GSM TDMA Frame Rate Internal Active Noise Cancellation

Ingvar Claesson
Blekinge Institute of Technology, SE 372 25 Ronneby, Sweden

Andreas Nilsson
Ericsson Mobile Platform, SE 221 00 Lund, Sweden

(Received 2 April 2002; accepted 11 August 2003)

A common problem in the world's most widely-used cellular telephone system, the GSM system, is the interfering signal generated by the switching nature of TDMA cellular telephony in handheld and other terminals. Signals are sent as chunks of data, speech frames, equivalent to 160 samples of data corresponding to 20 ms at sampling rate of 8 kHz. This paper describes a study of two different software solutions designed to suppress such interference internally in the mobile handset. The methods are 1) notch filtering, which is multiplicative in frequency, and 2) subtractive noise cancellation, which is an alternative method employing correlators. The latter solution is a straigtforward, although somewhat unorthodox, application of "in-wire" active noise control. Since subtraction is performed directly in the time domain, and we have access to the state of the mobile, it is also possible to consider a recurring pause in the interference caused by the idle frame in the transmission, when the mobile listens to other base stations communicating. More complex control algorithms, based on the state of the communication between the handset and the base station, can be utilised.

1. INTRODUCTION

In GSM mobile telephony, it is a common problem that an interfering signal is introduced into the microphone signal when the mobile telephone is transmitting. This interfering signal is transmitted along with the speech signal to the receiver. Due to the humming sound of the interfering signal it is commonly referred to as the '*bumblebee*.'

Since interleaving of data are utilised and since control data transmission is also necessary, the connection between transmitter/receiver frames and speech frames is somewhat complicated. Data from a speech frame of 20 ms are sent in several bursts, each occupying 1/8 of a transmitting frame. The radio circuits are switched on and off with the radio access rate frequency. An electromagnetic field pulsating with this frequency and its harmonics disturbs its own microphone signal, as well as electronic equipment in the vicinity, producing in some cases an annoying periodic humming noise in the uplink speech from the handset to the base station.

The bumblebee is generated by the switching nature of TDMA cellular telephony, where the radio circuits are switched on and off. During the time the radio is switched on, known as the 'time slot,' the mobile telephone transmits its information by sending electromagnetic impulses. These impulses are induced in the microphone path and generate interference, which consists of the fundamental frequency and its harmonics. The fundamental switching rate is approximately 217 Hz, or more specifically, $5200/(3 \cdot 8)$ Hz, according to the GSM standard.[1]

Since the frequency components of the disturbing periodic humming noise are crystal-generated and accurately known, it is possible to estimate the cosine and the sine parts of these with correlators. This is easily done by correlating the microphone signal with sinusoids having the same crystal-generated frequencies as the disturbing frequencies. By generating the cosine and sine signals with correct signed amplitudes and then subtracting these from the microphone signal, the humming bumblebee is almost perfectly suppressed in the microphone signal. This is a classical example where in-wire subtractive active noise control is beneficial.[2,3]

Depending on the power level the mobile telephone is transmitting, how it is held and if one uses portable hands-free equipment or not, the amplitudes and phases of the fundamental and its harmonics will vary. When the mobile changes time slot, i.e., during a hand-over between base stations, the amplitudes and phases will also change abruptly.

Earlier solutions of this problem have utilised different hardware constructions, i.e., better placement of the components, usage of special electronics and microphones, reconstruction of analog parts, etc. However, this is expensive, time consuming and becomes increasingly difficult since the mobile telephones are constantly shrinking in size, thus causing the microphone to be situated closer to the transmitting antenna.

The solution to the problem presented in this paper makes use of the fact that the disturbance, after a Fourier series expansion, can be accurately described by a sum of sinusoids with well-defined frequencies. Two time domain software solutions are evaluated in this study to attenuate these frequency components of the digitised microphone signal directly in the base band by using synchronised correlators and notch filtering.

The best results were achieved by estimating the amount of the different sinusoids with correlators, and then subtracting these sinusoidal estimates from the microphone signal, as opposed to conventional notch filtering. This is an illustrative example of an application in which subtraction of disturbances, typical for active noise control,[2,4] is suitable.

2. PROBLEM BACKGROUND AND SIGNAL MODEL

The humming bumblebee disturbance is a result of the transmitting technique used in GSM, Time Division Multiple Access (TDMA). The handheld mobile telephone, formally denoted the mobile station (MS), sends information during the time slot that it is assigned. Eight time slots make one TDMA-frame, in which the time slots are numbered 0-7. A mobile uses the same time slot in every TDMA-frame until the network orders it to another time slot, i.e., when the traffic is rerouted via another base station, a handover. The duration of a time slot is 3/5200 seconds, and the period time of the TDMA-frames is $8 \cdot 3/5200$ seconds. During the assigned time slot, the mobile transmits its information by sending electromagnetic bursts. These are induced in the analog microphone path and produce an annoying periodic interference in the uplink speech. The fundamental frequency is $1/(8 \cdot (3/5200)) \approx 217$ Hz in full rate (FR).

There is another case that is not so common but still worth mentioning, half rate transmission (HR), where the radio access pattern differs considerably from FR. This communication scheme offers cheaper traffic with slightly decreased speech quality, but approximately twice as many connections in the ideal case. The period of the interference in this case is $1/(8 \cdot 2 \cdot (3/5200)) \approx 108$ Hz, which is half the frequency of the FR, since the mobile telephone is only transmitting during every other time slot.

Some mobile networks supports a feature known as discontinuous transmission (DTX), which is a mechanism, which allows the radio transmitter to be switched off most of the time during speech pauses. During these pauses the background noise is averaged and only silent descriptor (SID) frames are transmitted to the receiver. A SID frames thus contains no disturbing frequencies, and consequently, the algorithm is not allowed to run during DTX.

2.1. Analysis of the Bumblebee

A typical recorded disturbed signal from a silent room can be seen in Fig. 1. The interfering signal is periodic but somewhat complicated since, in the case of FR, there is no transmission when the mobile telephone is listening to other base stations. Such silent frames occur once every 26 TDMA-frames and are denoted as *idle* frames. Idle frames are illustrated in Figs. 2 and 3. In the HR case, the disturbance pattern is even more complex, but we refrain from detailed analysis here. We observe that since the state of the communication between the mobile telephone and base station is known, sufficient information to ascertain whether estimation and/or cancellation should take place or not is always at hand.

The simple radio access pattern for FR as well as the more complex pattern for an even HR channel can also be seen in Figs. 2 and 3, respectively.

Obviously, the idle frame should be considered when eliminating interference. Since the disturbance is periodic, it can be viewed as a Fourier series expansion

$$x_p(n) = \sum_{k=1}^{K} C_k \sin\left(2\pi k (f_0/f_s) n + \theta_k\right), \tag{1}$$

where K denotes the number of tones (fundamental plus harmonics), f_s is the sample frequency, and f_0 represents the frequency of the fundamental tone.

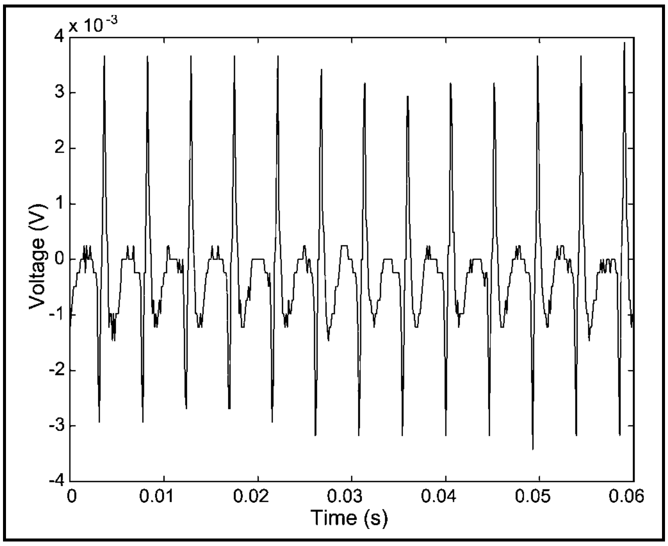

Figure 1. Interfering signal at the microphone A/D converter recorded in a silent room with no speech.

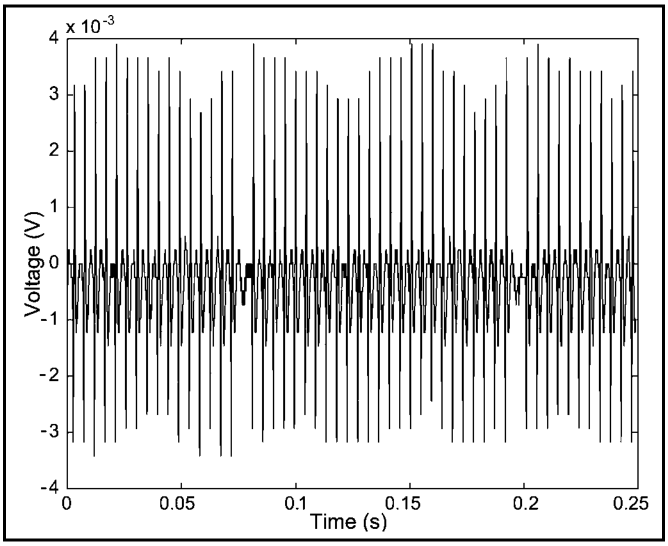

Figure 2. Pattern for interfering signal recorded in a silent room, full rate.

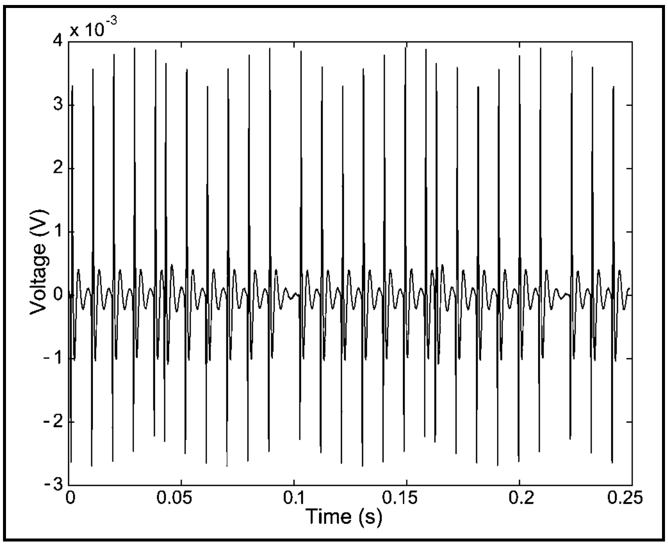

Figure 3. Pattern for interfering signal recorded in a silent room, half rate.

The number of tonal components K that are needed to represent the disturbance are limited by the sampling rate of the signal, which is 8 kHz. Consequently, the interfering signal after sampling will only consist of frequencies below 4 kHz since aliasing is carefully avoided in the mobile telephone. Further filters connected to the A/D conversion and the speech coder also band limit all signals, including bumblebee disturbance, to approximately 300-3400 Hz. Hence, the fundamental tone and the 15th harmonic will be slightly attenuated, as shown in Fig. 4.

A similar Fourier series expansion can of course be carried out for the HR case but the details are omitted in this paper. However, we observe that in this case the fundamental frequency, f_0, equals half the fundamental frequency in the full rate case. Hence, almost the double amount of harmonics is needed within the telephone frequency range to represent the disturbance. A comprehensive description illustrating the transmission patterns for both FR and HR transmission is given in Fig. 5.

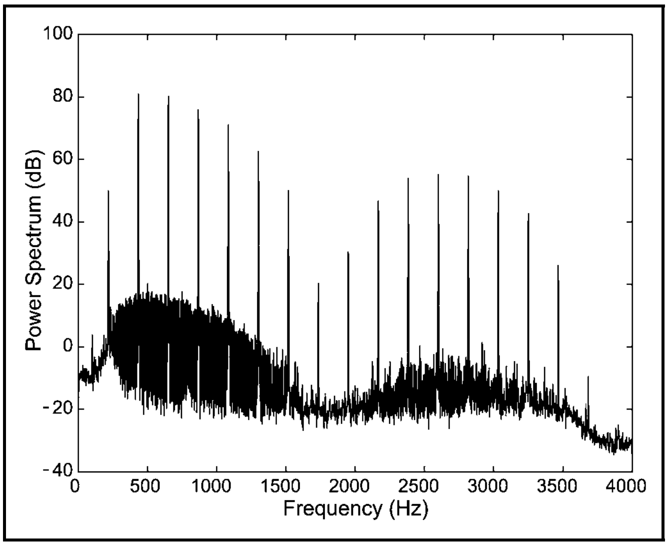

Figure 4. Spectrum of periodic bumblebee disturbance in random noise background.

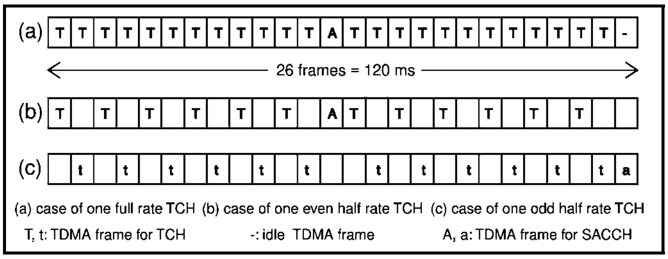

Figure 5. Transmission patterns in GSM full-rate and half-rate. TCH denotes traffic channel, and SACCH slow associated control channel.

3. SOLUTION PROPOSALS

Two different methods to eliminate the bumblebee disturbance are proposed, both working in the time domain.

These methods are linear time-invariant notch filters, which work on a sample-by-sample basis, and noise cancelling correlators, which work frame-wise on 160 samples in each time slot of 20 ms duration, i.e. the standardised slot duration in GSM at 8 kHz sampling rate.

3.1. Notch Filters

A notch filter consists of a number of deep notches, or ideally nulls, in its frequency response (see Fig. 6). Such a filter is useful when specific frequency components of known frequencies must be eliminated.[5,6] To eliminate the frequencies at $\omega_n, n = [1, \ldots, N]$, pairs of complex-conjugated zeros are placed on the unit circle at the angles ω_n

$$z_{n_{1,2}} = r_b e^{\pm j\omega_n}, \quad r_b = 1. \tag{2}$$

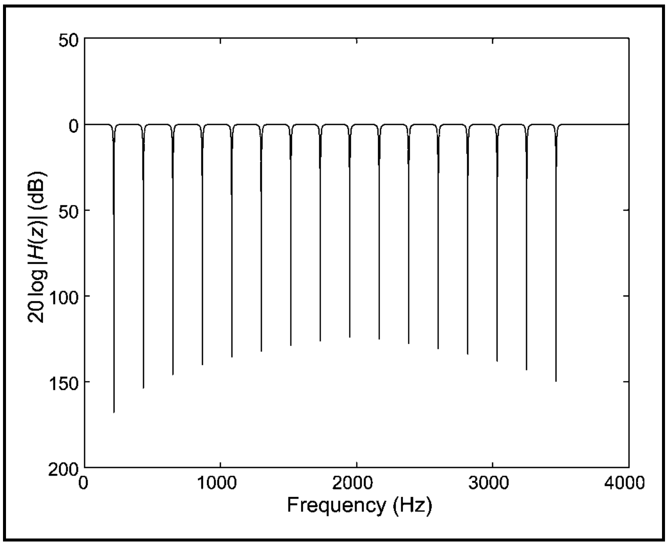

Figure 6. Frequency response of an FIR notch filter with $r_b = 1$, $N = 16$ and $\omega_n = n \cdot 2\pi \cdot (5200/(8 \cdot 3))$.

This results in a crude FIR notch filter with the system function

$$H(z) = B(z) = b_o \prod_{n=1}^{N} (1 - r_b e^{j\omega_n} z^{-1})(1 - r_b e^{-j\omega_n} z^{-1}). \tag{3}$$

The b_0 constant is chosen as

$$b_0 = 1 \Big/ \sum_{n=1}^{N} b_n \ , \tag{4}$$

to normalise the gain. To control the bandwidth of the FIR notches, poles are placed at the same angle as the zeros but with a slightly smaller magnitude. The positions of the poles are thus

$$p_{n_{1,2}} = r_a e^{\pm j\omega_n}, \quad 0 \ll r_a < r_b. \tag{5}$$

Consequently, the system function of the resulting notch filter is

$$H(z) = \frac{B(z)}{A(z)} = b_o \prod_{n=1}^{N} \frac{(1 - r_b e^{j\omega_n} z^{-1})(1 - r_b e^{-j\omega_n} z^{-1})}{(1 - r_a e^{j\omega_n} z^{-1})(1 - r_a e^{-j\omega_n} z^{-1})}, \tag{6}$$

where

$$b_0 = \sum_{n=1}^{N} a_n \Big/ \sum_{n=1}^{N} b_n \ . \tag{7}$$

The frequency response of the filter in Eq. (6) is plotted in Fig. 6. However, even sharp IIR notch filters have a non-negligible bandwidth, which leads to signal attenuation at frequencies also in the vicinity of the notches.

3.2. Orthogonal Correlators or Length-480 FFT Coefficients

Any band-limited periodic signal can be represented by a finite sum of sinusoids. Since we have periodic disturbance $x_p(n)$ superimposed on aperiodic speech $w(n)$, the model assumption for the input signal is given by

$$x(n) = x_p(n) + w(n) = \sum_{k=1}^{K} C_k \sin(nk\omega_0 + \theta_k) + w(n), \quad (8)$$

or alternatively

$$x(n) = \sum_{k=1}^{K} R_k \cos(2\pi f_k n) + I_k \sin(2\pi f_k n) + w(n), \quad (9)$$

where $f_k = k \cdot f_0$ and f_0 is the fundamental frequency of the disturbance. Since the disturbance frequencies are known, only the coefficients of the cosine and sine parts, R_k and I_k, need to be estimated.

The maximum likelihood (ML) estimate of known sinusoids in white noise background is given by correlation or matched filtering. This is equivalent, in our situation, to finding the Fourier Expansion coefficients, or in the discrete-time case, the Fast Fourier Transform (FFT) coefficients at the exact frequencies where the periodic disturbances are. Even if speech cannot be regarded as a white disturbance, it is still an attractive least squares (LS) solution to correlate out the sinusoidals.[7-9]

In order to inherently achieve unbiased LS estimates, correlation can be made over a whole number of periods for each sinusoidal. This corresponds to the fact that each disturbing frequency is situated exactly at an FFT bin. This is achieved if the correlation (FFT bin calculation) is made over 480 samples (3 frames) in the full-rate situation, and 960 samples (6 frames) in the half-rate case. Performing a pruned FFT with lengths of other lengths than factors of 2:s (2^M), in this case $N = 480$ or $N = 960$ is certainly not straightforward. Neither is it desirable in the present context, since we are only interested in the FFT bins where the periodic disturbance is present, typically only in 16 of the bins. Hence, an FFT is not the most efficient way to calculate the correlations in this case.

A sinusoidal correlator estimator consists mainly of a bank of dual product-adders, one for each frequency, one for each cosine and sine part, in total $2 \cdot K$ ($K = 16$) correlators of length $N = 480$ in the full-rate case. This makes it easy to estimate and compensate the bumblebee disturbance in 'real time,' frame by frame, by adding the correlation contribution of the most recent 160 samples, the present frame, and subtracting the correlation contribution of the 160 samples (3 frames back) in the frame leaving the estimation interval, i.e., the most recent 480 samples. To do this, the cosine and sine parts of the different frequencies are estimated by correlation in accordance with Fig. 7, yielding the estimates $\hat{R}_k$ and $\hat{I}_k$, respectively, in the two branches.

These signals are then subtracted from the input signal yielding

$$y(k) = x(k) - \sum_{k=1}^{K} \hat{R}_k \sin(2\pi f_k n) + \hat{I}_k \sin(2\pi f_k n). \quad (10)$$

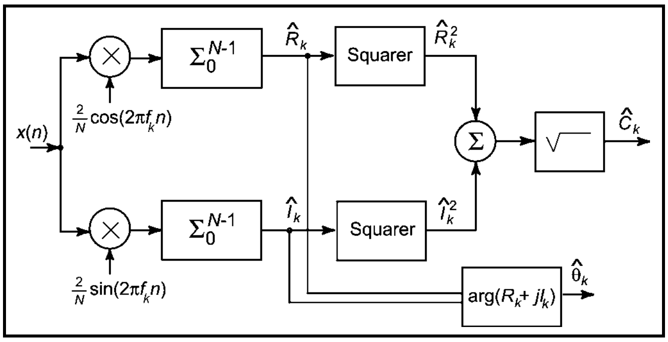

Figure 7. Sinusoidal estimation with correlators.

If the amplitude and phase are required instead, we proceed by

$$\sqrt{(\hat{R}_k \cos(\theta))^2 + (\hat{I}_k \sin(\theta))^2} = \hat{C}_k, \quad (11)$$

and the corresponding phase estimate of θ_k by calculating the four-quadrant angle

$$\hat{\theta}_k = \arg\left(\hat{R}_k, \hat{I}_k\right). \quad (12)$$

3.3. Implementation Aspects

This estimation is carried out block-wise using correlators. The amount of data in each block that is used for the estimation should preferably be done over an integer number of fundamental periods in order to avoid bias from incomplete periods. For the fundamental tone, which has the lowest frequency and thus requires the most samples, we need 480 samples to fulfil the requirement in the FR case. This is easily derived, since the frequency of the fundamental tone is $1/(8 \cdot (3/5200) \cdot 8000)$, and the sampling rate is 8 kHz. This gives $f_0/f_s = 13/480$, implying that 480 samples are needed to represent an integer number (13) of the fundamental periods with an integer number (3) of slots of 160 samples. In other words, to fulfil the bias-free requirement of whole periods, 13 fundamental periods are required which makes the block size 480, which is also the equivalent of 3 GSM frames, each with 160 samples. Since the length is given by 480 samples and only 16 tonal components are to be calculated, there is no need to use FFT algorithms. Instead, a more straightforward route is taken.

In discrete time, we simply correlate the received signal with the $16 \cdot 2$ basis functions of the correlators (cosines and sines) in order to obtain the coefficients for the cosines and sines. These estimates are subsequently used as coefficients for the amount each sinusoid should be subtracted from the received signal.

If estimation is performed during speech, the estimate of the bumblebee disturbance will be incorrect, since the speech contains high energy at the same frequencies as the disturbance. This problem is solved by only making estimates during speech pauses, a voice activity detector (VAD) is thus required. Fortunately, the mobile is already equipped with a VAD, which therefore can easily be utilised, as shown in Fig. 8.

The VAD information is further elaborated for several GSM frames, since a VAD algorithm works on 160-sample frames. A flag is set to one if speech is present. To consider the present frame as nonspeech, the three most recent frames (480 samples) and the following frame must all have $\text{VAD} = 0$. The reason for this is that even if $\text{VAD} = 0$ for the past three frames, it is wise to check the subsequent frame (n), since there may be the beginning of speech at the end of the present, most recent tentative estimation frame ($n-1$) which otherwise would destroy the estimation. As a result, the correlation estimate will be one frame older (delayed), but this is still a better solution. If the VAD conditions are not fulfilled, it is often much better to keep an old estimate than an erroneous one partially disturbed by speech, since the coefficients of the cosines and sines normally only vary slowly during operation.

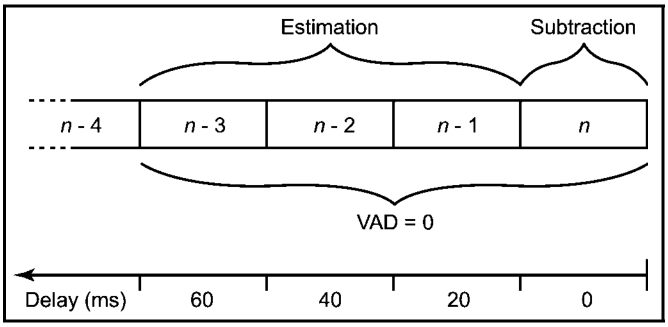

Figure 8. Estimation and subtraction when the VAD algorithm is used.

A more important observation is that the speech will not be delayed. To avoid any signal delay of the speech, we only estimate/correlate sinusoids on the three previous GSM frames, with delayed samples, though the subtraction is performed on the present GSM frame (Fig. 8).

The idle and silent states should also be considered. This is done by inhibiting disturbance subtraction during idle mode and preventing it from estimation/correlation during silent frames. Since the transmission state is locally known in the mobile telephone as well as the structure of the frames, Fig. 5, these states are easily handled in a software implementation.

4. CANCELLATION RESULTS ON RECORDED SIGNALS

The problem with the bumblebee disturbance is not just to eliminate it, but to do so without impairing speech quality. The following analysis is based on data recorded from the digital audio interface (DAI) in an Ericsson mobile. The DAI is the interface after the A/D-converter where the signal is pulse code modulated. This is the signal that enters the DSP, which is processed by the algorithm.

The frequencies that will be attenuated in the tests are: $k \cdot \omega_0$, $k = [1, \ldots, 16]$ and ω_0 is the fundamental tone of the bumblebee disturbance, $5200/(8 \cdot 3)$ Hz. With $K = 16$, the fundamental tone and 15 of its harmonics will be eliminated. This will span a range up to 3467 Hz which covers the frequency range of the telephone.

4.1. Notch Filter

Since the frequencies which constitute the bumblebee are well defined, we first apply a notch filter directly in the signal path to reduce the interference.

Implementation. The notches are made as deep as possible, so that ideally the frequencies in question are totally eliminated. This results in the following system function:

$$H(z) = \frac{B(z)}{A(z)} = \frac{\sum_{k=1}^{16} a_k}{\sum_{k=1}^{16} b_k} \prod_{k=1}^{16} \frac{(1 - r_b e^{jk\omega_0} z^{-1})(1 - r_b e^{-jk\omega_0} z^{-1})}{(1 - r_a e^{jk\omega_0} z^{-1})(1 - r_a e^{-jk\omega_0} z^{-1})}. \tag{13}$$

The calculations are made recursively on the whole data set. This will result in a convergence period at the start up and also when a handover between base stations occurs. Unfortunately, the notch filter is active also under idle frames, a drawback resulting from the fact that it works sample-by-sample and recursively, leading to unwanted artefacts during idle frames, when it tries to subtract a disturbance that is not present, i.e. a negative disturbance is added (Fig. 9).

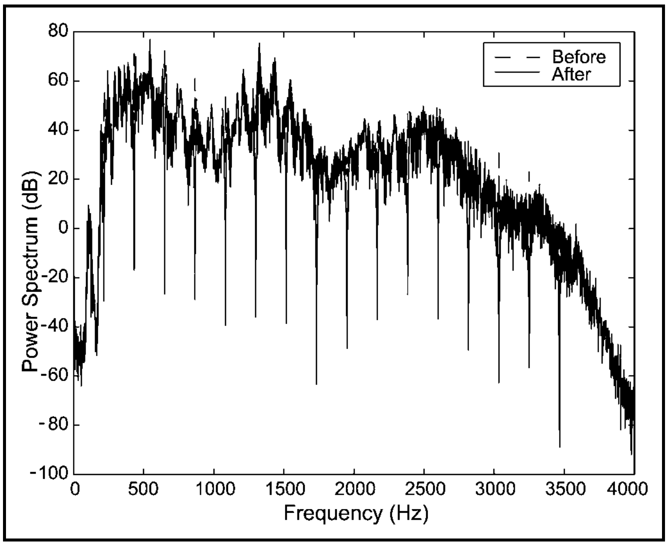

Figure 9. Cancellation of the bumblebee with notch filter in speech [S3]. Full rate, with speech.

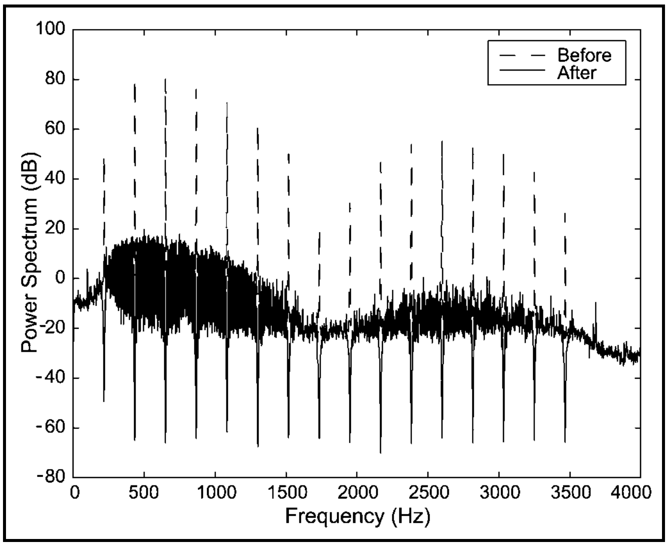

Figure 10. Cancellation of the bumblebee with notch filter. The bumblebee was recorded in a silent room [S2]. Full rate, no speech.

It can be seen that the bumblebee disturbance is considerably attenuated. However, this solution does not give a satisfactory result, since a portion of the speech is also attenuated, resulting in a 'canned' or metallic sound. This can be seen in Figs. 9 and 10. Another problem with this solution is that the periodic idle frame cannot be handled, resulting, in a new periodic interference 26 times lower in frequency (Fig. 11). The reason for this is that the notch filter consists of poles (autoregressive) which give feedback of the output signal ($y(t)$) continuously. Consequently, the bumblebee is added during the idle frame, according to the tails of the impulse responses of IIR filters.

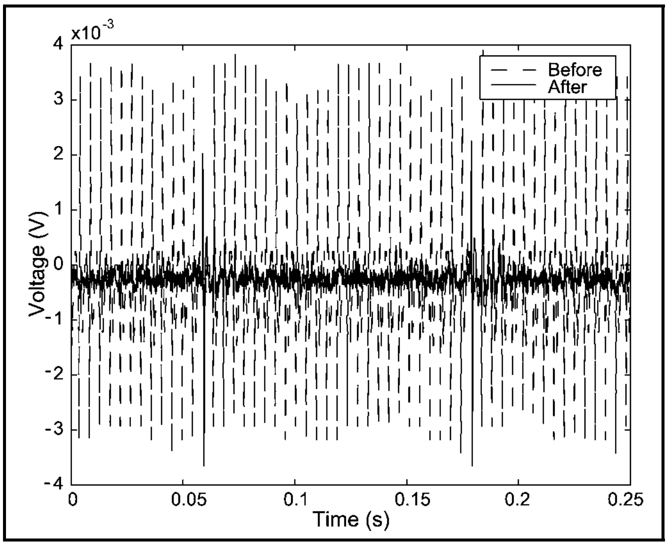

Figure 11. Time signal of the notched bumblebee.

4.2. Correlators

The data set that has been used is identical to that used when evaluating the notch filter. That is, the first test is done on data recorded both with speech and in a silent room (Figs. 12 and 13). The metallic sound and the periodic interference that appeared in the notch tests from the idle frame are also avoided, thanks to the time-limited subtractive nature of block correlation cancelling, thus avoiding long-tailed (recursive) impulse responses. This gives a very satisfactory result. Observe in Fig. 13 that only the bumblebee disturbance is attenuated.

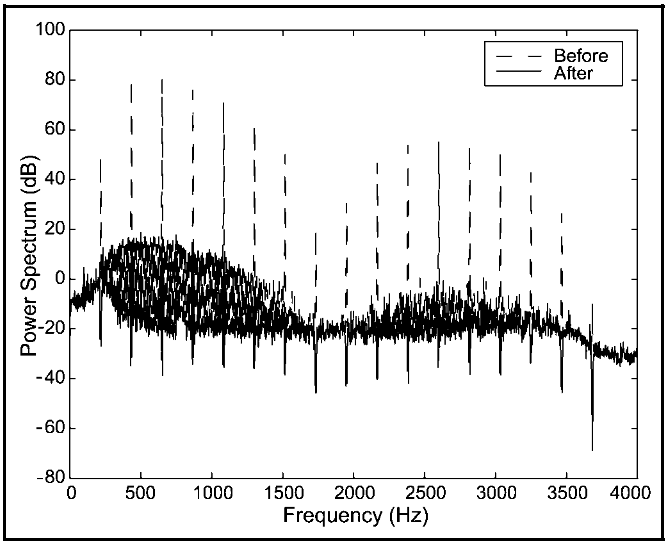

Figure 12. Cancellation of the bumblebee with correlators where idle mode has been taken into consideration. The bumblebee was recorded in a silent room [S4]. Full rate, no speech.

A corresponding and even more impressive result is also presented for the HR case in Fig. 14.

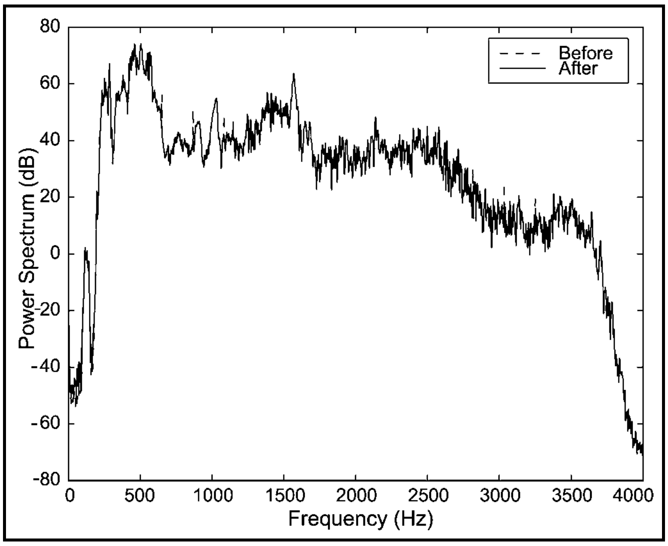

Figure 13. Cancellation of the bumblebee in speech with correlators where VAD and idle mode have been taken into consideration [S5]. Full rate, with speech.

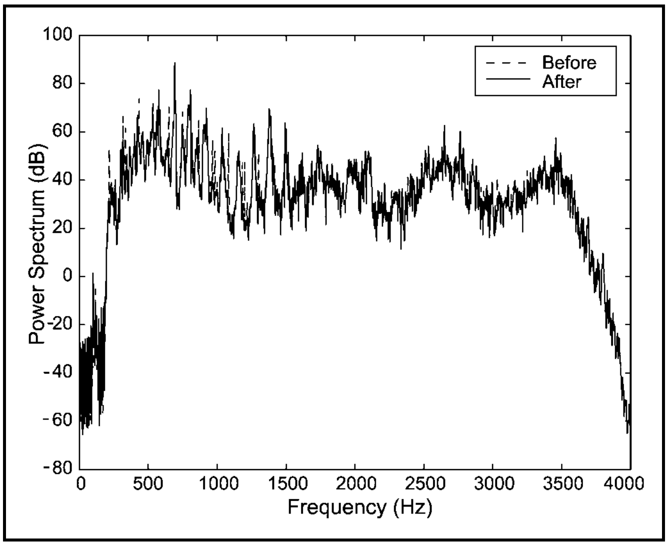

Figure 14. Cancellation of the bumblebee in speech with correlators in the half rate case where the VAD and idle frame have been taken into consideration [S7]. With speech.

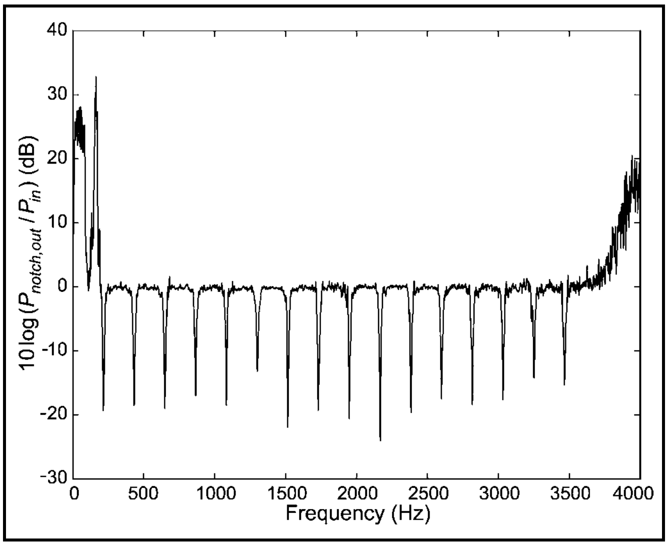

Figure 15. Divided power estimates with notch filter, no speech.

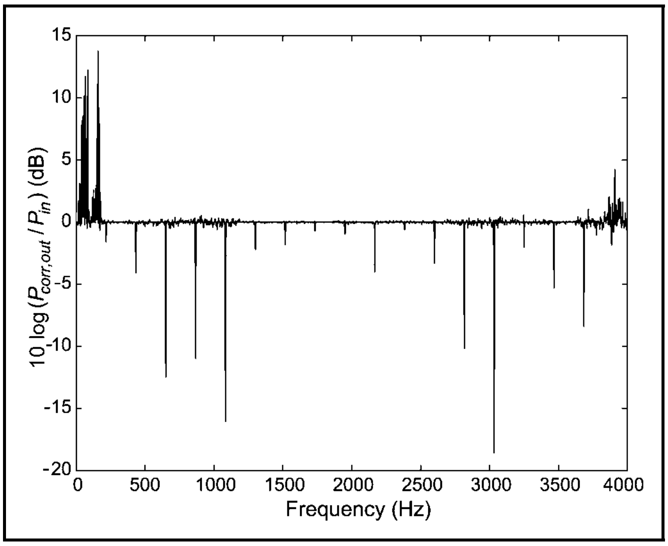

Figure 16. Divided power estimates with correlators, no speech.

Finally, an alternative type of comparison is introduced in Figs. 15 and 16 illustrating P_{out}/P_{in}, which gives the overall system attenuation both for the notch filter and the correlator. It can be observed that the notch filter gives both a deeper and wider attenuation, which explains the metallic sound and inferior quality as compared to the use of correlators.

5. COMPLEXITY AND IMPLEMENTATION ASPECTS

Complexity estimates have only been made for the correlators, since this solution was preferred. The most commonly used unit when performing complexity estimates is MIPS (millions of instructions per second). However, this can be a misleading measure because of the varying amounts of work done by an instruction. That is, an instruction on one processor may accomplish far more work than an instruction on another. This is especially important for DSP processors, which often have highly specialised instruction sets.

Similarly, MOPS (millions of operations per second) suffer from related problems: what counts as an operation and the number of operations needed to accomplish useful work varies greatly from processor to processor.

A third performance unit that can be used is MACS (multiply accumulates per second). Most DSP processors can complete one MACS per instruction cycle, making this unit equivalent to MIPS for DSPs. Furthermore, MACS estimates disregard the important data movement and processing required before and after.

After considering the various drawbacks, we selected the MIPS measure, which is given in Table 1.

Table 1. Complexity of the table approach.

Task	Instructions / 20 ms	MIPS
Correlation	$16\times2\times480$	0.768
Building $\hat{b}$	$16\times2\times2\times160$	0.512
Subtracting	160	0.008
Total	25760	1.288

The complexity calculations are based on the attenuation of 16 sinusoids. The estimation is performed on 480 samples, and the subtraction of the estimated signal on 160 samples. This is the way it should be done in the mobile to avoid a delay. The sinusoids and the cosinusoids are stored in a read-only memory (ROM) as a table. Another solution could be to use a digital sinusoidal oscillator. Such a solution does not require as much ROM memory as the table approach, but it is much more complex and does not generate the sinusoids and the cosinusoids perfectly.

To build up the 480 samples of the sinusoids, the table should contain an integer number of periods for each frequency. That is, $480/k$ samples with the exception of the frequencies stated in Table 2. If $K = 16$, a ROM of 6452 words, is required and the complexity is approximately 1.3 MIPS (Table 1). Control code and data transfers will also be needed. A very conservative estimation of the total complexity is 2 MIPS.

Table 2. Samples needed for the frequencies $k \cdot f_0$.

k	7	9	11	13	14
Samples needed	480	160	480	480	240

As mentioned before, the fundamental tone (and the first harmonic for HR) are already severely attenuated because of the filter, A/D converter and the speech coder. This makes it possible to also ignore these tones without degrading the result.

Symmetries in sinusoidal base functions and recursive estimation where the estimates are updated with the recent frame data of 160 samples can reduce the computational load by more than 50%. With this in mind we conclude that correlation cancelling is a cheap and convenient way of coping with the problem of humming bumblebee noise in GSM cellular telephony.

6. SUMMARY, CONCLUSIONS AND FUTURE WORK

In this paper we have compared two methods for eliminating an annoying self-disturbance in mobile telephone microphone signals originating from the telephones's own antenna. Such disturbance is caused by TDMA switching in GSM cellular telephones. The active noise control approach, which subtracts disturbances instead of filtering them out, has shown great potential. The aim is now to implement the algorithm in fixed-point precision.

REFERENCES

[1] GSM Standard (GSM 05.01 version 7.0.0 Release 1998), *Digital cellular telecommunications system* (*Phase 2+*), *Physical layer on the radio path* (*General description*).

[2] Kuo, M., and Morgan, D.R. *Active Noise Control Systems*, John Wiley & Sons, Inc., (1996).

[3] Chaplin, G.B.B., and Smith, R.A. Method and apparatus for cancelling vibrations, United States Patent no 4,490,841 Chaplin, December 25, (1984).

[4] Widrow, B, and Stearns, S.D. *Adaptive Signal Processing*, Prentice Hall, (1985).

[5] Proakis, J.G., and Manolakis, D.G. *Digital Signal Processing*, Prentice-Hall Inc., 343-345, (1996).

[6] Haykin, S. *Digital Communications*, John Wiley & Sons Inc., (1988).

[7] Peebles, P.Z., Jr. *Probability*, *Random Variables*, *and Random Signal Principles*, McGraw-Hill Inc., (1993).

[8] Kay, S.M. *Fundamentals of Statistical Signal Processing: Estimation Theory*, Prentice-Hall Inc., 183-198, (1993).

[9] Eriksson, P. On estimation of the amplitude and the phase function, Technical Report TR-148), University of Lund, Sweden, (1981).

APPENDIX

A number of demonstration sound wav-files can be reached by searching the more extensive *Bumblebee killer* report on-line at *http://www.bth.se/fou/*. Just search the archive with the keyword 'bumblebee.' All files are stored under mnemonic names.

[S1] Interfering signal, *Bumblebee*, recorded in a silent room.

[S2] Before and after Cancellation of the Bumblebee with Notch filter. No speech.

[S3] Before and after Cancellation of the Bumblebee with Notch filter in speech.

[S4] Before and after Cancellation of the Bumblebee with Correlators where the idle mode has been considered. No speech.

[S5] Before and after Cancellation of the Bumblebee with Correlators in speech where VAD and idle have been taken into consideration.

[S6] Before and after Cancellation of the Bumblebee with Correlators in the Half Rate case where the idle mode has been taken into consideration. No speech.

[S7] Before and after Cancellation of the Bumblebee with Correlators in speech in the Half Rate case where VAD and idle mode have been taken into consideration.