Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Sound and Vibration 286 (2005) 37–54
JOURNAL OF SOUND AND VIBRATION
www.elsevier.com/locate/jsvi

A comparison of two adaptive algorithms for the control of active engine mounts

A.J. Hillis*, A.J.L. Harrison, D.P. Stoten

Department of Mechanical Engineering, University of Bristol, Queens Building, University Walk, Bristol BS8 1TR, UK

Received 19 March 2004; received in revised form 27 September 2004; accepted 27 September 2004
Available online 22 December 2004

Abstract

This paper describes work conducted in order to control automotive active engine mounts, consisting of a conventional passive mount and an internal electromagnetic actuator. Active engine mounts seek to cancel the oscillatory forces generated by the rotation of out-of-balance masses within the engine. The actuator generates a force dependent on a control signal from an algorithm implemented with a real-time DSP. The filtered-x least-mean-square (FXLMS) adaptive filter is used as a benchmark for comparison with a new implementation of the error-driven minimal controller synthesis (Er-MCSI) adaptive controller. Both algorithms are applied to an active mount fitted to a saloon car equipped with a four-cylinder turbo-diesel engine, and have no a priori knowledge of the system dynamics. The steady-state and transient performance of the two algorithms are compared and the relative merits of the two approaches are discussed. The Er-MCSI strategy offers significant computational advantages as it requires no cancellation path modelling. The Er-MCSI controller is found to perform in a fashion similar to the FXLMS filter—typically reducing chassis vibration by 50–90% under normal driving conditions.

1. Introduction

Automotive engine mounts are required to constrain motion of the engine resulting from low-frequency road inputs. It is also necessary for the chassis to be isolated from high-frequency engine vibrations resulting from the rotation of out-of-balance masses. The frequency range of the

*Corresponding author.
E-mail address: a.j.hillis@bristol.ac.uk (A.J. Hillis).

0022-460X/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jsv.2004.09.023

engine vibrations is such that they can excite components in contact with occupants, and can also generate significant acoustic noise. The constraint of engine motion requires a stiff, highly damped mount, whereas isolation requires a flexible mount with low damping. Hence, a passive solution (a tuned mass/spring/damper system) will always be a compromise.

Active engine mounts seek to overcome this compromise. The example under development at Bristol University in conjunction with Avon VMS combines a passive hydraulic mount with an internal electromagnetic actuator capable of transmitting additional forces to the chassis [1]. A diagram of the mount is shown in Fig. 1.

The passive components consist of a rubber cone spring, and fluid chambers and channels. The spring provides engine support and acts as a piston upon the main fluid chamber. Deflection of the spring forces fluid through the channels and into the compensation chamber. The effect of this is to provide a high level of damping at low frequencies (<20 Hz). The diaphragm partially bounds the main fluid chamber and determines the dynamic stiffness of the passive mount. The passive section of the mount is designed to have characteristics very similar to a standard passive hydraulic mount. The active components consist of a permanent magnet and a moving coil, behaving very similar to a loudspeaker. The coil acts upon the diaphragm and changes the pressure in the main fluid chamber and thus will transmit a force to the chassis that is dependent upon an actuator signal supplied by a control algorithm.

The chassis vibration is measured by an accelerometer and is considered to be nearly sinusoidal. An adaptive algorithm is desirable in order to generate the appropriate actuator signal to make the chassis stationary relative to the high-frequency engine vibrations, since the frequency of the vibrations and the dynamics of the system are time-varying. Additionally, it is desirable for the mount to be applicable to different vehicles with minimal adjustment; hence, no knowledge of

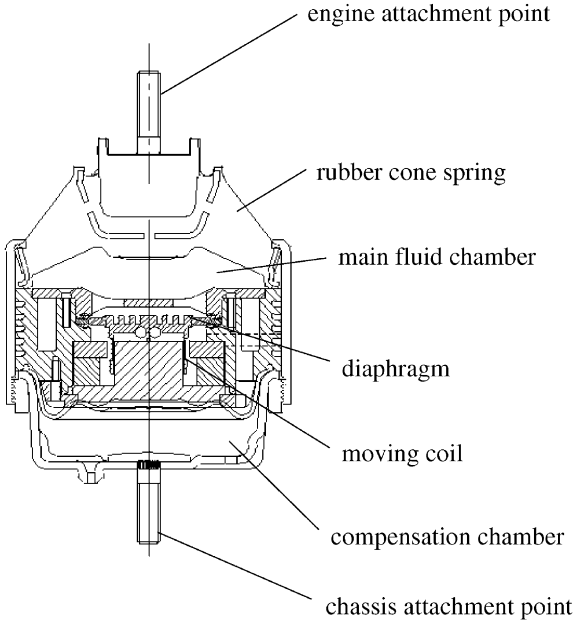

Fig. 1. Cross-section of active engine mount.

the system parameters can be assumed. Should the system parameters be known a priori, robust control or gain-scheduling schemes could be employed. The combination of unknown and time-varying system parameters, however, requires the use of a fully adaptive control strategy (i.e. the algorithms must operate in the absence of any parametric information).

Much work has been carried out in the past decade in the field of active noise and vibration cancellation (ANVC). Due largely to its simplicity, the most popular cancellation strategy for this type of application is the least-mean-square (LMS) adaptive filter, described in Section 2. A detailed description of the LMS algorithm and its applications may be found in Refs. [2,3]. The filtered-x LMS (FXLMS) version of the algorithm has been applied to automotive active engine mounts in a variety of forms, for example in Refs. [4–6].

This paper describes the implementation of the error-driven minimal controller synthesis (Er-MCSI) adaptive controller [7] with the active engine mount. The Er-MCSI controller is described in Section 3. The FXLMS algorithm is used as a benchmark for comparison. For this work, Er-MCSI- and FXLMS-based algorithms are applied to an active engine mount fitted to a saloon car equipped with a four-cylinder turbo-diesel engine. The performances of the two strategies are compared and the relative merits of the two approaches are discussed.

2. The LMS adaptive filter

The LMS algorithm was originally developed by Widrow and Hoff in 1960. Since then it has been widely applied in the field of active noise control. A detailed analysis of the properties of the algorithm may be found in Ref. [3].

The version of the algorithm used here is the narrow-band FXLMS for cancellation of periodic noise [2], which takes the form of an adaptive notch filter. Fig. 2 shows the FXLMS arrangement for the case where it is required to cancel a single frequency.

The engine vibration $v(n)$ is transmitted through the primary path dynamics $P(z)$, representing the passive component of the engine mount. This results in the primary vibration $d(n)$ at the chassis, which is to be cancelled. The cancelling signal $u(n)$ is modified by the secondary path

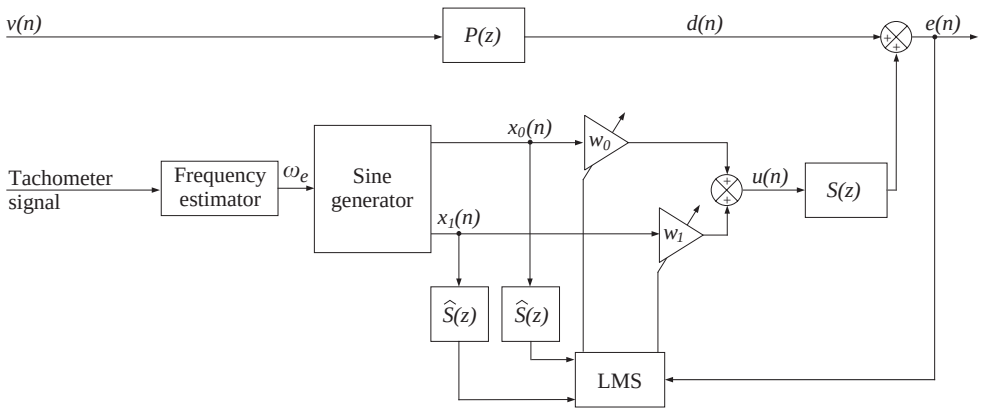

Fig. 2. Implementation of the FXLMS algorithm for cancelling periodic noise.

dynamics $S(z)$. For this application the secondary path represents the dynamics of the actuator, the vibration path to the feedback sensor and an anti-aliasing filter. The resultant chassis vibration is the measured error signal $e(n)$.

A synthesised reference signal $[x_0\ x_1]$ is generated, which contains only the frequencies that require cancellation. The filter will only respond to components of the error signal that are correlated with the reference input. The total measured engine vibration signal $v(n)$ could be used as a reference, but this signal will contain low-frequency components resulting from the overall movement of the vehicle (road inputs). The forces required to cancel these low-frequency components are too large for an active mount to generate; so it is undesirable for the filter to respond to them as saturation of the actuator would occur. For this application, a tachometer generates a square wave at the frequency of the crankshaft. This signal is used to estimate the engine frequency and to generate a sinusoidal reference signal at the appropriate frequency (e.g. twice the engine frequency to cancel the dominant vibration from a four-cylinder engine). The reference signal is defined by the equations:

$$\begin{aligned} x_0(n) &= \sin(\mathrm{i}\omega_e n\Delta), \\ x_1(n) &= \cos(\mathrm{i}\omega_e n\Delta), \end{aligned} \tag{1}$$

where ω_e is the estimated engine frequency, i is the harmonic order corresponding to the particular frequency to be cancelled (e.g. 2 to cancel the second-order vibration from a four-cylinder engine), and Δ is the fixed sampling interval.

For every frequency present in the reference signal, two adaptive weights are required for cancellation. The weights are updated by the equations:

$$\begin{aligned} w_0(n+1) &= \gamma w_0(n) - \mu x_0'(n)e(n), \\ w_1(n+1) &= \gamma w_1(n) - \mu x_1'(n)e(n), \end{aligned} \tag{2}$$

where μ is the step-size parameter—a positive real constant used to determine the convergence rate of the algorithm, and γ is a leakage factor applied to the tap weights to avoid the accumulation of numerical rounding errors. x_0' and x_1' are the reference signals x_0 and x_1 filtered by the secondary path estimate $\hat{S}(z)$, described by the length $L+1$ vector of time-domain impulse response values:

$$\hat{\mathbf{s}}(n) = [\hat{s}_0(n)\ \ \hat{s}_1(n)\ \ \hat{s}_2(n)\ \ \ldots\ \ \hat{s}_L(n)]^{\mathrm{T}}, \tag{3}$$

where L is chosen to be sufficiently large to capture the impulse response of the system dynamics at a given sampling rate.

Filtered versions of the reference signals must be used to compensate for the secondary path dynamics in order to achieve successful cancellation.

Finally, the filter output (control signal) is given by

$$u(n) = w_0(n)x_0(n) + w_1(n)x_1(n). \tag{4}$$

The FXLMS algorithm will converge, providing the difference in phase between $S(z)$ and $\hat{S}(z)$ is not greater than 90° [8]. $\hat{S}(z)$ is usually estimated off-line and then fixed in the real-time implementation of the filter. For this application, $S(z)$ is initially unknown, and is also time-varying, so it is necessary to perform simultaneous on-line identification and cancellation.

On-line identification is achieved using an additional LMS filter to model $S(z)$ [9]. The reference input to this filter is low-level broadband noise $b(n)$ which is added to the actuator signal, and the

error signal is the measured chassis vibration $e(n)$ minus the output from $\hat{S}(z)$. The use of a time-domain filter can cause problems with the on-line system identification, since the estimate $\hat{S}(z)$ can be corrupted by error signal content uncorrelated with $b(n)$ [9]. This can increase the convergence time of the cancellation filter and result in instability. Road inputs and uncancelled engine vibration components will be present within the measured error signal, thus a time-domain estimate would be corrupted. The robustness of the system identification for this application is substantially improved by implementing the fast-block LMS (FBLMS) algorithm, operating in the frequency domain [10]. The FBLMS algorithm is shown in Fig. 3.

The weight update equation for block number k encompassing the length L vector of samples from $n-L+1:n$ is defined as

$$\hat{\mathbf{s}}(k+1) = \gamma\hat{\mathbf{s}}(k) + \mathscr{F}^{-1}[\boldsymbol{\mu}_{\mathrm{ID}}(z_k)\mathbf{B}^{*}(z_k)\mathbf{E}_{\mathrm{ID}}(z_k)]_{1:L}, \tag{5}$$

where (*) denotes complex conjugation and $\mathscr{F}$ is the fast fourier transform (FFT). $\mathbf{B}(z_k)^{2L\times 2L}$ is a diagonal matrix with the diagonal terms consisting of

$$\mathrm{diag}[\mathbf{B}(z_k)] = \mathscr{F}\begin{bmatrix}\mathbf{b}(k-1)\\ \mathbf{b}(k)\end{bmatrix} \tag{6}$$

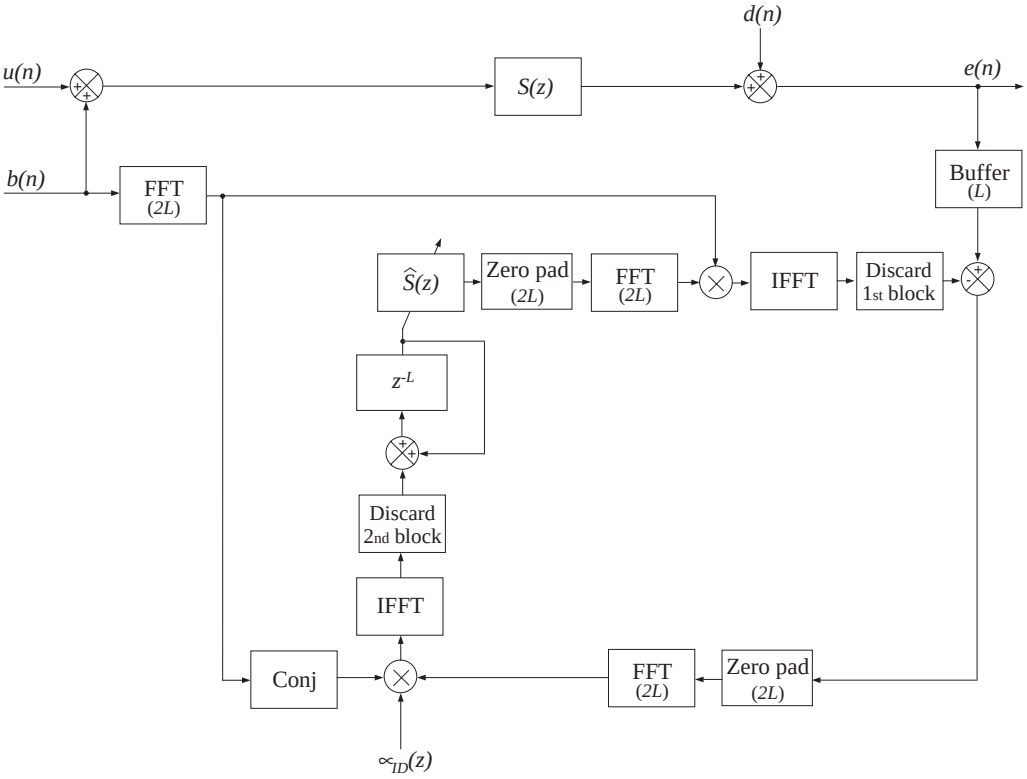

Fig. 3. Frequency-domain FBLMS algorithm for on-line secondary path estimation.

and the block error signal for the identification filter is

$$\mathbf{E}_{\mathrm{ID}}(z_k) = \mathscr{F}\left[\begin{array}{c}\mathbf{0}_{L\times 1}\\ \mathbf{e}(k) - \mathscr{F}^{-1}\left[\mathbf{B}(z_k)\mathscr{F}\left[\begin{array}{c}\hat{\mathbf{s}}(k)\\ \mathbf{0}_{L\times 1}\end{array}\right]\right]_{L+1:2L}\end{array}\right]. \tag{7}$$

The FBLMS algorithm is usually implemented in order to achieve equal convergence rates across the frequency spectrum. The step-size parameter becomes a function of the power spectral density (PSD) of the reference signal. For this application, however, it is a function of the PSD of the measured error signal, and is a $2L \times 2L$ diagonal matrix with the diagonal terms:

$$\mathrm{diag}[\boldsymbol{\mu}_{\mathrm{ID}}(z_k)] = \frac{\mu_0}{\mathbf{P}'_{ee}(z_k)}, \tag{8}$$

where $\mathbf{P}'_{ee}$ is the first-order low-pass filtered version of the length $2L$ PSD of the block error signal, given by

$$\mathbf{P}'_{ee}(z_k) = \tau\mathbf{P}_{ee}(z_k) + (1-\tau)\mathbf{P}_{ee}(z_{k-1}) \tag{9}$$

with block number k. The degree of filtering is determined by the constant τ.

Thus, a lower adaption rate is applied to the frequency bins containing signal power relating to the uncancelled engine vibration components and corruption of the estimate $\hat{S}(z)$ is reduced.

A filter length of 32 was used and the algorithm was implemented on a dSPACE DS1104 board sampling at a rate of 4 kHz.

3. The Er-MCSI adaptive controller

The minimal controller synthesis (MCS) algorithm [11] is an extension of the model reference adaptive controller (MRAC). The important distinction between MCS and MRAC is that MCS requires no a priori knowledge of the plant dynamics. The algorithm used is a derivation of the MCS controller known as error-driven MCS with integral action (Er-MCSI) [7], and is shown in Fig. 4 for a single-input/single-output (SISO) system. This algorithm is a form of proportional-plus-integral (P + I) adaption.

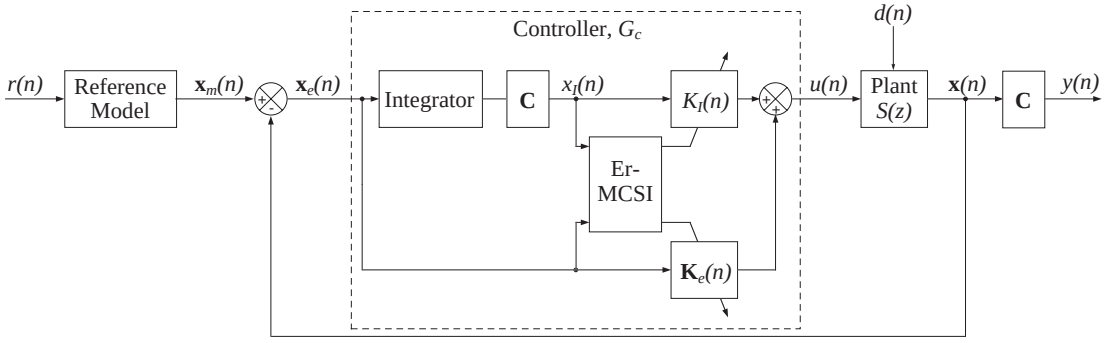

Fig. 4. Block diagram of the Er-MCSI adaptive controller.

The dynamics of the jth-order time-varying plant may be represented by the state equation

$$\mathbf{x}(n+1) = \mathbf{A}(n)\mathbf{x}(n) + \mathbf{B}(n)u(n) + \mathbf{d}(n), \tag{10}$$

where $\mathbf{A}^{j\times j}$ and $\mathbf{B}^{j\times 1}$ are the plant and input matrices, $\mathbf{x}(n)^j$ and $\mathbf{d}(n)^j$ are the state and disturbance vectors, respectively, and $u(n)$ is the scalar control signal.

The scalar controlled output variable, $y(n)$, is given by

$$y(n) = \mathbf{C}\mathbf{x}(n), \tag{11}$$

where $\mathbf{C}^{1\times j}$ is the fixed output matrix.

Similarly, the reference model, describing the desired closed-loop system response (in terms of a desired settling time and level of damping), is given by

$$\mathbf{x}_m(n+1) = \mathbf{A}_m\mathbf{x}_m(n) + \mathbf{B}_m r(n), \tag{12}$$

$$y_m(n) = \mathbf{C}\mathbf{x}_m(n), \tag{13}$$

where the dimensions are the same as those for the plant state equation.

The reference model acts upon the reference signal $r(n)$, and $\mathbf{x}_e(n)$ is the error between the plant state $\mathbf{x}(n)$ and the desired state $\mathbf{x}_m(n)$. The gains of the feedback controller, $\{\mathbf{K}_e(n)^{1\times j}, K_I(n)\}$, are adapted to minimise $\mathbf{x}_e(n)$ via the control signal $u(n)$.

The control law for Er-MCSI is chosen as

$$u(n) = \mathbf{K}_e(n)\mathbf{x}_e(n) + K_I(n)x_I(n), \tag{14}$$

where $x_I(n)$ is the scalar discrete-time integral of the output error signal, defined as

$$y_e(n) = \mathbf{C}_e\mathbf{x}_e(n), \tag{15}$$

where $\mathbf{C}_e^{1\times j}$ is the output error matrix, which is determined in order to satisfy the strictly positive real (SPR) condition [7].

The controller gains are initially zero and are updated according to

$$\begin{aligned}\mathbf{K}_e(n) &= \mathbf{K}_e(n-1) + \beta\mathbf{q}_e(n) - \sigma\mathbf{q}_e(n-1),\\ K_I(n) &= K_I(n-1) + \beta q_I(n) - \sigma q_I(n-1),\end{aligned} \tag{16}$$

where

$$\begin{aligned}\mathbf{q}_e &= y_e\mathbf{x}_e^{\mathrm{T}},\\ q_I &= y_e x_I,\end{aligned} \tag{17}$$

$$\sigma = \beta - \alpha\Delta \tag{18}$$

and α and β are positive scalar adaption weights and Δ is the sampling interval.

The Er-MCSI controller implementation for this application is shown in Fig. 5, where the G_c block represents the controller blocks from Fig. 4. As before, $v(n)$ represents the engine vibration and $P(z)$ and $S(z)$ represent the primary path and secondary path (or plant) dynamics, respectively.

The plant is first order, since only the acceleration state is used to update the controller gains. Additionally, this state is directly measured, so $\mathbf{C} = 1$. The desired state $\mathbf{x}_m(n)$ is zero since the chassis is required to be stationary; hence, no reference model is required and $\mathbf{x}_e(n)$ is equal to

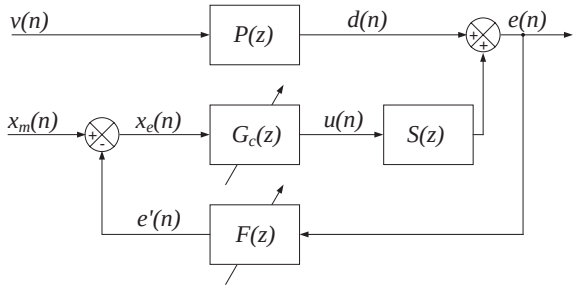

Fig. 5. Implementation of the Er-MCSI algorithm for cancelling periodic noise.

$e'(n)$, a filtered version of the scalar measured error signal $e(n)$. $F(z)$ is a sharp band-pass filter with a variable central pass frequency equal to the known engine speed. The purpose of this filter is twofold:

1. It reduces the low-frequency content of the measured chassis vibration $e(n)$. As already stated, it is undesirable for the controller to respond to this input.
2. It reduces the high-frequency content of $e(n)$. This is required to prevent the controller gains from winding up and inducing instability.

By adjusting the filter to pass frequencies near the engine-running speed, the band-width of the error signal remains narrow. This improves the stability of the controller and permits the use of higher adaption weights $\{\alpha, \beta\}$, which in turn improves the cancellation performance. The filter $F(z)$ is described in detail in Section 4.

The algorithm was implemented on a dPSACE DS1104 board operating at 10 kHz. This higher sampling rate compared to the LMS-based algorithm is afforded by substantial savings in computational effort.

4. Filtering with minimal phase shift

Filters are used to attenuate or amplify specific frequency components of a signal. Commonly filters are designed purely to produce a particular magnitude response, but will also affect the phase of the signal. The introduction of a phase lag to control signals can result in instability and is thus highly undesirable.

For the application described here, filtering is necessary both to attenuate noise and also to attenuate high-amplitude low-frequency components resulting from road inputs. Additionally, the road input frequency and the lowest frequency of interest are typically around 10 and 30 Hz, respectively. The attenuation of the 10 Hz component without affecting the 30 Hz component requires a sharp filter, and standard filter designs will inevitably introduce substantial phase shift to the frequencies of interest. This section describes the design of a filtering strategy that will separate the frequencies of interest and not introduce significant phase shift.

An initial approach to this problem may be to implement a fixed finite impulse response (FIR) filter, with the tap weights designed to give a band-pass response encompassing the frequencies of

interest. The advantage of the FIR filter is guaranteed stability. The disadvantage is that, in order to create a sharp filter a very large number of weights may be required, which introduces a large phase shift to the signal.

The strategy employed here is to measure the instantaneous engine frequency and use a very narrow low-order infinite impulse response (IIR) band-pass filter centred on that frequency. Thus, the centre frequency of the filter would move with the measured engine speed. This would enable a filter with very small phase shift to be designed.

To design the discrete-time variable filter, we begin by taking a standard normalised second-order band-pass filter with the following continuous-time Laplace transfer function:

$$H_{\mathrm{bp}}(s) = \frac{E'(s)}{E(s)} = \frac{s/Q}{s^2 + s/Q + 1}, \tag{19}$$

where $E(s)$ and $E'(s)$ are the Laplace transforms of the continuous-time error signal $e(t)$ and the filtered error signal $e'(t)$, and Q determines the damping of the filter (i.e. the notch width and depth).

The discrete-time transfer function is obtained by mapping to the z-domain using the bi-linear transform [12]

$$s = c\left(\frac{z-1}{z+1}\right), \tag{20}$$

where

$$c = \cot\left(\frac{\mathrm{i}\omega_e \Delta}{2}\right) \tag{21}$$

is the frequency warping coefficient, which compensates for the inherent inaccuracy in the bilinear transform, $\mathrm{i}\omega_e$ is the ith order of the measured instantaneous engine frequency, and Δ is the fixed sampling interval.

Thus, we obtain the discrete-time transfer function:

$$H_{\mathrm{bp}}(z) = \frac{c/Q(1-z^{-2})}{z^{-2}(c^2 - c/Q + 1) + z^{-1}(2 - 2c^2) + (c^2 + c/Q + 1)}, \tag{22}$$

whose co-efficients are updated at each time step according to the instantaneous measured frequency.

Fig. 6 shows the magnitude and phase responses for H_{bp} with a centre frequency of 30 Hz and a Q-value of 40. The filter is designed to pass only frequencies near its centre frequency, achieving 40 dB attenuation at 10 Hz and 0 dB at 30 Hz. The phase at the centre frequency is 0°, so providing the frequency is exactly known no phase shift will be introduced to the frequency component of interest.

The measured engine frequency will never be known exactly, however. Fig. 6 shows that the phase changes very rapidly around the centre frequency, so even small errors in the frequency estimate could introduce large amounts of phase shift. The frequency is estimated using a simple edge detector and counter, operating on a signal from a tachometer (commonly an accurate estimate of the engine speed would be available from an engine control unit signal).

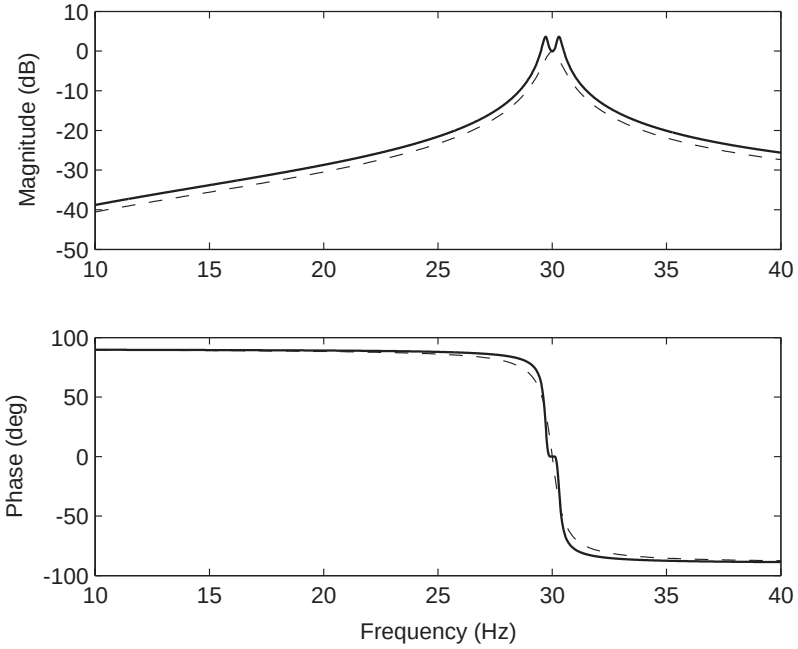

Fig. 6. Magnitude and phase responses of $H_{bp}(z)$ with $Q = 40$ (- -) and composite filter $F(z)$ (—).

To overcome this problem, three of these band-pass filters are arranged in parallel. Two are very sharp filters with centre frequencies on either side of the desired overall centre frequency. The third is a broader filter centred on the desired overall centre frequency. The composite filter is normalised to give unity gain around the centre frequency. Fig. 6 also shows a composite filter designed to have broadly similar characteristics to H_{bp} with $Q = 40$. Fig. 7 shows the magnitude and phase responses of the individual filters and the composite filter for a narrow frequency range around the centre frequency.

In the range around the centre frequency encompassing the maximum predicted frequency error, the magnitude is very nearly unity and there is almost zero phase shift.

The performance of the filter was established by comparison with H_{bp}. A function generator was used to simulate engine noise as a sinusoid with a swept frequency of 25–200 Hz. The frequency estimator was driven from the function generator square-wave output. A large-amplitude 10 Hz signal was added to the engine noise to simulate road inputs. A single band-pass and a composite band-pass filter were used to filter out the 10 Hz component. For interest, an LMS adaptive notch filter (similar to Fig. 2 but with no secondary path dynamics) designed to give similar levels of attenuation to the band-pass filters across the frequency sweep was also tested. Fig. 8 compares the magnitude and phase modifications to the signal of interest, which are introduced by the three filters.

It is clear that the composite filter substantially outperforms the single band-pass filter and the LMS notch filter in terms of phase response. The maximum phase shift introduced to the signal of interest using the composite filter is approximately 0.5°, compared to over 20° with the others. The

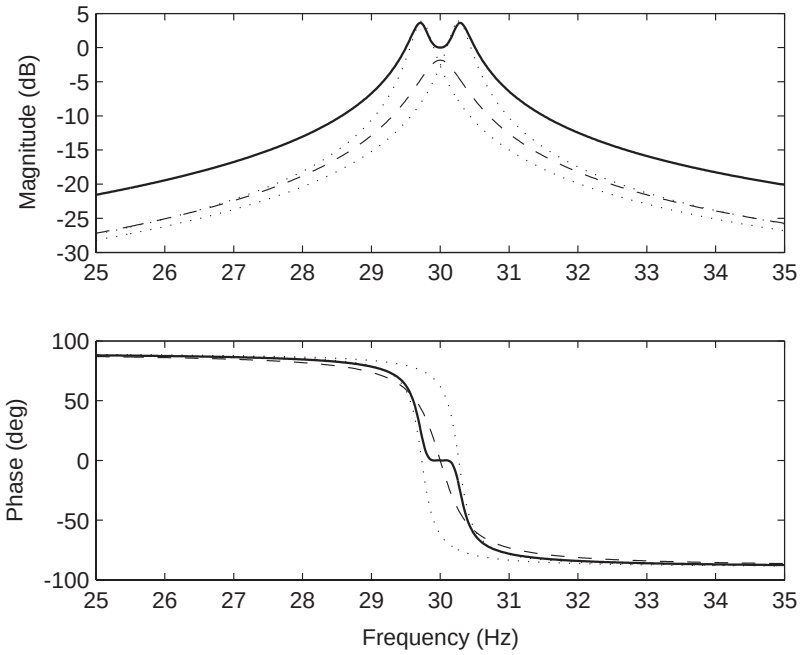

Fig. 7. Magnitude and phase responses of individual band-pass filter components (- -), (···) and composite band-pass filter $F(z)$ response (—).

magnitude accuracy is less important than the phase accuracy for this application since it does not affect the stability of the algorithm.

5. In vehicle testing

The algorithms described in Sections 2 and 3 were applied to a single active mount fitted to a saloon car equipped with a four-cylinder 2 l turbo-diesel engine. For both algorithms the adaption weights were experimentally determined as the maximum permissible values to achieve stable operation across all tests. The feed-back accelerometer was located near the chassis mounting point. Typically for a four-cylinder engine the second-order component of the vibration (twice the engine speed) is dominant, and this was the case for the vehicle in question. In particular, at around 2700 rpm the second-order vibration was manifest as substantial vibration transmitted through the accelerator pedal and driver's seat. The algorithms were thus set up to cancel this component (though they can be modified to cancel several harmonics if required).

The algorithms were tested while the car was in motion under normal driving conditions. In particular, the following tests were conducted:

1. second gear acceleration;
2. third gear acceleration;
3. constant 2700 rpm at motorway speed.

The algorithms were evaluated in terms of speed of response, transient performance and stability.

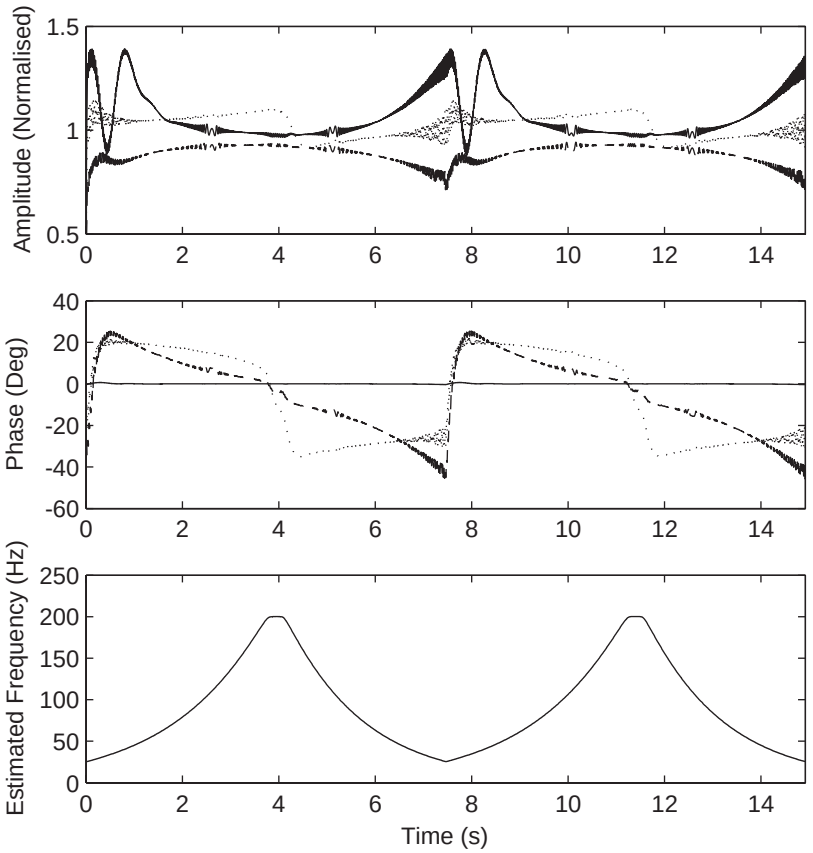

Fig. 8. Magnitude and phase modifications to 30 Hz signal component when filtered from large amplitude 10 Hz signal using: LMS adaptive notch filter (· · ·), second-order band-pass filter $H_{bp}(z)$ with $Q = 40$ (- -), composite band-pass filter $F(z)$ (—).

6. Results

6.1. Speed of response

The algorithms were switched on with the engine running at a constant speed of approximately 2700 rpm with the car stationary. Fig. 9 shows the magnitude of the second-order vibration component following algorithm switch-on. It is seen that the Er-MCSI and FXLMS algorithms perform similarly, achieving approximately 90% reduction in less than 500 samples.

6.2. Transient performance

The transient performance of the algorithms was assessed with the car in motion. Fig. 10 shows the results for second gear acceleration over approximately 5 s for the three cases where the system is uncontrolled, and controlled by the Er-MCSI and FXLMS algorithms, respectively. The

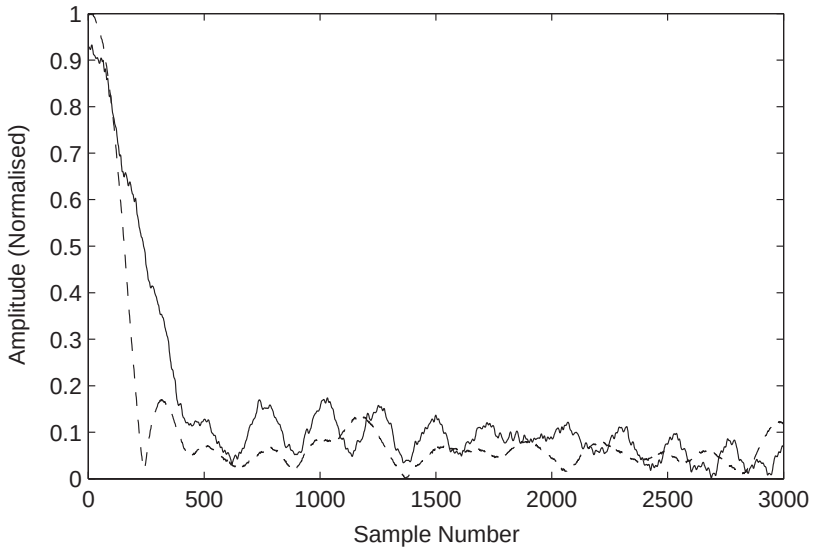

Fig. 9. Amplitude of error signal following controller switch-on for Er-MCSI (—) and narrow-band FXLMS (- -).

dominant second-order vibration component appears in the frequency range 50–100 Hz (and is the only component to be cancelled), and several higher-order harmonics are also visible. The plots were generated using a moving window FFT.

It is observed that both the Er-MCSI and FXLMS algorithms achieve similarly high levels of cancellation across the transient.

Fig. 11 shows similar results for third gear acceleration over approximately 8–10 s.

Again, the algorithms have performed very similarly. In both figures, the effect of road inputs is clearly visible as low-frequency content.

6.3. Stability

Fig. 12 shows the trajectories of the Er-MCSI gains during a transient up and down the frequency range of interest (50–100 Hz for the dominant second-order vibration). Over time the gains were observed to behave in a stable fashion, exhibiting no wind-up and taking similar values at a given frequency during transients. The differences in gain values shown in Fig. 12 are attributable to the nonlinear behaviour of the system, different rates of change of frequency during the transients and errors in frequency estimation.

The measure of stability of the FXLMS algorithm is the quality of the identification of $\hat{S}$. Fig. 13 shows the gain and phase representation of the filter impulse response generated by the FBLMS identification. The results are shown for three cases:

1. frequency-domain identification off-line (i.e. engine switched off);
2. frequency-domain identification on-line (engine running at constant 2700 rpm);
3. time-domain identification on-line (engine running at constant 2700 rpm).

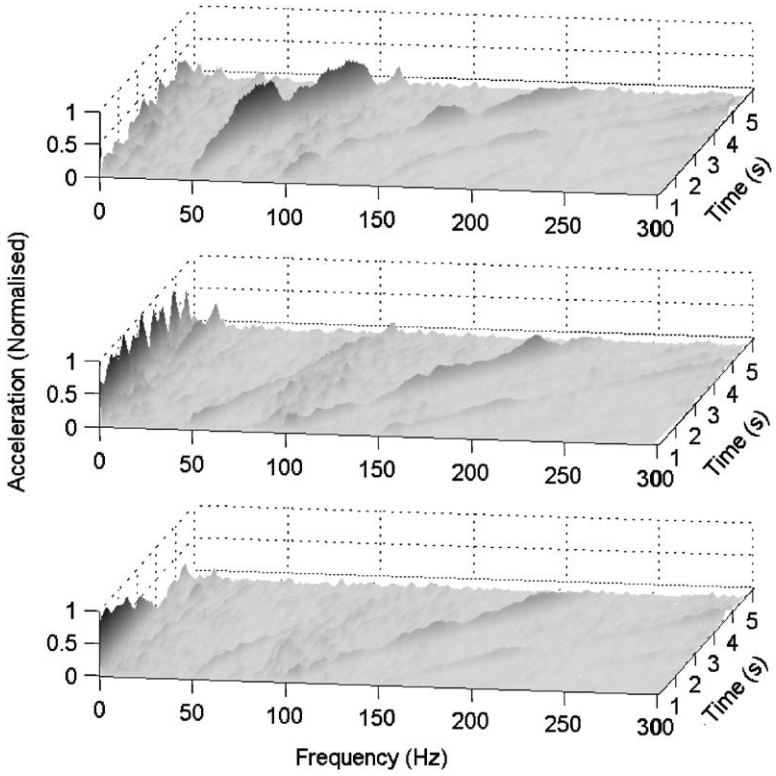

Fig. 10. Spectrogram showing second-order cancellation (50–100 Hz) during second gear acceleration (top: uncontrolled; middle: Er-MCSI; bottom: FXLMS).

The time-domain identification is included to highlight the effects of signal corruption. For the off-line case, time-domain and frequency-domain identification generate very similar filters, and they are sufficiently accurate to guarantee the overall stability of the algorithm (assuming sufficiently small μ). For the on-line case, it is observed that the frequency-domain identification remains sufficiently accurate but the time-domain identification does not. In particular, the anti-resonance at 200 Hz is not present and consequently the phase estimate in the higher portion of the frequency range of interest is very poor. This is because the estimate was corrupted by strong signal content at approximately 180 Hz, corresponding to the un-cancelled fourth-order vibration component.

6.4. Cancellation level

Fig. 14 shows the levels of cancellation achieved across the frequency range of interest during acceleration in third gear.

The vibration levels have been reduced to the level of background tyre noise across the spectrum, with the algorithms typically achieving between 50% and 90% reduction.

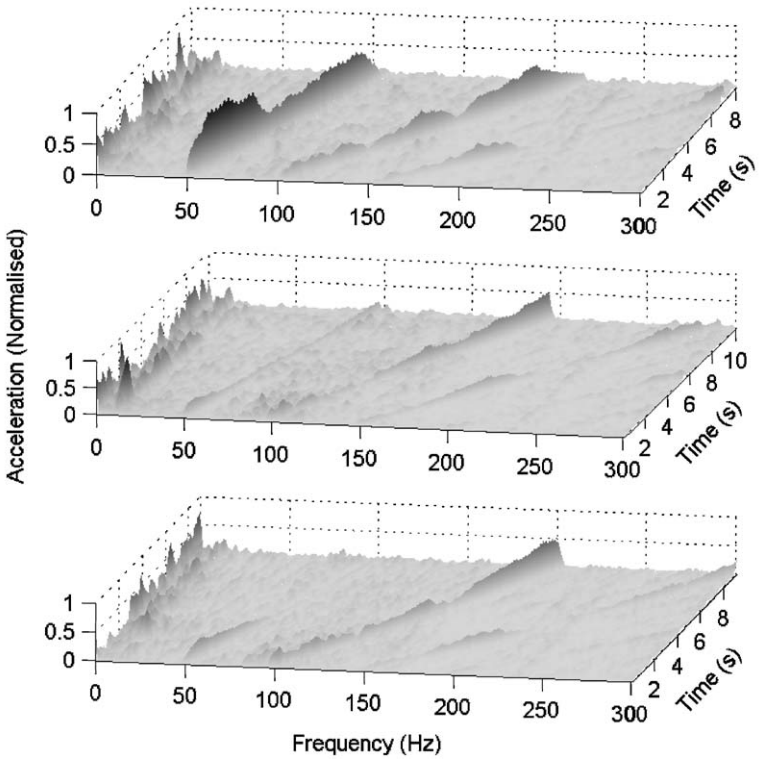

Fig. 11. Spectrogram showing second-order cancellation (50–100 Hz) during third gear acceleration (top: uncontrolled; middle: Er-MCSI; bottom: FXLMS).

7. Computational requirements

The computational requirements of the algorithms are compared in terms of their execution times within a time-step. The FXLMS algorithm was implemented as a multi-tasking system (a cancellation task and an on-line system identification task), with system identification executed in the background. Hence, the identification task execution time is within a length L block of time-steps. The execution times are summarised in Table 1.

The execution times for the FXLMS cancellation task and the Er-MCSI algorithm are similar, but the FXLMS system identification task is comparatively intensive. Furthermore, for a two-input/two-output coupled system, four identification tasks would be required for the FXLMS algorithm. This would further increase its computational requirements compared to the equivalent Er-MCSI algorithm.

8. Conclusions

Two active vibration cancellation algorithms have been tested with an active engine mount in conjunction with a saloon car equipped with a four-cylinder turbo-diesel engine. Neither

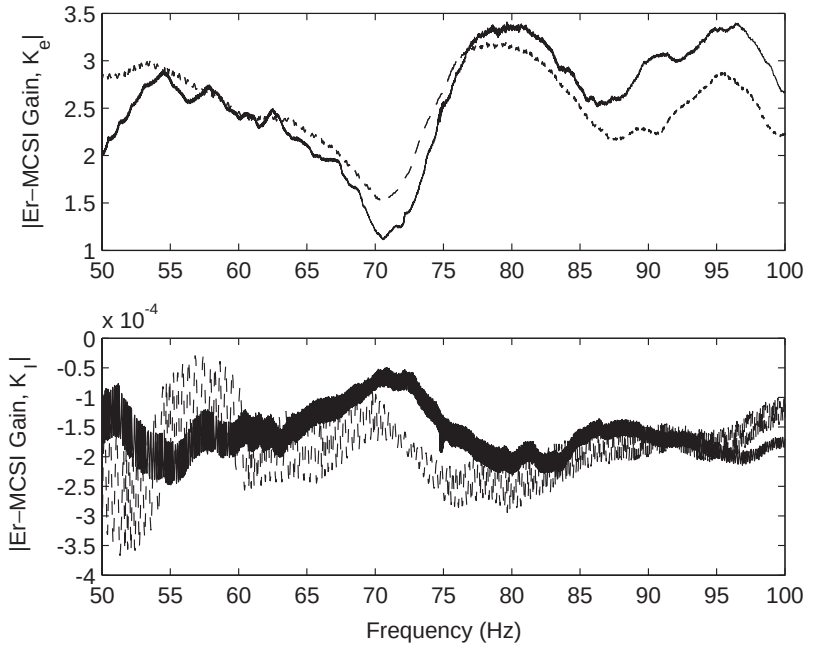

Fig. 12. Er-MCSI gains during engine speed transient from 25–50–25 Hz over approx. 20 s corresponding to a second-order transient of 50–100–50 Hz. Frequency increasing (—), frequency decreasing (- -).

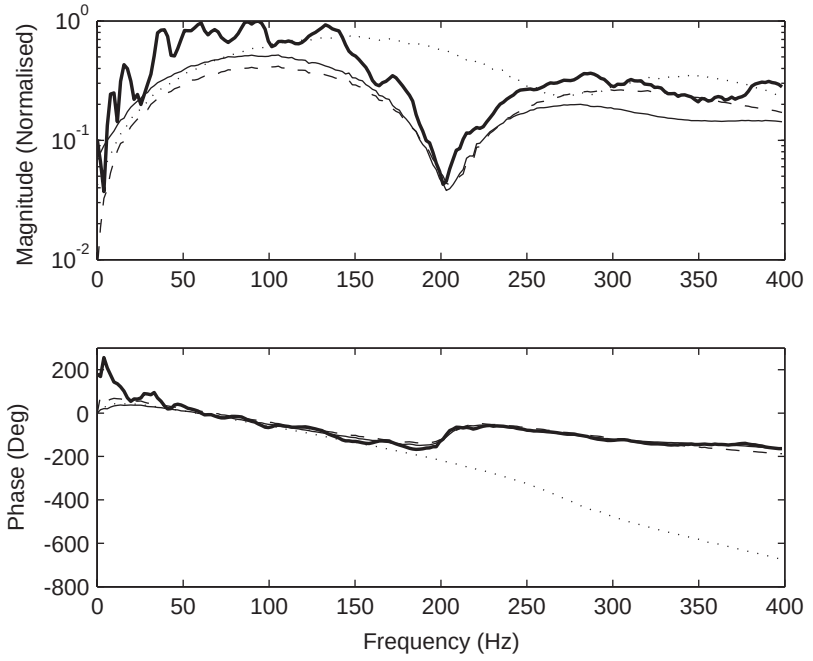

Fig. 13. Magnitude and phase responses of secondary path estimates using: spectral analysis (▬), FBLMS off-line (—), FBLMS on-line (- -), FXLMS time-domain (···).

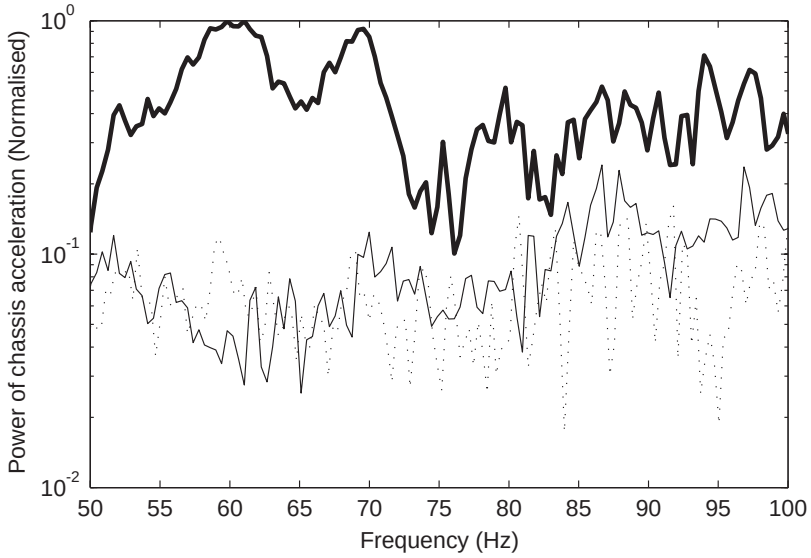

Fig. 14. Cancellation levels achieved during normal driving. Uncontrolled (—), Er-MCSI (—), FXLMS with frequency-domain ID (···).

Table 1
Algorithm computational requirements

Algorithm	Task	Execution time (s)
FXLMS	Cancellation	6.0e−5
	System identification	1.9e−3
Er-MCSI	Single task	6.4e−5

algorithm requires a priori knowledge of the system dynamics—an essential property for this application. The two algorithms were:

1. FXLMS with frequency-domain identification of secondary path dynamics;
2. Er-MCSI with a variable band-pass filter.

The Er-MCSI algorithm was found to perform very similarly to the FXLMS algorithm with frequency-domain identification, both in terms of speed of convergence and level of cancellation. The operation of the Er-MCSI algorithm in this application is made possible by using a variable band-pass filter to pick out a particular frequency for cancellation without introducing significant phase shift to the filtered signal.

The Er-MCSI algorithm provides a significant computational advantage over the FXLMS algorithm since it requires no on-line system identification. This advantage will become more significant when the algorithms are extended to control multiple mount-sensor systems.

Acknowledgements

The authors gratefully acknowledge the financial and technical support of Avon VMS, and in particular Dr. M. Fursdon, during the course of this work.

References

[1] P.M.T. Fursdon, A.J. Harrison, D.P. Stoten, The design and development of a self-tuning active engine mount, *IMechE (C577/018/2000)* (2000) 21–32.
[2] B. Widrow, J.R. Glover, J.M. McCool, J. Kaunitz, C.S. Williams, R.H. Hearn, J.R. Zeidler, E. Dong, R. Goodlin, Adaptive noise cancelling: principles and applications, *Proceedings of the IEEE* 63 (12) (1975) 1692–1716.
[3] S. Haykin, *Adaptive Filter Theory*, third ed., Prentice-Hall, Englewood Cliffs, NJ, 1996.
[4] Y. Nakaji, S. Satoh, T. Kimura, T. Hamabe, Y. Akatsu, H. Kawazoe, Development of an active control engine mount system, *Vehicle System Dynamics* 32 (2–3) (1999) 185–198.
[5] B. Riley, M. Bodie, J. Hoying, K. Majeed, S. Tewani, Active Vibration and Noise Cancellation Control of Four Cylinder Engines—An Application, *Society of Automotive Engineers*, No. 951299, 1995.
[6] D.C. Quinn, Automotive active engine mount systems, *Proceedings of the AVL Conference Engine and Environment*, 1992.
[7] D.P. Stoten, S.A. Neild, The error-based minimal control synthesis algorithm with integral action, *Journal of Systems and Control Engineering, Part I, Mechanical Engineering* 217 (2003) 187–201.
[8] S.J. Elliott, I.M. Stothers, P.A. Nelson, A multiple error LMS algorithm and its application to the active control of sound and vibration, *IEEE Transactions on Acoustics, Speech, and Signal Processing ASSP-35* (10) (1987) 1423–1434.
[9] S.M. Kuo, D.R. Morgan, *Active Noise Control Systems*, Wiley, New York, 1996.
[10] J.J. Shynk, Frequency-domain and multirate adaptive filtering, *IEEE Signal Processing Magazine*, January 1992, pp. 14–37.
[11] D.P. Stoten, H. Benchoubane, Robustness of a minimal controller synthesis algorithm, *International Journal of Control* 51 (4) (1990) 851–861.
[12] A.V. Oppenheim, R.W. Schafer, *Discrete-time Signal Processing*, Prentice-Hall, New York, 1989.