

Adaptively Canceling Server Fan Noise

Principles and Experiments with a Short Duct and the AD73522 dspConverter

by Paschal Minogue, Neil Rankin, Jim Ryan

INTRODUCTION

In the past, when one thought of noise in the workplace, heavy industrial noise usually came to mind. Excessive noise of that type can be damaging to the health of the worker. Today, at a much lower level, though not as severe a health hazard in office environments, noise from equipment such as personal computers, workstations, servers, printers, fax machines, etc., can be distracting, impairing performance and productivity. In the case of personal computers, workstations and servers, the noise usually emanates from the disk drive and the cooling fans.

This article is about the problem of noise from server cooling fans, but the principles can be applied to other applications with similar features. Noise from server cooling fans can be annoying, especially when the server is located near the user. Usually, greater server computing power means higher power dissipation, calling for bigger/faster fans, which produce louder fan noise. With effective fan noise cancellation, bigger cooling fans could be used, permitting more power dissipation and a greater concentration of computer power in a given area.

DESCRIPTION OF PROBLEM

Typically, server cooling-fan noise has both a random and a repetitive component. A spectrum plot of the fan noise of the Dell Poweredge 2200 server illustrates this (Figure 1). Also, the profile of the fan noise can change with time and conditions; for example, an obstruction close to the fan will affect its speed and thus the noise that it generates.




Figure 1. Profile of server fan noise.

ONE SOLUTION

One solution is to confine the propagation of much of the fan noise to a duct, and then use *active noise control* (ANC) to reduce the strength of the fan noise leaving the duct [1].

A block diagram of a basic ANC system as applied to noise propagating inside a duct is illustrated in Figure 2. The noise propagating down the duct is sampled by the upstream reference microphone and adaptively altered in the electronic feed-forward path to produce the antinoise to minimize the acoustic energy at the downstream error microphone. However, the antinoise can also propagate upstream and can disrupt the action of the adaptive feedforward path, especially if the speaker is near the reference microphone (as is always the case in a short duct). To counteract this, electronic feedback is used to neutralize the acoustic feedback. This neutralization path is normally determined off-line, in the absence of the primary disturbance, and then fixed when the primary noise source is present. This is done because the primary noise is highly correlated with the antinoise.

There are many problems associated with short-duct noise cancellation [2, 3, 4]. Acoustic feedback from antinoise speaker to reference microphone is more pronounced; the number of acoustic modes increases exponentially; duct resonances can cause harmonic distortion; and the group delay through the analog-to-digital converters, processing unit, and digital-to-analog converters can become significant [5]. This paper concentrates on the latter problem in particular.

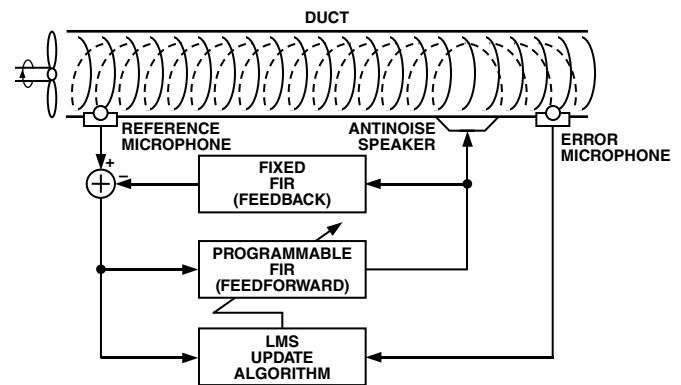


Figure 2. Basic duct ANC system.

THE IMPORTANCE OF GROUP DELAY

To provide an unobtrusive and viable solution in terms of size and cost, the smaller the duct is made the better; ideally, it should fit inside the server box—which would result in a very short acoustic path. To maintain causal relationships, the delay through the entire (mostly electronic) feedforward path has to be less than or equal to the delay in the forward acoustic path if broadband primary noise is to be successfully cancelled.

$$\delta_{ff} \leq \delta_{ap} \quad (1)$$

where δ_{ff} is the delay through the feedforward path and δ_{ap} is the primary acoustic delay.

In cases where the secondary speaker is recessed into its own short duct, the feedforward path will also include the acoustic delay through this secondary duct.

$$\delta_{ff} = \delta_e + \delta_{as} \quad (2)$$

where δ_e is the electronic part of the feedforward path and δ_{as} is the acoustic delay in the secondary duct.

The electronic delay in the feedforward path consists of group delay through the microphone, anti-aliasing filter and A/D

converter (ADC); processing delay (+ digital filter group delay) in the DSP; group delay through the D/A converter (DAC) and anti-imaging filter; and finally the delay through the secondary speaker.

$$\delta_e = \delta_{mic} + \delta_{adc} + \delta_{dsp} + \delta_{dac} + \delta_{spkr} \quad (3)$$

Therefore, from causality:

$$\delta_{mic} + \delta_{adc} + \delta_{dsp} + \delta_{dac} + \delta_{spkr} + \delta_{as} \leq \delta_{ap} \quad (4)$$

To minimize δ_{ap} , and thereby the length of the duct, δ_{ff} (and all its various components) should be made as small as possible. Let's assume that delays through the microphone, speaker and secondary duct path have already been minimized. Then, the remaining quantity to be minimized is $\delta_{adc} + \delta_{dsp} + \delta_{dac}$.

δ_{adc} can be minimized by using an oversampled ADC with low group-delay-decimation filtering; δ_{dsp} can be minimized by using a DSP with sufficiently high-MIPS and an efficient instruction set; δ_{dac} can be minimized by using an oversampled DAC with interpolation filtering having low group delay. The latter should be capable of being bypassed to further minimize the group delay.

Furthermore, the processor should use the most recent ADC sample as soon as it becomes available, and the DAC should use the latest DSP output result as soon as it becomes available. To achieve this, it must be possible to advance the timing of the DAC relative to the ADC in some fashion.

High-speed gain taps in parallel with the main processing path can help to ease the situation, especially in the case of practical short ducts where the noise can flank the acoustic path via the duct hardware itself.

Although a feedforward cancellation technique with a very low group delay is required to cancel the random component when using a relatively short duct, a feedback approach can be applied to cancel the repetitive component using an rpm sync signal as the reference input.

ANC ARCHITECTURE

Less group delay through the cancellation system means that a shorter duct can be used, making the approach more feasible and acceptable. To achieve very low group delay, a heavily oversampled sigma-delta converter technique is employed at the *analog front-end* (AFE) section of the system. Furthermore, both an *analog gain tap* (AGT) and a *digital gain tap* (DGT) can be utilized to provide even lower group delay in the processing path.

With no high-speed gain taps, a 2-ADC/1-DAC configuration (Figure 3) can be used, as all the processing is done digitally and at a relatively low rate. In Figure 3, each conversion channel is shown with its sample-rate conversion in a separate block: a decimator block in the case of the ADC channel and an interpolator block for the DAC channel.

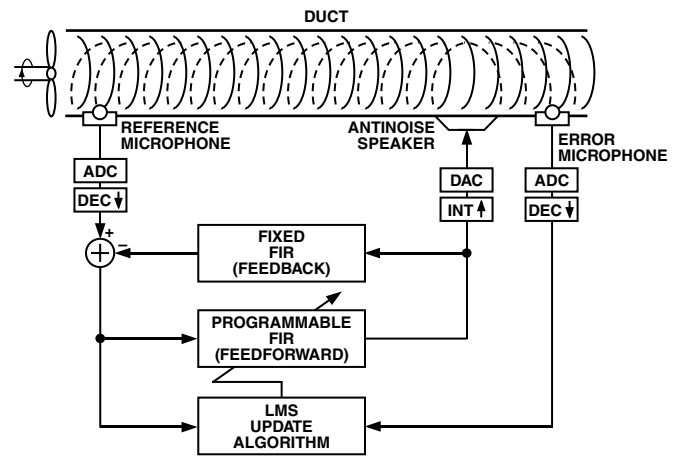


Figure 3. A 2-ADC/1-DAC configuration with no high-speed gain taps.

The introduction of gain taps is illustrated in Figure 4. Filters with gain taps can be thought of as just single-tap FIR filters. The feedforward tap is programmable and adapted during cancellation; the feedback tap is fixed and determined off-line.

Note that the DGTs act on the high-rate output of the ADC and their outputs are combined with the high-rate input to the DAC.




Figure 4. ANC system with high-speed analog and digital gain taps.

ANC ALGORITHM

The standard filtered-x LMS (FXLMS) algorithm was used to update the ANC coefficient for the feedforward cancellation,

$$h_{k+1} = h_k + 2 \mu \times e_k \times x'_k \quad (5)$$

where x'_k is filtered by the secondary path model. Other adaptive algorithms have been suggested for improved performance on fixed-point DSPs [6].

The modeling for the secondary path and feedback neutralization path is done off-line; then fixed versions are used in the active-cancellation mode. In addition, each microphone input is processed through an adaptive dc tap, and a leakage component is optionally part of the feedforward-path coefficient update algorithm.

ANC HARDWARE AND SOFTWARE REQUIREMENTS

The short-duct ANC hardware should include an AFE with at least two ADC channels and one DAC channel. The reference signal ADC and the antinoise DAC need to have inherently high sample rates and low group delay. The sampling timing of the antinoise DAC should be capable of being advanced relative to the sample timing of the reference ADC. The AFE should also have both high-speed analog and digital gain taps to provide an ultra-short delay path. The error-signal ADC also needs to be low in group delay, as its delay contributes to the delay through the secondary path as seen by the processor from the antinoise speaker to the error microphone. As this secondary path model has to be run by the processing block as well as the main feedforward path, it should be as short as possible. The main processing block should have as high a MIPS rate as possible (with an efficient instruction set) to reduce the delay, keeping within the general requirement for a low-cost solution. Finally, a single-package embodiment of the main signal conversion and processing functions should make the ANC solution more flexible yet cost-effective.

One such ANC solution with a single integrated circuit package can be obtained using the AD73522 dspConverter.

AD73522 PRODUCT INFORMATION

The AD73522 (Figure 5) is a single-device *dspConverter* incorporating a dual analog front end (AFE), a microcomputer optimized for digital signal processing (DSP) and a *flash*-based boot memory for the DSP.

The AFE section features two 16-bit ADC channels and two 16-bit DAC channels. Each channel provides 77-dB signal-to-noise ratio over a voiceband signal bandwidth with a maximum sample rate of 64 ksp/s. It also features an input-to-output gain network in both the analog (AGT) and digital (DGT) domains. The low group-delay characteristic (typically 25 μ s per ADC channel and 50 μ s per DAC channel) of the AFE makes it suitable for single- or multi-channel active control applications. The ADC and DAC channels feature programmable input/output gains with ranges of 38 dB and 21 dB respectively. An on-chip reference voltage is included to allow single-supply operation.

The AD73522's 52-MIPS DSP engine combines the *ADSP-2100 family* base architecture (three computational units, data address generators and a program sequencer) with two serial ports, a 16-bit internal DMA port, a byte DMA port, a programmable timer, flag I/O, extensive interrupt capabilities, and on-chip program- and data memory.

The AD73522-80 integrates 80 Kbytes of on-chip memory configured as 16K 24-bit words of program RAM and 16K 16-bit words of data RAM. The AD73522-40 integrates 40K bytes of on-chip memory configured as 8K words of 24-bit program RAM and 16-bit data RAM.

Both devices feature a Flash memory array of 64 Kbytes (512 Kbits) connected to the DSP's byte-wide DMA port (BDMA). This allows nonvolatile storage of the DSP's boot code and system data parameters. The AD73522 runs from a 3.3-V power supply. Power-down circuitry is inherent to meet the low-power needs of battery-operated portable equipment.

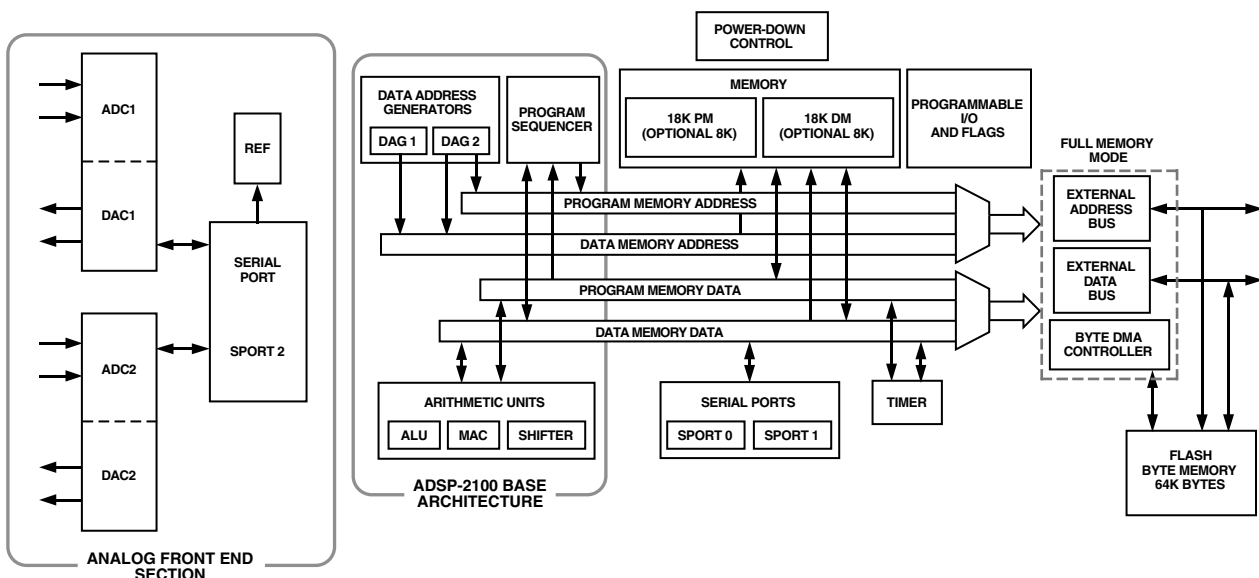


Figure 5. AD73522 dspConverter.

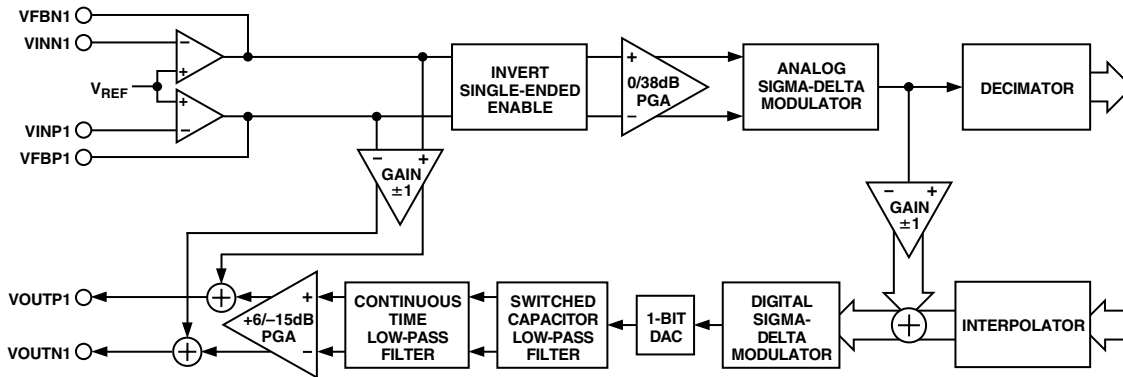


Figure 6. Analog front-end subsection of the AD73522 dspConverter.

SIGMA-DELTA ADC AND DAC ARCHITECTURE

The conversion technique adopted in the AFE is of the sigma-delta type. An analog sigma-delta modulator is used in the ADC channel and a digital sigma-delta modulator is employed in the DAC channel. A sigma-delta modulator is a heavily oversampled system that uses a low-resolution converter in a noise-shaping loop. The quantization noise of the low-resolution, high-speed converter is inherently high-pass filtered and “shaped” out of band. The output of the modulator or noise shaper is then low-pass filtered to reduce the sample rate and remove the out-of-band noise.

The AFE conversion channels used in the AD73522 are shown in Figure 6. The ADC section consists of an analog second-order, $32\times$ to $256\times$ oversampling, 1-bit sigma-delta modulator, followed by a digital sinc-cubed decimator (divide-by-32 to divide-by-256). The DAC section contains a digital sinc-cubed interpolator, a digital, second-order, $32\times$ to $256\times$ oversampling, 1-bit sigma-delta modulator, followed by an analog third-order switched-capacitor LPF and a second-order continuous-time LPF.

The group delay through the ADC channel is dominated by the group delay through the sinc-cubed decimator and is given by the following relationship:

$$\delta_{dec} = Order \times \frac{M - 1}{2} \times \delta_{ds} \quad (6)$$

where *Order* is the order of the decimator (= 3), *M* is the decimation factor (= 32 for 64-ksps output sample rate) and δ_{ds} is the decimation sample interval (= $1/2.048E6$ s)

$$\Rightarrow \delta_{dec} = 3 \times \frac{32 - 1}{2} \times \frac{1}{2.048E6} \quad (7)$$

$$\Rightarrow \delta_{dec} = 22.7 \mu\text{s}$$

for a 64-ksps output sample rate.

The group delay through the DAC channel is primarily determined by the group delay through the sinc-cubed interpolator and the group delay through the third-order switched-capacitor LPF. The inherent group delay through the interpolator is identical to that through the decimator and equals $22.7 \mu\text{s}$ for 64-ksps input sample rate. However, the interpolator can be optionally bypassed to

avoid this inherent group delay at the expense of reduced out-of-band rejection.

The z-transform of both the sinc-cubed decimator and interpolator is given by:

$$\left[\frac{1 - z^{-M}}{1 - z^{-1}} \right]^3 \quad (9)$$

The group delay through the analog section of the DAC is approximately $22.7 \mu\text{s}$.


Notice that with sample rates of only 8 ksps the inherent group delays through both the decimator and interpolator increase to $186.8 \mu\text{s}$. Therefore, it is very important to run the converters at as high a rate as possible to reduce the inherent group delay.

The AFE features high-speed analog and digital feedforward paths from ADC input to DAC output via the AGT and DGT respectively. The AGT is configured as a differential amplifier with gain programmable from -1 to $+1$ in 32 steps and a separate *mute* control. The gain increment per step is 0.0625. The group delay through the AGT feedforward path is only $0.5 \mu\text{s}$. The DGT is a programmable gain block whose input is tapped off from the bit-stream output of the ADC’s analog sigma-delta modulator. This single-bit input is used to add or subtract the digital gain tap setting, a 16-bit programmable value, to the output of the DAC’s interpolator. The group delay through the DGT feedforward path is only $25 \mu\text{s}$.


The loading of the DAC is normally internally synchronized with the unloading of the ADC data in each sampling interval. However, this DAC load position can be advanced in time by up to $15 \mu\text{s}$ in $0.5\text{-}\mu\text{s}$ steps. This facility can be used to further minimize the feedforward delay from analog input to analog output via the DSP.

AD73522 PACKAGING

The three main processing elements (AFE, DSP and Flash memory) are combined in a single package to give a cost-effective, self-contained solution. This single package is a 119-ball plastic ball grid array (PBGA), as shown in Figure 7. It measures $14 \text{ mm} \times 22 \text{ mm} \times 2.1 \text{ mm}$, and the solder balls are arranged in a 7×17 array with a 1.27-mm (50 mil) pitch.



a.



7 ROWS × 17 COLUMNS OF SOLDER BALLS

b.

Figure 7. AD73522 plastic ball grid array (PBGA) packaging.

AD73522 EVALUATION BOARD

The AD73522 dspConverter evaluation board (Figures 8 and 9) combines all of the front-end analog signal conditioning with a

user-friendly programming platform that allows quick and easy development. The board, interfacing to the serial port of a PC, comes with Windows® 95-compatible interface software that allows the transfer of data to and from all memory, including the Flash sections. All of the dspConverter pins are available at the output connectors. The board has an EZ-ICE® connector for advanced software development. Other features include a microphone with conditioning circuitry on one input channel and speaker amplifiers on the output channels.




Figure 8. AD73522 dspConverter evaluation board.

EXPERIMENTAL SETUP

The experimental setup (Figure 10) consists of a server box (containing just a fan and power supply), a plastic duct (with reference and error microphones and secondary loudspeaker), and the AD73522 evaluation board. The server fan was 5 inches (about 13 cm) in diameter. The T-shaped duct and the speaker measured 6 inches (about 15 cm) in diameter. The duct length was adjustable down to a minimum of 12 inches (30.5 cm).

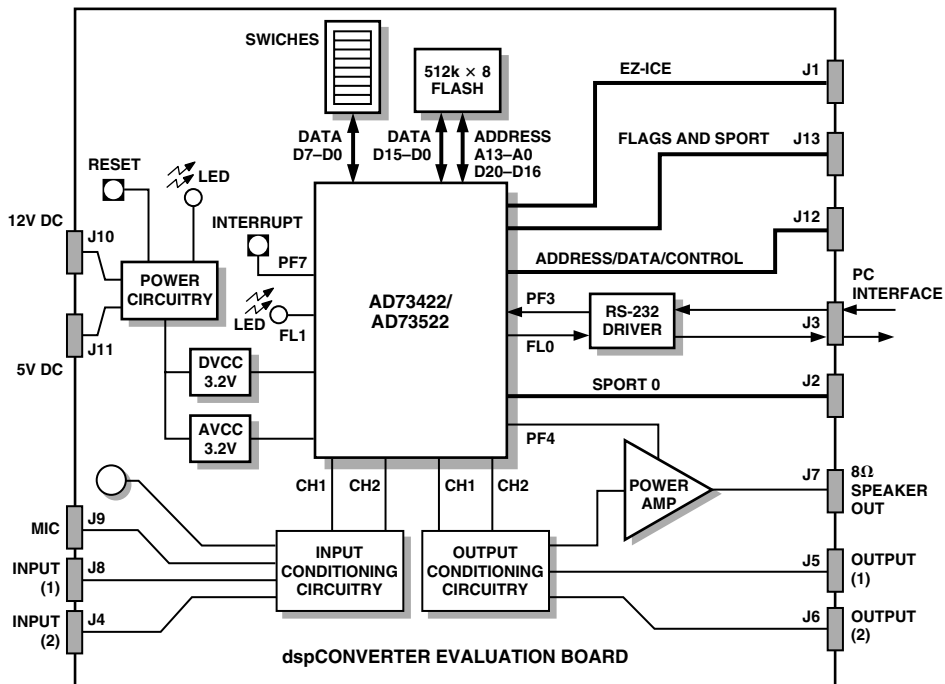


Figure 9. Block diagram of the AD73522 dspConverter evaluation board.

Windows is a registered trademark of Microsoft Corporation.
EZ-ICE is a registered trademark of Analog Devices, Inc.

During experimentation, the AD73522 evaluation board was hooked up to a PC for debug. Also, internal variables were written out to unused DAC channels for monitoring. Initially, the system was set up using a primary speaker instead of the actual server fan to allow testing with programmable tones and broadband signals.





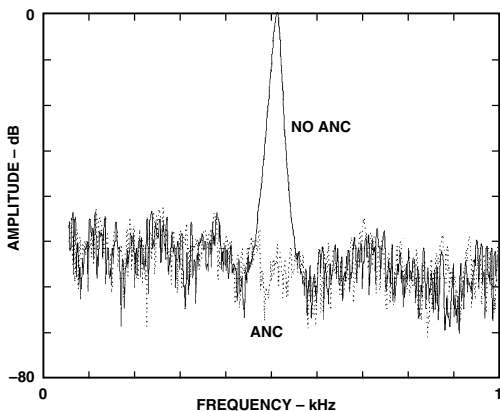
Figure 10. Server fan experimental setup.

RESULTS

The performance of the experimental setup with a single-tone disturbance from the primary speaker is shown in Figure 11. The main tone is reduced by a factor of 30 dB. When the primary speaker delivers a broadband disturbance, the reduction factor is about 20 dB, as illustrated in Figure 12.



a. 136-Hz performance, with and without ANC.



b. 510-Hz performance, with and without ANC.

Figure 11. Noise spectra with single-tone disturbance.




Figure 12. Performance with broadband disturbance.

CONCLUSIONS

An approach combining an analog gain tap (AGT) and digital gain tap (DGT) allows the use of sigma-delta techniques in low-group-delay ANC applications. A single-package embodiment combining analog and digital functions, like the AD73522 dspConverter, should provide an ANC solution that is both flexible and cost-effective.

ACKNOWLEDGEMENTS

Dell Computer, Limerick, Ireland, for the Poweredge 2200 server box used in the experimental setup.

Cork Institute of Technology, Cork, Ireland, for the duct construction used in the experimental setup and a kick-start in the code development.

The Institute of Sound and Vibration Research (ISVR) and Digisonix, Inc., for an excellent introduction to active noise control.

REFERENCES

1. S.M. Kuo & D.R. Morgan, *Active Noise Control Systems: Algorithms and DSP Implementations* (Wiley, New York, 1996)
2. "A Primer on Active Sound and Vibration Control," L.J. Eriksson, Digisonix brochure
3. "Application of Active Noise Control for Small-Sized Electronic Equipments," T. Hoshino, K. Fujii, T. Ohashi, A. Yamaguchi, H. Furuya, J. Ohga, *Inter-noise 94*, 1409-1412 (1994)
4. "Study of Active Noise Control System for Duct with High Speed Flow," T. Hoshino, T. Ohashi, J. Ohga, H. Furuya, A. Yamaguchi, K. Fujii, *Active95* (International Symposium on Active Control of Sound and Vibration), 423-430 (1995)
5. "Active Control of Fan Noise in a Personal Computer," M. O'Brien, P. Pratt, *IEEE ISSC 99 Digest*, pp. 385-392 (1999)
6. "An Adaptive Algorithm for Active Noise Control System by a Fixed Point Processing Type DSP," K. Fujii, A. Yamaguchi, H. Furuya, J. Ohga, *Active97*, 1163-1170 (1997)
7. "Short Duct Server Fan Noise Cancellation," P. Minogue, N. Rankin, J. Ryan, *Active99*, 539-550 (1999)