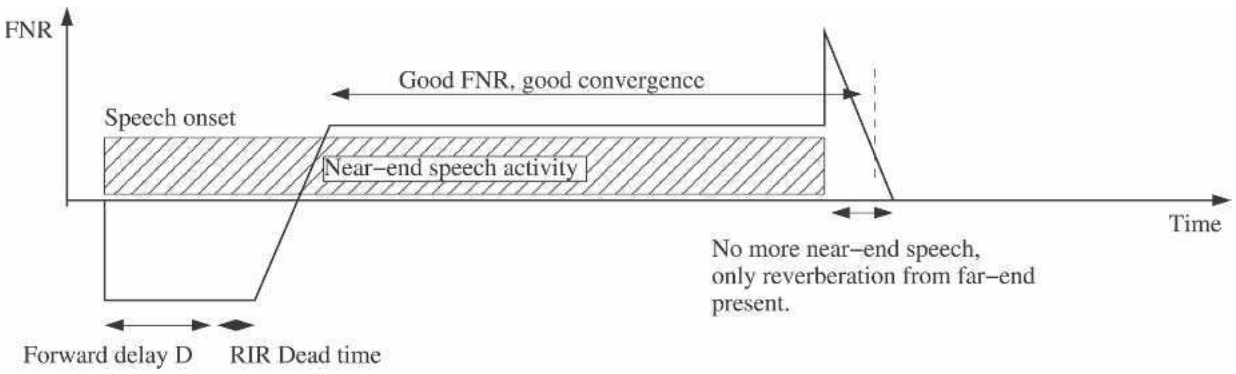

Fig. 3. Time line of far-end to near-end ratio (FNR) starting from near-end speech onset.

lution. This will often lead to instability of the whole system.

The control algorithm we propose is based on the observation that a level increase in $y(k)$ is always followed by a time interval with low FNR, until the level in $x(k)$ increases correspondingly. If an increase in $y(k)$ (such as a speech onset) is detected, adaptation is switched off for a predefined time interval Δ. It is important that the onset be detected instantaneously, and hence time averaging to obtain an energy estimate of $y(k)$ is impossible, because the lag introduced by the averaging would already allow the adaptive filter coefficients to drift substantially.

In order to achieve instantaneous detection, a Gaussian distribution of the (whitened) near-end signal $\hat{w}(k)$ is assumed, and the short-term variance $\sigma_w^2(k)$ of this distribution is estimated over an exponential window,

$$\sigma_w^2(k) = \lambda\sigma_w^2(k-1) + (1-\lambda)\hat{w}^2(k)$$

with $0 \ll \lambda < 1$ a forgetting factor. A speech onset is detected at time k_0 when an (instantaneous) value of $\hat{w}(k)$ occurs that is greater than $c\sigma_w(k)$, with c a constant.

If an onset is detected, adaptation is disabled from time k_0 to time $k_0 + \Delta$. A conservative setting for Δ would be the sum of the forward delay and the number of taps in the room impulse response. Sometimes most of the energy of the room impulse response is concentrated in the beginning of the room impulse response, and Δ can be smaller. It is important that after the delayΔ, the FNR assume a sufficiently large value.

1.2 Prior Knowledge of Feedback Path

The near-end speech source signal takes on the role of measurement noise for the identification process of the room impulse response by an adaptive filter. The PEM–AFROW algorithm contains an adaptive filter, which identifies the room impulse response, and a prewhitening stage, which whitens the measurement noise. The whitened measurement noise will result in a variance on the estimate of the room impulse response.

In order to provide robustness, the cost function can be modified to incorporate prior knowledge. This will lead to regularized versions [12], [13] of the adaptive filtering algorithm, in which the variance on the estimate of the room impulse response is reduced at the expense of a (hopefully) minor bias. We now introduce other regularized versions of the NLMS algorithm starting from a mean square error (MSE) cost function,

$$\begin{aligned} J(\hat{\boldsymbol{f}}(k), k) &= \varepsilon(\hat{w}^2(k)\} \\ &= \varepsilon\{[y_w(k) - \boldsymbol{u}_w^{\mathrm{T}}(k)\hat{\boldsymbol{f}}(k)]^2\} \end{aligned}$$

with $y_w(k) = \hat{\boldsymbol{a}}^{\mathrm{T}}(k)\boldsymbol{y}(k)$ the whitened microphone signal, where

$$\boldsymbol{y}(k) = \begin{pmatrix} y(k) \\ y(k-1) \\ \vdots \\ y(k-P+1) \end{pmatrix}$$

and

$$\boldsymbol{u}_w(k) = \begin{pmatrix} u_w(k) \\ u_w(k-1) \\ \vdots \\ u_w(k-L+1) \end{pmatrix}$$

with L the adaptive filter length and $u_w(k) = \hat{\boldsymbol{a}}(k)^{\mathrm{T}}\boldsymbol{u}(k)$ the whitened loudspeaker signal, where

$$\boldsymbol{u}(k) = \begin{pmatrix} u(k) \\ u(k-1) \\ \vdots \\ u(k-P+1) \end{pmatrix}.$$

From this minimization problem the well-known LMS or stochastic gradient descent algorithm is derived,

$$\begin{aligned} \hat{\boldsymbol{f}}(k) &= \hat{\boldsymbol{f}}(k-1) - \bar{\mu}\left.\frac{\partial J(\hat{\boldsymbol{f}}, k)}{\partial \hat{\boldsymbol{f}}}\right|_{\hat{f}=\hat{f}(k-1)} \\ &= \hat{\boldsymbol{f}}(k-1) + \bar{\mu}\,\varepsilon\{\boldsymbol{u}_w^{\mathrm{T}}(k)\hat{w}(k)\} \\ &\cong \hat{\boldsymbol{f}}(k-1) + \bar{\mu}\boldsymbol{u}_w(k)\hat{w}(k) \end{aligned}$$

where $\bar{\mu}$ is the step-size parameter. With a normalized step size

$$\bar{\mu}(k) = \frac{\mu}{\boldsymbol{u}_w^{\mathrm{T}}(k)\boldsymbol{u}_w(k)}$$

the update formula for NLMS is obtained as

$$\hat{\boldsymbol{f}}(k) = \hat{\boldsymbol{f}}(k-1) + \mu\,\frac{1}{\boldsymbol{u}_w^{\mathrm{T}}(k)\boldsymbol{u}_w(k)}\,\boldsymbol{u}_w(k)\hat{w}(k). \tag{3}$$

If a good estimate $\hat{\boldsymbol{f}}_p$ for $\hat{\boldsymbol{f}}$ has been obtained off line (for example, by a measurement of the impulse response at system setup, or by an off-line procedure based on recursive least squares (RLS), which does not exhibit excess mismatch as NLMS does, but which is computationally more expensive), then an alternative cost function can be defined,

$$J(\hat{\boldsymbol{f}}(k), k) = \varepsilon\{\hat{w}^2(k) + \beta\|\hat{\boldsymbol{f}}(k) - \hat{\boldsymbol{f}}_p\|^2\}.$$

Here β is a parameter that can be used to weight the importance of both terms. Along the lines of the derivation of the LMS and NLMS algorithms, regularized LMS and NLMS algorithms can be obtained. The regularized LMS update is

$$\begin{aligned} \hat{\boldsymbol{f}}(k) &= \hat{\boldsymbol{f}}(k-1) + \bar{\mu}(\boldsymbol{u}_w(k)\hat{w}(k) - \beta[\hat{\boldsymbol{f}}(k-1) - \hat{\boldsymbol{f}}_p)] \\ &= (1 - \beta\bar{\mu})\hat{\boldsymbol{f}}(k-1) + \bar{\mu}(\boldsymbol{u}_w(k)\hat{w}(k) + \beta\hat{\boldsymbol{f}}_p) \end{aligned}$$

from which the regularized NLMS update follows by choosing $\bar{\mu}$ to minimize the a posteriori error, and then introducing an additional step-size parameter,

$$\begin{aligned} \hat{\boldsymbol{f}}(k) &= \hat{\boldsymbol{f}}(k-1) + \mu\,\frac{\hat{w}(k)}{\bar{\boldsymbol{u}}_w^{\mathrm{T}}(k)\bar{\boldsymbol{u}}_w(k)}\,\bar{\boldsymbol{u}}_w(k) \\ \bar{\boldsymbol{u}}_w(k) &= \bar{\boldsymbol{u}}_w(k)\hat{w}(k) - \beta[\hat{\boldsymbol{f}}(k-1) - \hat{\boldsymbol{f}}_p]. \end{aligned} \tag{4}$$

Experiments show that using the update formula, Eq. (4), instead of Eq. (3) in the RIR identification part of the PEM–AFROW algorithm is especially useful at system startup. So at startup a relatively large β can be used, which can then be decreased gradually. A similar approach was used in [14].

1.3 Foreground/Background Filter

In order to minimize the variance on the estimate and (for NLMS algorithms) the excess mismatch due to the presence of the (whitened) speech source signal (measurement noise), a small step size should be chosen. While this provides robustness against measurement noise, it results in slow convergence. Under variations in the room impulse response, however, fast tracking is required in order to avoid instability.

In the acoustic echo cancellation context a twin filter structure is often used [15], [16] to overcome problems in double-talk situations. A similar idea can now be applied in acoustic feedback cancellation. Two adaptive filters estimate the RIR from the whitened microphone signal and the whitened loudspeaker signal. The so-called background adaptive filter $\hat{\boldsymbol{b}}(k)$ is operated with a larger step size (μ_b) than the foreground filter $\hat{f}(k)$ (step size μ). For example, $\mu_b = 10\mu$. The residual energies $\epsilon\{\hat{w}_b^2(k)\}$ of the whitened background filter residuals and $\epsilon\{\hat{w}^2(k)\}$ of the whitened foreground filter residuals are estimated as $\sigma_{t,b}^2(k)$ and $\sigma_t^2(k)$, respectively, over an exponential window,

$$\sigma_t^2(k) = \lambda_t \sigma_t^2(k-1) + (1-\lambda_t)\hat{w}^2(k)$$

$$\sigma_{t,b}^2(k) = \lambda_t \sigma_{t,b}^2(k-1) + (1-\lambda_t)\hat{w}_b{}^2(k)$$

and similarly the energy $\varepsilon\{y_w^2(k)\}$ of the whitened microphone signal $y_w(\mathrm{k}) = \hat{\boldsymbol{a}}(k)\boldsymbol{y}(k)$ is estimated as

$$\sigma_{t,y_w}^2(k) = \lambda_t \sigma_{t,y_w}^2(k-1) + (1-\lambda_t)\hat{y}_w{}^2(k).$$

Here $0 \ll \lambda_t < 1$ is a forgetting factor. At descrete time intervals $k = \Delta_t n$, with n integer, these energy estimates are compared, and the coefficients $\hat{\boldsymbol{b}}(k)$ are copied to $\hat{\boldsymbol{f}}(k)$ if the condition

$$\sigma_{t,b}^2(n\Delta_t) < \gamma_1 \sigma_t^2(n\Delta_t) \tag{5}$$

is met. (Optionally an additional condition

$$\sigma_{t,b}^2(n\Delta_t) < \gamma_2 \sigma_{t,y_w}^2(n\Delta_t) \tag{6}$$

can be imposed.) If on the other hand

$$\sigma_{t,b}^2(n\Delta_t) > \gamma_3 \sigma_t^2(n\Delta_t) \tag{7}$$

the coefficients $\hat{\boldsymbol{f}}(k)$ are copied to $\hat{\boldsymbol{b}}(k)$ to avoid divergence of $\hat{\boldsymbol{b}}(k)$. In these conditions, $\gamma_{1,2,3}$ are parameters with $0 < \gamma_{1,2,3} < 1$ and $\gamma_1 \leq \gamma_3$.

This procedure effectively combines the advantages of a large step size (fast convergence in case of a large change in the room impulse response) and of a small step size (low-variance steady-state behavior). This strategy can be applied in a time-domain NLMS approach, in a subband approach (see Section 2.1), as well as in a frequency-domain NLMS approach (Section 2.2).

1.4 Notch Filtering

1.4.1 Problem Statement

In the PEM–AFROW algorithm an unbiased estimate $\hat{\boldsymbol{f}}_0$ of the RIR is formed. The loudspeaker signal $u(k)$ is filtered by this RIR model, and the result is subtracted from the microphone signal $y(k)$, effectively canceling the transfer function $\boldsymbol{f}$ between loudspeaker and microphone. This means that essentially no signal distortion occurs in this setup (assuming a perfect estimate), more specifically that the signal is not attenuated in certain frequency regions, and that the gain level g, which is set on the amplifier, is effectively obtained, even when this means that the corresponding loop gain is larger than unity. An obvious issue is that if the tracking behavior of the PEM–AFROW algorithm is nonperfect, an unstable system may result as soon as the RIR changes.

Traditional approaches to feedback cancellation in public address applications use notch filtering. Disadvantages of this approach are the reactive character and the fact that signal distortion occurs. The gain is reduced in frequency regions where howling is likely to occur, which not only gives the signal an audible distortion, but which also results in the fact that the gain level that is set on the amplifier is not reached whenever the system becomes unstable.

PEM–AFROW was shown to be very robust against changes in the acoustic feedback path due to moving objects. These changes have the largest influence on the room impulse response (and require a good tracking behavior of the algorithm) when they occur in the vicinity of the loudspeaker or microphone, such as head movements of the speaker. Experiments were performed with a real-time system with loudspeaker and microphone spaced 2–3 m apart, in an office room environment, and it was shown that howling cannot be induced by fast head movements, by partially covering up the microphone or loudspeaker, or by moving large objects in the space between loudspeaker and microphone. Hence for these cases robustness is guaranteed by the PEM–AFROW algorithm.

On the other hand, when the microphone or loudspeaker position is changed, the RIR "shifts" over the time axis. This is a change of the RIR, which the PEM–AFROW algorithm cannot track fast enough. It is less robust than notch-filtering-based techniques when the loudspeaker or microphone itself is moved.

1.4.2 Proactive Notch Filtering

We have simulated an RIR [17] for a room of 10 by 5 by 4 m with reflection coefficients of the walls equal to 0.8, with a loudspeaker on position (3.21, 3.21, 2.5) m and a microphone on positions (1, 2, 1) m and (1, 1.95, 1) m, hence 50 mm apart. The RIR for the first microphone position and the difference between the RIRs for both microphones are plotted in Fig. 4. The difference has the same order of magnitude as the impulse responses them-

selves. It is obvious that a PEM–AFROW-based approach, which has to track the impulse response, may need a considerable amount of time before it has compensated for the difference between the impulse responses. On the other hand, the position of many of the peaks in the loop gain frequency spectrum shown in Fig. 5 does not change much, which indicates that a notch-filtering-based approach can be more robust against the microphone displacement. Depending on the room acoustics, only a limited number of modes are available that can give rise to instabilities, and so a small change in the distance between loudspeaker and microphone may lead to a small change in the center frequencies of the oscillations (unless a completely different mode becomes active), while it can lead to large changes in the room impulse response (see Fig. 4). Since notch filtering is robust against small changes in the center frequencies, it is desirable to combine the PEM–AFROW-provided estimates of the RIR (proactive) with the notch-filtering technique (robust).

The condition for instability is fulfilled when the loop gain $G_1(w) \geq 1$ on frequencies where the loop phase $\varphi_1 = n2\pi$, with n an integer. The frequencies w_i for which this condition holds, can be found from the estimate of the room impulse response, which is provided by PEM–AFROW, and the knowledge of the forward gain G, $G_l(w) = G$ abs (fft($\hat{f}$)), where abs(fft($\hat{f}$)) denotes the magnitudes of the discrete Fourier transform of $\hat{f}$.

The algorithm we have used to design the notch filters is given in Algorithm 1.

Algorithm 1 Notch filter design, starting from the PEM–AFROW room impulse response estimate:

1) For $i = 1$, NrActive,
   if $G_l(w_i) < T_{min}$, decrease depth on w_i
2) While maxima can be found,
   a) find $w_{max} = \text{argmax}G_l(w)$, where $w_{max} > T_{max}$
   b) if w_{max} close to existing notch w_F, $w_F = w_{max}$; increase depth on w_F
      elseif NrActive < max: new notch on w_{max}.

Algorithm I places notch filters on frequencies w_i, where $G_l(w_i) > T_{max}$ with threshold $T_{max} < 1$, independently of the phase characteristic (allthough this would be easy to incorporate), and it removes notch filters on frequencies where $G_l(w) < T_{min}$. The notch-filter design procedure is repeated regularly (a few times per second), and each time only a limited number of notch filters are designed. Then the PEM–AFROW algorithm is allowed to adapt to the new acoustic impulse response with the notch filters included. This results in an update of $G_l(w)$. The notch-filter depth on a certain frequency may then be increased or decreased, depending on the new value of $G_l(w)$.

The notch filters are IIR filters with two poles and two zeros each. From a numerical point of view it is important not to combine all of the notch filters into one IIR filter, but to apply all of the so-called biquads individually to the signal, as otherwise numerical instability may result if the poles are close together. A design procedure for notch filters is described in [18].

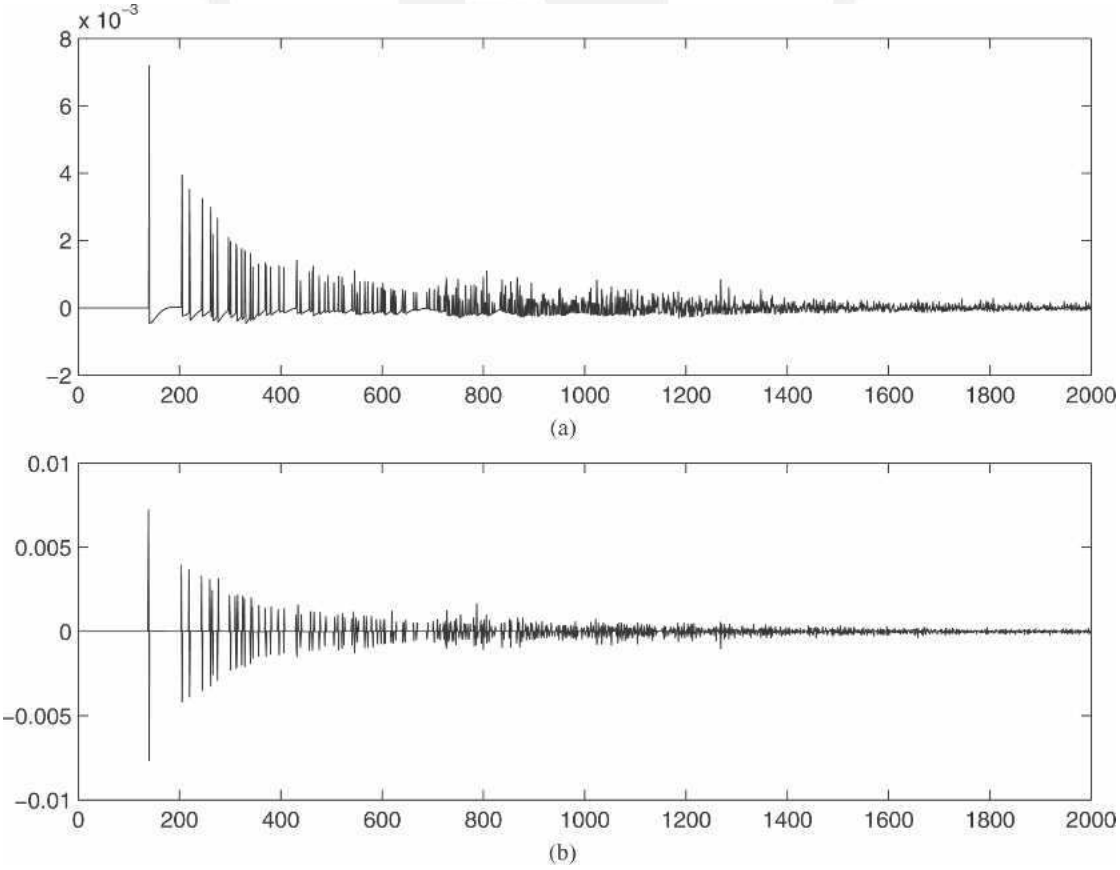

Fig. 4. RIR versus time index. (a) First microphone position. (b) RIR difference between two microphones.

At first sight it may seem advantageous not to include the notch filters in the feedback path (the part of the loop modeled by the PEM–AFROW algorithm), but instead place them in the forward path (the part containing the delay). That way the PEM–AFROW adaptive filter would not have to model the changing notch-filter characteristics. But on the other hand the FNR would become very low on the notch frequencies when the notch filters are placed in the forward path (the loudspeaker signal is attenuated while the near-end signal is not), and a large variance on the estimation of the loop gain would occur on the notch frequencies. Simulations show that this leads to wrong decisions on where to place (or remove) notch filters.

The notch-filtering-based acoustic feedback controller is now combined with a PEM–AFROW-based controller, as shown in Fig 6. The PEM–AFROW algorithm estimates the cascade of the notch filters with the RIR, and this estimate is then used in a controller $\hat{f}_0$, which cancels the remaining loudspeaker component in the microphone signal. The advantage compared to the PEM–AFROW–only algorithm [4], [5] is that the notch filters—when chosen appropriately wide—guarantee robustness for path changes (in exchange for some distortion), and the controller removes the residual reverberations like sounds.

If $T_{\max}$ is chosen smaller than 1, the preceding technique leads to a proactive system because new notch filters will be designed before the system becomes unstable at the corresponding frequencies. The system can be made robust to path changes by choosing an appropriate width for the notch filters. The parameters for the notch filters should be chosen to avoid audible distortion. It is important to note that due to the additional PEM–AFROW-based controller $\hat{f}_0$, which removes residual loudspeaker components from the microphone signal, the loop gain on the frequencies w_i is not limited to $T_{\max}$, as in traditional notch-filtering-only-based techniques.

1.5 Nonlinearities

If large signal levels are applied to the loudspeaker/amplifier (for example, in case the Larssen effect occurs), the loudspeaker or amplifier may exhibit nonlinear behavior (clipping). A PEM–AFROW-based acoustic feedback cancellation scheme provides a linear estimate of the transfer function for the system with input signal $u(k)$ and output signal $x(k)$ (see Fig. 1). This system comprises the amplifier (which is not explicitly drawn in the figure, but in practice is placed just before the loudspeaker), the loud-

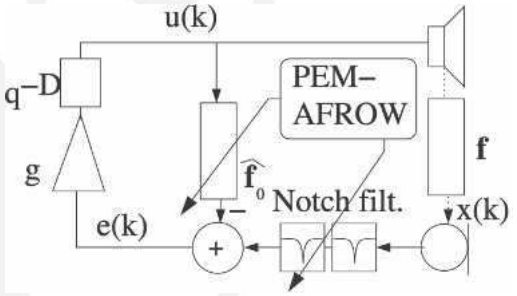

Fig. 6. Combined system with proactive notch filtering and controller $\hat{f}_0$.

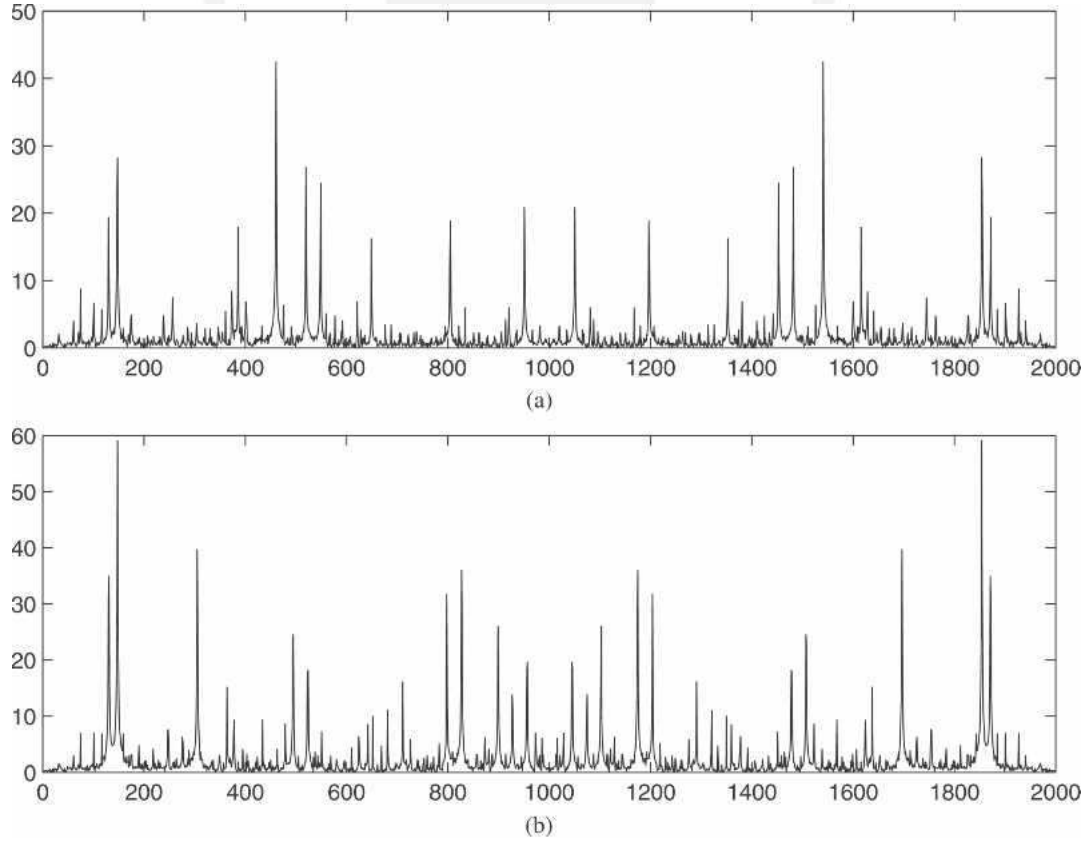

Fig. 5. Loop gain versus frequency index for two slightly different microphone positions (100 mm apart).

speaker, the acoustic feedback path, and the microphone. (We will refer to this system as the room impulse response.) The linear estimate, which is used in the controller $\hat{f}_0$, will not be a correct model when nonlinearities due to loudspeaker or amplifier clipping occur in the system. Because of the mismatch between the (linear) controller and the (nonlinear) system, instabilities (howling) will occur.

A simple and efficient method to avoid clipping in the RIR is by limiting the maximum amplitude of $u(k)$, as shown in Fig. 7, that is, by adding a limiter or clipper on the left-hand side of the controller, before its input.

In several applications nonlinear operations are required in the signal path (such as dynamics processing). These operations should be applied before the input of both the controller and the RIR system, such that the identification algorithm does not see these nonlinearities. Note that also the gain control g should preferably be placed in the left-hand part of the scheme, as shown in the figure, such that variations and nonlinearities (such as dynamics compression) in g do not have to be tracked by the adaptive algorithm.

1.6 High-Pass Filtering

Often the low cutoff frequency of the loudspeaker or amplifier is higher than that of the microphone. This means that for low-frequency near-end signals (such as generated by displacing the microphone stand), the FNR may become very small. Due to the large variance in the estimation resulting from this, the system may become unstable. A simple solution is to use a high-pass filter on the microphone signal with a cutoff frequency (depending on the loudspeaker characteristics) of, for example, 100 Hz. The influence on the speech quality will be very small, while low-frequency near-end signals are suppressed.

2 SUBBAND AND FREQUENCY-DOMAIN IMPLEMENTATIONS

In public-address applications typically long impulse responses are involved, and high sample frequencies are desired. This leads to a high computational complexity in time-domain-based algorithms. We propose two different frequency-domain implementations. The first one (Section 2.1) is based on a subband implementation of the time-domain algorithm. A filter bank is used to split the time-domain signals into independent frequency bands, and processing is done independently in each of these bands.

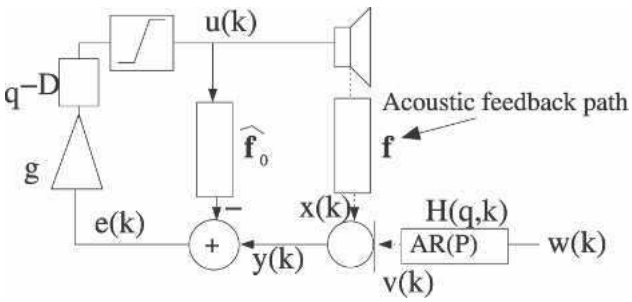

Fig. 7. Limiter or clipper should be added to avoid clipping in loudspeaker (amplifier).

The advantage (apart from a reduction in complexity) will be that the filter length can be chosen independently for each of the frequency bands. The downside of this approach is the latency introduced by the filter-bank operation.

The second frequency-domain implementation reduces this latency since it is based on the PBFDAF filter (Section 2.2).

2.1 Oversampled DFT-Modulated Filter-Bank Implementation

We first propose an implementation based on a subband implementation of the NLMS algorithm. A DFT-modulated perfect reconstruction filter bank [19] divides the signals in subbands, and then the PEM–AFROW-based algorithm is applied in each subband. This is advantageous from a performance and complexity point of view.

It is well known that the performance of the NLMS-based algorithms, which run in each band, is better when the eigenvalue spread of their input signal is smaller. Since for a typical speech signal the band-filtered signals in the subbands are whiter than the full-band signal, the eigenvalue spread is indeed smaller.

The complexity is also reduced compared to time-domain processing. In an M-band filter bank with Nfold downsampling the filter length of the adaptive filters in each band is roughly N times smaller than for the corresponding time-domain setup, and the adaptive filters also operate at an N times lower sampling rate. The adaptive filters in only $M/2 + 1$ of the bands have to be computed, since the other bands are complex conjugated. So, neglecting the complexity of the filter-bank operations, the complexity in the subband scheme is $M/2N^2$ smaller compared to the corresponding time-domain scheme.

In order to avoid aliasing, it is important to use an oversampled filter bank (which means that $N < M$). The filter-bank prototype filter will never have ideal properties (it is not a brick wall filter), and hence for $M = N$, frequency components above the Nyquist frequency would be present in the subband signals before the downsampling. The "perfect reconstruction" property of the filter bank would indeed compensate for the aliasing even in that case, but only if no processing occurs in the subbands. This is not the case in our scheme. So N should be chosen small enough such that no aliasing occurs even with non-ideal bandpass filters in the filter bank.

In the subband processing context, complexity can be reduced even further thanks to the properties of the physical environment. It can be observed that for a typical room (or hall) the absorption of the high-frequency components is much higher than for the low-frequency components. We propose to estimate the maximum required number of filter taps for each subband individually in a calibration stage. It should be noted that calibration should be performed in an empty room, and that (due to absorption in the clothing) the impulse response of rooms is always shortened when people (the audience) enter the room. (Note that a similar approach could also be used in acoustic echo canceling scenarios.)

Fig. 8 shows the setup of this system. Straightforward implementation is shown in Fig. 8(a), where a PEM–AFROW algorithm is performed in each of the subbands independently. In Fig. 8(b) we show how one analysis filter bank can be removed by directly reusing the subband signals as inputs for the adaptive filters. Since the synthesis/analysis filter-bank pair in Fig. 8(b) introduces a delay $q^{-D_{fb}}$, and the subband filters have to estimate the subband path between loudspeaker and microphone, this delay must also be included before the subband filters.

Since the individual PEM–AFROW algorithms in each of the subbands run at an N times lower sampling frequency, their parameters should be adjusted compared to the full-band version of the algorithm. This means that the autoregressive order, the frame length, and the forward delay are all N times smaller compared to the full-band version.

The resulting algorithm, implemented in C++ without any specific optimization effort, runs in real time on a Pentium Mobile 1.7 GHz with a Linux operating system, and Jack (www.jackaudio.org) as the sound device interface, for a 16-band filter-bank, with two-fold oversampling, and with 600 filter taps per band.

2.2 Frequency-Domain Implementation

A major disadvantage of the filter-bank approach in Section 2.1 is the delay introduced by the filter banks. This delay becomes larger when the number of subbands (and hence the complexity reduction) is larger. Allthough a certain amount of delay is required in the PEM–AFROW algorithm, the sum of a large filter-bank delay and the delay due to the AD/DA buffering may become prohibitive for certain applications. In [20] the authors use a partitioned-block frequency-domain-adaptive filtering algorithm (PBFDAF, also known as the MDF algorithm [21]) to overcome this disadvantage.

It should be noted though that in the PEM–AFROW algorithm the input vector of the adaptive filter $\hat{\boldsymbol{f}}$ has to be recalculated each time when the estimated AR coefficients $\hat{\boldsymbol{a}}$ change. This amounts to a matrix–vector multiplication with dimension $L \times (P + 1)$, which occurs at least every speech frame, but which may occur more often due to the nondeterministic nature of the long-term predictor [5]. In an implementation with PBFDAF this results in the FFTs of all of the input vector partitions having to be recalculated each time this occurs. The computational complexity for such a PBFDAF implementation will hence become comparable to the complexity of an FDAF implementation with the FFT size equal to $2L$ and a shift size equal to the frame length.

In practical applications it is often advisable to use the PBFDAF or FDAF algorithm to achieve a low latency, instead of the subband approach of Section 2.1 PBFDAF and FDAF combined with PEM–AFROW both have comparable complexity, but the FDAF algorithm results in a more simple implementation.

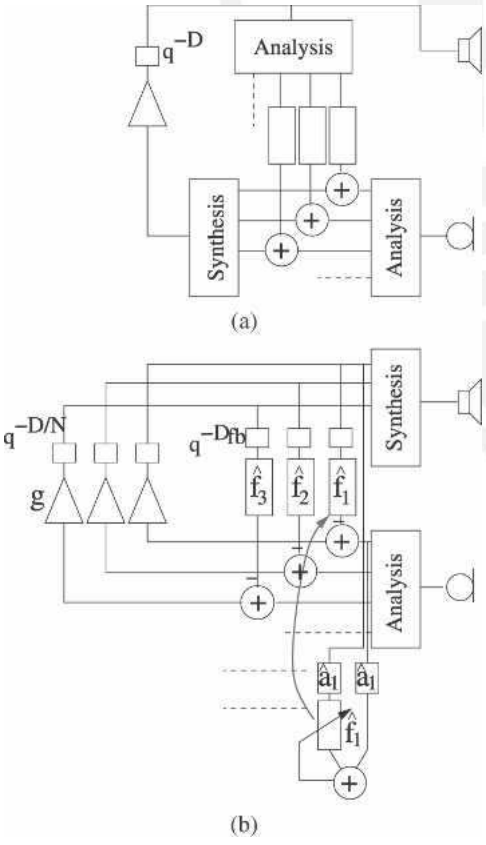

Fig. 8. Filter-bank implementation. (a) Straightforward implementation. (b) Implementation with one analysis bank less (identification part shown only for one subband).

3 SIMULATIONS

In this section we will first focus on the individual benefits achieved by onset detection, the use of both a foreground and a background filter, and proactive notch filtering, and then we will review a combined system which incorporates all of the techniques described.

3.1 Onset Detection

The simulation environment consists of an artificial RIR of 2000 taps (obtained with the mirror image method [17]) for a room of 10 by 5 by 4 m, the reflection coefficients of the walls are 0.8, with a loudspeaker on position (3.21, 3.21, 2.5) m and a microphone on position (1, 2, 1) m. A sample frequency of 16 kHz is used. A speech signal contains several different types of onsets, and hence in order to obtain reproducible results, we use a white noise input signal with intermittently a level of −46 dB during 1 s and a level of 0 dB during 1 s. The adaptive filter in the PEM–AFROW algorithm has a (sufficient) order of $L = 2000$ taps. The frame length for energy comparison (Section 1.1.2) is 20 ms, parameter c defined in Section 1.1.2 is 5, and the adaptation inhibition time after detection of an onset (Section 1.1.2) is $\Delta = 2L$. The added gain is set to 5 dB.

For the algorithm without onset detection, audible (transient) instabilities are present at the onsets of the louder signal blocks, even with small step sizes for the adaptive NLMS algorithm ($\mu = 0.005$), which corresponds to an unacceptably low convergence and tracking speed. With

$\mu = 0.05$ these transients even take more than a second to disappear. With the onset detection switched on, the step-size can be as large as $\mu = 0.2$ (which obviously results in an acceptable tracking speed), without audible instabilities.

These results are confirmed when a speech input signal is used instead of a white noise signal. The values for the maximum of μ when onset detection is off may vary, because of the different types of onsets in a speech signal. But the value of $\mu = 0.2$ as a maximum value when onset detection is switched on, is consistent with the white noise experiment.

3.2 Foreground/Background Filter

The same simulation setup is used, without the onset detection algorithm. The step size for the foreground filter is chosen 10 times smaller than that for the background filter. The values for the parameters mentioned in Section 1.3 are $\gamma_1 = \gamma_3 = 0.86$ and the condition in Eq. (6) is omitted. We evaluate the behavior for the previously mentioned input signal. Large step sizes can be used for the background filter (for example, $\mu = 0.9$), which allows for very fast convergence, while a small variance can be obtained due to the behavior of the foreground filter with the smaller step size. When μ is chosen not too large, then for high loop gains (high FNR), such as with an added gain of 10 dB, the foreground/background filter structure does not exhibit instabilities at onsets, because the slow foreground filter "bounds" the divergence of the faster background filter. For low loop gains (as with negative added gains) or higher values of μ, onset detection should be used to prevent coefficients from drifting from the optimal solution.

The foreground/background filter structure is not only important for robustness related to signal properties, but also for robustness against system changes (that is, changes of the RIR). The simulations use the room and geometry setup as detailed in Section 1.4.2. In the simulations, abrupt switches between two artificial impulse responses are performed, one corresponding to microphone position (1, 2, 1) m and one corresponding to microphone position (1, 1.95, 1) m. In this case a real speech input signal is used. As a reference, we consider a system with a single adaptive filter with $\mu = 0.09$, resulting in the same variance as the foreground filter in the two–filter structure. In the two–filter setup, for an added gain of 5 dB, no audible transient is present, while a clear transient is audible in the reference system. For an added gain of 10 dB, a short transient is audible during 0.5 s. For an added gain of 13 dB, a transient instability is audible during 1 s.

All these results were also verified with a real-time (16-kHz) implementation with a real speech input signal, in a real acoustic environment (office room of 7 by 3.5 by 3.6 m, reverberation time (RT_{60}) 120 ms, distance between loudspeaker and microphone 2 m).

3.3 Proactive Notch Filtering

Real-time experiments at a 16-kHz sampling rate in a real environment were performed with added gains of about 5 dB, using only the notch-filtering section (proactive, designed based on the PEM–AFROW estimate), and without the additional controller $\hat{f}_0$. It is known that this is a typical stable added gain for notch-filtering-based feedback suppression. Fast movements (about 1 Hz) of the microphone over a distance of 300 mm, at 1-m distance from the loudspeaker, do not result in instability for this proactive notch-filtering approach, while for the PEM–AFROW-based feedback suppression algorithm in [4], [5], instability does indeed sometimes occur under these conditions. Note, however, that for constant microphone–loudspeaker distances much higher added gains can be achieved with the approach in [4], [5] (up to 15 dB), even with large changes in the RIR due to movements of, for example, people in the vicinity of the loudspeaker or microphone.

In addition to this experiment, and in order to provide reproducible results, simulations were performed with a room geometry as detailed in Section 1.4.2 for a combined system as shown in Fig. 6. Sudden changes between the two RIRs are applied (which are worse conditions than the 1-Hz movements in the real-environment experiment described here). Then it is evaluated for which maximum gain level no instability or audible artifacts occur upon changes of the impulse responses. The results are given in Table 1. The step size for the NLMS filter in the PEM–AFROW algorithm is 0.05. For comparison, the stable added gain when only the notch-filtering section is used is 5 dB.

3.4 High-Pass Filter

Though difficult to quantify numerically, real-time experiments show a huge improvement in stability in case the microphone is displaced when a high-pass filter is available (see Section 1.6).

3.5 Combined Real-Time System

The combined frequency-domain system is implemented for use in a real-time environment, in C++ on a Linux PC (Pentium M, 1.7 MHz). A graphical user interface is available for setting the parameters of different building blocks, and for switching on and off the different modules. The filter length is 2048 taps, and the sampling frequency is 16 kHz. In this way an RIR of 128 ms can be modeled. The system is set up in a real environment (office room of 7 by 3.5 by 3.6 m, reverberation time (RT_{60}) 120 ms, distance between loudspeaker and microphone 2 m), and maintains stability far better than the PEM–AFROW-based system without the robustness add-ons.

As a general result from real-time experiments we can state that the most important elements for increased ro-

Table 1. Simulation results for different notch-filter parameters.

Bandwidth*	Depth	Added Gain
3 dB BW = 1/20	12 dB	10 dB
3 dB BW = 1/10	12 dB	13.9 dB
1 dB BW = 1/5	3 dB	12.9 dB

* Bandwidths are normalized (1 = Nyquist frequency).

bustness are the high-pass filter, the foreground/background filter structure, and onset detection. Proactive notch filtering can increase robustness somewhat, but has the disadvantage of introducing signal distortion.

4 CONCLUSIONS

In this paper we have shown a number of procedures to increase the robustness and reduce the complexity of a PEM–AFROW-based feedback cancellation system. Concerning robustness, we have shown how onset detection avoids instability at onsets of the near-end signal. We have derived an NLMS-type algorithm which allows for prior knowledge of the RIR to be incorporated in the adaptation rule, which is particularly useful at system startup. Then a twin-filter structure (adopted from acoustic echo cancellation) was shown to be also effective in acoustic feedback cancellation. On the one hand room impulse response changes can be tracked quickly while on the other hand a small variance on the estimate of the RIR can be maintained during steady-state operation. A high-pass filter is included to avoid high variance on the low-frequency part of the RIR estimate. Finally the combination of notch filtering and PEM–AFROW was shown to provide a more robust solution to echo path changes if some distortion is allowable. Concerning computational complexity we have demonstrated how the PEM–AFROW algorithm can be implemented in subbands, where extra complexity reduction can be obtained due to the physical characteristics of reverberation in a typical environment, and as an alternative a low-latency frequency-domain implementation can be used.

5 ACKNOWLEDGMENT

This research work was carried out at the ESAT laboratory of the Katholieke Universiteit Leuven, in the frame of the Belgian Programme on Interuniversity Attraction Poles, initiated by the Belgian Federal Science Policy Office IUAP P5/22 (Dynamical Systems and Control: Computation, Identification and Modelling) and P5/11 (Mobile Multimedia Communication Systems and Networks), the Concerted Research Action GOA-AMBioRICS, IWT project 040803, SMS4PA-II (Sound Management System for Public Address Systems). The scientific responsibility is assumed by its authors.

6 REFERENCES

[1] G. Nicollini, Y. Mazoyer, C. Crippa, S. Pernici, and P. Confalonieri, "A CMOS Fully Integrated Antilarsen System for Digital Telephones," *IEEE J. Solid-State Circ.,* pp. 605–609 (1993).

[2] R. Porayath and D. J. Mapes-Riordan, "Acoustic Feedback Elimination Using Adaptive Notch Filter Algorithm," U.S. patent 5,999,631 (1999 Dec.).

[3] M. P. Lewis, T. J. Tucker, and D. M. Oster, "Method and Apparatus for Adaptive Audio Resonant Frequency Filtering," U.S. patent 5,245,665 (1993 Sept.).

[4] G. Rombouts, T. van Waterschoot, K. Struyve, and M. Moonen, "Acoustic Feedback Suppression for Long Acoustic Paths Using a Nonstationary Source Model," in *Proc. Eusipco 2005* (Antalya, Turkey, 2005 Sept.).

[5] G. Rombouts, T. van Waterschoot, K. Struyve, and M. Moonen, "Acoustic Feedback Suppression for Long Acoustic Paths Using a Nonstationary Source Model," *IEEE Trans. Signal Process.,* vol. 54, pp. 3426–3434 (2006 Sept.).

[6] F. G. van der Meulen, S. Kamerling, and K. Janse, "A New Way of Suppressing Acoustic Feedback," presented at the 104th Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts),* vol. 46, pp. 582, 583 (1988 June), preprint 4735.

[7] C. P. Janse and H. J. W. Belt, "Sound Reinforcement System Having an Echo Suppressor and Loudspeaker Beamformer," patent WIPO WO 03/010996 A3 (2002 June 24).

[8] C. P. Janse and P. A. A. Timmermans, "Signal Amplifier System with Improved Echo Cancellation," U.S. patent 5,748,751 (1998 May).

[9] M. G. Siqueira and A. Alwan, "Steady-State Analysis of Continuous Adaptation in Acoustic Feedback Reduction Systems for Hearing-Aids," *IEEE Trans. Speech Audio Process.,* vol. 8, pp. 443–453 (2000 July).

[10] J. Hellgren and U. Forssell, "Bias of Feedback Cancellation Algorithms in Hearing Aids Based on Direct Closed Loop Identification," *IEEE Trans. Speech Audio Process.,* vol. 9, pp. 906–913 (2001 Nov.).

[11] A. Spriet, I. Proudler, M. Moonen, and J. Wouters, "Adaptive Feedback Cancellation in Hearing Aids with Linear Prediction of the Desired Signal," *IEEE Trans. Signal Process.,* vol. 53, pp. 3749–3763 (2005 Oct.).

[12] T. van Waterschoot, G. Rombouts, and M. Moonen, "Optimally Regularized Adaptive Filtering Algorithms for Room Acoustic Signal Enhancement," ESAT-SISTA Tech. Rep. TR 05-184, Katholieke Universiteit Leuven, Belgium, 2005; accepted for publication in Elsevier Signal Processing.

[13] T. van Waterschoot, G. Rombouts, and M. Moonen, "MSE Optimal Regularization of APA and NLMS Algorithms in Room Acoustic Applications," in *Proc. 2006 Int. Workshop on Acoustic Echo Noise Control* (IWAENC '06), (Paris, France, 2006 Sept.).

[14] J. M. Kates, "Constrained Adaptation for Feedback Cancellation in Hearing Aids," *J. Acoust. Soc. Am.,* vol. 106, pp. 1010–1019 (1999 Aug.).

[15] K. Ochiai, T. Araseki, and T. Ogihara, "Echo Canceler with Two Echo Path Models," *IEEE Trans. Commun.,* vol. 25, pp. 589–595 (1977 June).

[16] J. Prado and E. Moulines, "Frequency Domain Adaptive Filtering with Applications to Acoustic Echo Cancellation," *Ann. Telecommun.,* vol. 49, pp. 414–428 (1994).

[17] J. Allen and D. Berkley, "Image Method for Efficiently Simulating Small-Room Acoustics," *J Acoust. Soc. Am.,* vol. 65, pp. 943–950 (1979 Apr.).

[18] T. van Waterschoot and M. Moonen, "A Pole–Zero Placement Technique for Designing Second-Order IIR

Parametric Equalizer Filters," *IEEE Trans. Audio, Speech, Language Process.,* vol. 15, pp. 2561–2565 (2007 Nov.).

[19] P. P. Vaidyanathan, *Multirate Systems and Filter Banks* (Prentice-Hall, Englewood Cliffs, NJ, 1993).

[20] A. Spriet, G. Rombouts, M. Moonen, and J. Wouters, "Adaptive Feedback Cancellation in Hearing Aids," *J. Franklin Inst.,* vol. 343, pp. 545–573 (2006 Sept.).

[21] J. Soo and K. Pang, "Multidelay Block Frequency Domain Adaptive Filter," *IEEE Trans. Acoust., Speech, Signal Process.,* vol. 38, pp. 373–376 (1990 Feb.).

THE AUTHORS



G. Rombouts



T. van Waterschoot



M. Moonen

Geert Rombouts was born in Turnhout, Belgium, in 1973. He received an electrical engineering degree in 1997 and a Ph.D. degree in applied sciences in 2003, both from the Katholieke Universiteit Leuven, Belgium.

At present he is a postdoctoral researcher in the Electrical Engineering Department of the Katholieke Universiteit Leuven, working in the field of audio signal processing.

●

Toon van Waterschoot was born in Lier, Belgium, on June 11, 1979. He received a master's degree from Katholieke Universiteit Leuven (K. U. Leuven), Belgium, in 2001, where he is currently working toward a Ph.D. degree.

In 2002 he spent a year as a teaching assistant with the Hogere Zeevaartschool Antwerpen, Belgium, after which he returned to K. U. Leuven to start a Ph.D. research project on signal processing applied to audio and acoustics. His current research is focused on the development of adaptive filtering algorithms for acoustic echo and feedback cancellation.

●

Marc Moonen received an electrical engineering degree and a Ph.D. degree in applied sciences from Katholieke Universiteit Leuven, Belgium, in 1986 and 1990, respectively.

Since 2004 he has been a full professor in the Electrical Engineering Department of Katholieke Universiteit Leuven, where he is heading a research team working in the area of numerical algorithms and signal processing for digital communications, wireless communications, DSL, and audio signal processing.

Dr. Moonen is a fellow of the IEEE. He received the 1994 K. U. Leuven Research Council Award and the 1997 Alcatel Bell (Belgium) Award (with Piet Vandaele).