Digital Signal Processing 33 (2014) 71–82

Contents lists available at ScienceDirect

Digital Signal Processing

www.elsevier.com/locate/dsp

Self-tuning adaptive frequency tracker

Michał Meller [1]

Department of Automatic Control, Gdańsk University of Technology, Faculty of Electronics, Telecommunications and Computer Science, Narutowicza 11/12, 80-233 Gdańsk, Poland

ARTICLE INFO

Article history:
Available online 19 June 2014

Keywords:
Automatic tuning
Adaptive notch filtering
Adaptive signal processing
Frequency tracking

ABSTRACT

An automatic gain tuning algorithm is proposed for a recently introduced adaptive notch filter. Theoretical analysis and simulations show that, under Gaussian random-walk type assumptions, the proposed extension is capable of adjusting adaptation gains of the filter so as to minimize the mean-squared frequency tracking error without prior knowledge of the true frequency trajectory. A simplified one degree of freedom version of the filter, recommended for practical applications, is proposed as well.

1. Introduction

Adaptive notch filters (ANFs) are devices capable of suppressing or enhancing nonstationary narrowband signals buried in noise. The vector of state variables of ANFs often includes, or allows one to determine, instantaneous frequency of the tracked narrowband signal. Such filters, e.g. [1–4], can be used to estimate time-varying frequency of the signal of interest.

A typical ANF employs one or more user-dependent adaptation gains which must be judiciously tuned to optimize the filters' performance. Since tuning of the filter is a time-consuming process, it is desirable to provide some means of automatic adjustment of these gains, e.g. in a form of a supervisory self-optimization layer. Such a combination could equip an ANF with a capability to deliver nearly-optimal performance, despite possible variations of tracking conditions, e.g. the signal to noise ratio or a rate at which the frequency changes.

Adaptation gains are most often adjusted with an aim of optimizing signal tracking performance. Popularity of this approach results primarily from two facts. Firstly, the underlying application of an ANF may require tight signal tracking. Such is the case of, among others, filtering power signal from electrocardiogram (ECG) recordings [5], tracking harmonic currents in power applications [6–8] or active control of narrowband acoustic noise [9]. Secondly, in case of signal tracking it is rather simple to distinguish the poorly performing ANF from the one which performs well. This can be done, without any prior knowledge of the true values of the narrowband signal, by evaluating prediction errors yielded by the filter [10]. Due to this relative simplicity of measuring the filter's performance, several self-tuning ANF solutions were proposed in signal processing literature, see e.g. [10,11].

The problem of optimizing frequency tracking performance of an ANF has received substantially less attention – existing results are generally limited to evaluations of theoretically achievable mean-square errors and experimental comparisons of various notch filtering approaches [12–16]. This stems from the fact it is actually quite difficult to tell how well the filter tracks the signal's frequency without a prior knowledge of its true trajectory – it occurs that settings which minimize signal tracking errors are usually different from those which minimize frequency tracking errors [17]. For this reason, existing solutions often underperform in terms of frequency tracking when they are run in a real-world scenario.

The problems outlined above are encountered in other approaches to frequency estimation as well. For instance, estimators such as the short time Fourier transform (STFT) or ESPRIT [18] require one to choose local analysis window length, which should be adjusted to the signal's characteristics. Similarly, trackers based on the extended Kalman filter (EKF) are known to be highly accurate [19–21,11], but difficult to tune – even though some very good rules were pointed in [19].

Recently, new results on ANF application to frequency tracking were established in [22]. It was shown that one can quantify frequency tracking accuracy of an ANF with a special auxiliary predictor. A parallel adaptive frequency tracker was also proposed in [22]. It employs a bank of notch filters and uses the proposed predictor to select the filter which offers the most accurate tracking. The new scheme outperformed several existing algorithms, including a recent self-tuning one.

However, the parallel scheme introduced in [22] is not free of drawbacks. First, it has rather high computational complexity, because it involves a bank of notch filters. Second, its accuracy

E-mail address: michal.meller@eti.pg.gda.pl.
[1] Fax: +48 58 347 15 55.

http://dx.doi.org/10.1016/j.dsp.2014.06.004
1051-2004/

grossly depends on how representative the setting of the filters composing the bank are. If there is no filter 'matched' to tracking conditions, the parallel scheme will yield degraded performance. This means that the filterbank must be quite large, which further increases computational complexity.

The contribution of the paper is threefold. First, a novel automatic gain adjustment mechanism is proposed for a recently introduced adaptive notch filter. Second, theoretical analysis of the algorithm's behavior is performed. It is shown that, under Gaussian assumptions, the resulting estimator is locally convergent in mean to the optimal tracker, even in the case of unknown and time-varying signal characteristics. This is, most likely, the first solution with such a capability – as it will be argued later, the parallel tracker from [22] cannot provide performance guarantees similar to those of the sequential approach. Third, the analysis reveals that the preliminary version of the self-adjustment mechanism suffers from slow convergence. Therefore, a normalized variant of the algorithm is proposed, which is free of this drawback. A simplified, one degree of freedom scheme is proposed as well. Finally, rich simulation material confirms good properties of the improved solutions and validates results of convergence analysis.

The remaining part of paper is organized as follows. Section 2 formulates the problem and reviews fundamental findings from [22]. Section 3 develops the proposed approach. Convergence of the proposed solution is analyzed in Section 4. Section 5 discusses extensions and compares the sequential approach with the parallel one. Section 6 presents simulation results. Section 7 concludes.

2. Problem formulation and related results

Consider the problem of estimating an unknown slowly time-varying frequency $\omega(t)$ of a narrowband complex sinusoid (cisoid) using noisy measurements

$$y(t) = s(t) + v(t), \tag{1}$$

where $t = 0,\ 1, \dots$ denotes discrete time, $v(t)$ is a wideband measurement noise,

$$s(t) = a(t)e^{j\sum_{\tau=1}^{t}\omega(\tau)} \tag{2}$$

is a nonstationary complex sinusoid with instantaneous frequency $\omega(t)$ and

$$a(t) = m(t)e^{j\phi_0}, \tag{3}$$

where $m(t)$ is a real valued slowly time varying amplitude and ϕ_0 denotes the initial phase.

Tracking of $\omega(t)$ may be accomplished e.g. using the following ANF algorithm, introduced in [17][2]

$$\begin{aligned}
\hat{f}(t) &= e^{j[\hat{\omega}(t-1)+\hat{\alpha}(t-1)]}\hat{f}(t-1) \\
\varepsilon(t) &= y(t) - \hat{a}(t-1)\hat{f}(t) \\
\hat{a}(t) &= \hat{a}(t-1) + \theta_3 \hat{f}^*(t)\varepsilon(t) \\
\hat{\alpha}(t) &= \hat{\alpha}(t-1) + \theta_1\delta(t) \\
\hat{\omega}(t) &= \hat{\omega}(t-1) + \hat{\alpha}(t-1) + \theta_2\delta(t) \\
\delta(t) &= \mathrm{Im}\left[\frac{\varepsilon(t)}{\hat{a}(t-1)\hat{f}(t)}\right] \\
\hat{s}(t) &= \hat{a}(t)\hat{f}(t),
\end{aligned} \tag{4}$$

where $\hat{f}(t)$ is a phase term, $\epsilon(t)$ is the prediction error, $*$ denotes complex conjugation, the quantities $\hat{a}(t)$, $\hat{s}(t)$, $\hat{\omega}(t)$ and $\hat{\alpha}(t)$ are the estimates of the signal's complex 'amplitude', instantaneous value, instantaneous frequency and instantaneous frequency rate $[\alpha(t-1) = \omega(t) - \omega(t-1)]$, respectively. Adaptation laws of (4) take a form typical to adaptive signal processing [23] and frequency estimation [2,19,20]: new values of variables are obtained by adding an update term, proportional to quantities which may be interpreted as errors ($\epsilon(t)$ and $\delta(t)$). The parameters $\theta_1 > 0$, $\theta_2 > 0$, $\theta_3 > 0$, $\theta_1 \ll \theta_2 \ll \theta_3$, are small adaptation gains, determining the rates of amplitude adaptation, frequency adaptation and frequency rate adaptation, respectively.

Tracking of frequency rate is quite uncommon among adaptive notch filters and deserves more justification. Many real world narrowband signals exhibit approximately piecewise linear frequency modulation. Such is the case, among others, of radar signals, acoustic engine noise encountered in active noise control systems or vibration generated by rotating machinery during transient states. In these cases estimation of both frequency and frequency rate can improve tracking accuracy considerably – even by an order of magnitude, see e.g. [17]. Finally, observe that setting $\hat{\alpha}(0) = 0$ and $\theta_1 = 0$ effectively disables the frequency rate tracking feature.

As shown in [17], the algorithm (4) has very good statistical properties. Under the following assumptions:

(A1) Instantaneous frequency drifts according to the 2-nd order random walk (also called integrated random walk)

$$\begin{aligned}
\omega(t) &= \omega(t-1) + \alpha(t-1) \\
\alpha(t) &= \alpha(t-1) + w(t),
\end{aligned} \tag{5}$$

where $\{w(t)\}$ forms a stationary zero-mean Gaussian white noise sequence, $w \sim \mathcal{N}(0, \sigma_w^2)$,

(A2) The sequence $\{v(t)\}$, independent of $\{w(t)\}$, is a circular complex Gaussian white noise, $v \sim \mathcal{CN}(0, \sigma_v^2)$,

(A3) The magnitude of the narrowband signal is constant, $|s(t)| \equiv a_0$,

a proper choice of the gains θ_1, θ_2, θ_3 can turn the algorithm (4) into a statistically efficient frequency tracker. In such case steady state mean squared frequency tracking errors, $\Delta\hat{\omega}(t) = \omega(t) - \hat{\omega}(t)$, will reach the fundamental lower bound (called posterior or Bayesian Cramér–Rao bound [24]) which limits mean-squared tracking performance of any estimation scheme. Although closed-form expressions for the optimal values of gains θ_1, θ_2, θ_3 do not exist, it can be shown that they depend only on the following 'normalized' measure of signal nonstationarity [17]

$$\kappa = \frac{a_0^2\sigma_w^2}{\sigma_v^2}. \tag{6}$$

Note that, from a practical point of view, statistical efficiency of (4) is of somewhat questionable value – the values of the parameters σ_v^2, σ_w^2 and a_0^2 are unlikely to be known, which makes tuning (4) a challenging task.

In [22] a novel way of evaluating frequency tracking performance of ANFs was proposed. It was shown that the settings which minimized the mean-squared frequency tracking error also minimized the mean-squared value of the following auxiliary sequence

$$\xi(t) = \frac{1 - 2q^{-1} + q^{-2}}{\theta_2 + (\theta_1 - \theta_2)q^{-1}}\hat{\omega}(t), \tag{7}$$

where q^{-1} denotes the backward shift operator, $q^{-1}u(t) = u(t-1)$. Note that $\xi(t)$ may be obtained without any prior knowledge of the true values of $\omega(t)$. This makes (7) particularly useful for on-line optimization.

[2] The original formulation used the symbols γ_α, γ_ω, μ, rather than θ_1, θ_2, and θ_3.