Application Note

AEROFLEX
A passion for performance.

Sigma delta IF A-D converters for digital radios

A.M. Thurston

The IFR 2309 incorporates 1 bit Sigma Delta IF A to D converter technology developed in conjunction with industry partners. This application note explains this technology and the implementations used in 2309.

This application note is a reprint of an article previously issued in The GEC Journal of Research 1995 vol 12 number 2.

The advantages of digitally encoding the IF signal in a heterodyne radio receiver are considered, and include improved discrimination between upper and lower sidebands, reduced spurious carrier component and the removal of flicker noise in the demodulated signal. Bandpass sigma delta converters are well suited to this task because of the narrow channel bandwidth to IF ratio and the requirement for high linearity. The concept of sigma delta conversion is explained at a fundamental level and examples of higher order advanced architectures are presented. A mapping between baseband and digital modulator architectures is presented, allowing bandpass converters to be designed from their baseband equivalents. Finally, a discrete component implementation of a converter is presented with an overview of the performance obtained in practice. The converters, implemented with discrete components and LC filter stages, have been successfully developed at the GEC-Marconi Research Centre and are currently being designed into a variety of receivers by several GEC product companies.

Keywords: bandpass; sigma delta; analogue-to-digital (A-D); intermediate frequency (IF)

INTRODUCTION

Almost inevitably, digital signal processing (DSP) techniques have found their way into the design of radio transmitters and receivers. Early applications were limited to the control circuitry, a good example being digitally-tuned synthesizer stages, allowing the removal of the rotary tuning dial and the inclusion of the numerical keypad or digital search facility. As DSP devices and dedicated ASICs have evolved, digital techniques have been used to perform the baseband signal processing functions, such as the final stages of channel filtering and demodulation. The early analogue-to-digital (A-D) converters had limited conversion speed and typically operated on either baseband or low-intermediate frequency (IF) signals. With the rapid development in signal processing capabilities over the last decade it is now possible to process higher IF signals digitally with a subsequent reduction in the analogue RF content within a receiver. Other advantages include: potentially enhanced performance through the use of digital filtering techniques; reduced production set-up time; and improved long-term stability. Although the location of a substantial number of high-speed digital switches alongside sensitive RF circuitry invites interference, the potential benefits are often considered to outweigh the new design difficulties.

Another problem introduced by the digital processing of IF signals is the need to perform high-speed A-D conversion; a problem compounded by the need for higher linearity in early stages of the receiver. Conventional multi-bit A-D converters have the property that the signal bandwidth available is equal to one half of the sampling frequency, less a margin to allow for anti-alias filtering. The product of the bandwidth and resolution of a converter is a measure of its performance, and this will typically be reflected in the difficulty of designing the device and also in its market price. Because a typical IF signal is narrowband compared to its carrier frequency, the use of wideband multi-bit converters does not represent an optimal coding solution to a very specific problem. Some reduction in the A-D converter's processing overhead can be achieved by operating it in a subsampled mode such that the carrier frequency is above the sampling frequency; however, this method requires enhanced channel filtering prior to the conversion to prevent other channels from aliasing into the passband.

Sigma delta converters, although traditionally operating on baseband signals - especially audio - exhibit attractive properties. First, they are an oversampled coding technique that achieves coding accuracy by fine temporal quantization rather than fine level quantization. Thus, for a given sampling frequency, the usable bandwidth is very much reduced compared with standard pulse code modulation (PCM) techniques, and this trade-off in requirements is reflected by a simplified design suited to low tolerance components. In a bandpass implementation the sampling frequency would typically be four times the IF, and the primary purpose of the anti-alias filter would be to attenuate the image signal at three times the IF, and those at other, higher, odd multiples. The analogue filtering required with such a converter is thus comparatively simple.

A second advantage of sigma delta coding is its inherent linearity. A multi-bit converter is very susceptible to component tolerances and a non-linear mapping between the analogue and digital domains is difficult to avoid. One very successful means of combating this effect is by the use of high-level additive dither, which effectively decorrelates the non-linearities from the input signal and reduces the effect to a benign noise source. This technique may be used to remove the non-linear effects from the coder, but the limiting performance is ultimately that of a PCM code, and this itself can introduce highly-correlated distortion, which in an application comprising evenly-spaced radio channels is likely to present difficulties. The decimated output of a sigma delta coder typically contains correlated distortion terms at a much lower level and is thus more suited to the receiver environment.

This article considers the advantages offered by the digital encoding of the IF signal in a radio receiver, and argues that a bandpass sigma delta A-D converter represents an optimal solution to narrowband IF coding. The fundamentals of sigma delta conversion are presented as an introduction before some of the more advanced coding techniques are reviewed. Sigma delta conversion has traditionally been performed at baseband, particularly for encoding audio signals, and hence a technique is described here to convert baseband architectures into their bandpass equivalents. Detailed description of the performance of the converters has been avoided in favour of a qualitative comparison with PCM coding techniques because this suffices to show the advantages of sigma delta when used in radio receivers, or similar IF coding applications.

DIGITALLY-IMPLEMENTED RECEIVER ARCHITECTURES

A heterodyne receiver architecture has been chosen to demonstrate the use of the bandpass sigma delta A-D converter in a digitally-implemented receiver. This is a commonly-used architecture, and serves as a good vehicle to demonstrate the use of sigma delta converters because it contains both high and low IF signals in addition to baseband in-phase and quadrature (I and Q) signals, allowing the relative merits of conversion at each frequency to be considered.

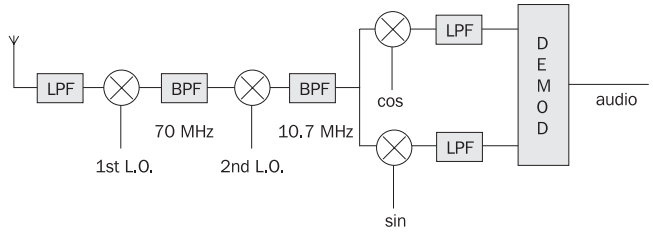

LPF=Low-pass filter

BPF=Band-pass filter

Fig. 1 Basic heterodyne receiver

Fig. 1. shows the basic heterodyne receiver architecture, configured here for use in an HF receiver. Signals are received at the antenna in the frequency range 3-30 MHz, with higher frequency signals attenuated by a low-pass filter. The required signal is mixed up to a first IF of 70.0 MHz with a variable-frequency local oscillator (LO), and the resulting signal is bandpass filtered prior to mixing with a fixed local oscillator down to the second IF of 2.5 MHz. More bandpass filtering is applied before the final mix to baseband orthogonal I and Q components. These are then low-pass filtered by the channel filters and applied to a demodulator.

To provide a degree of flexibility in the types of modulation that may be accommodated, the demodulation function is to be performed using a DSP, and a suitable analogue-to-digital interface is sought within the receiver. Three options will be considered: a pair of baseband converters operating on the I and Q baseband signals, and bandpass converters operating on each of the first and second IF signals.

The first option is depicted in fig. 2. The low-pass filters preceding the A-D converters are no longer performing the channel filtering function but are serving as anti-aliasing filters prior to the conversion process. The channel filtering is now performed digitally, giving generally superior performance and a degree of flexibility that is more difficult to achieve using analogue filters. The signal bandwidth required at each converter is one half of the channel bandwidth, and each converter is sampled at over twice this frequency in order to allow the anti-aliasing filter to roll off sufficiently. In such a scheme the dynamic range required of each converter could be made small by substantially rejecting adjacent channels using the second IF filter, combined with the effect of the anti-aliasing filters, and by providing suitable automatic gain control (AGC) earlier in the receiver.

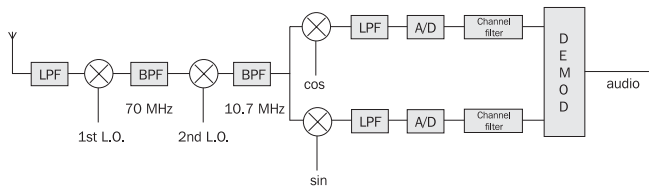

LPF=Low-pass filter

BPF=Band-pass filter

Fig. 2 Heterodyne receiver with baseband A-D converters

A number of sources of degradation exist within this receiver that will limit the available performance. Any phase error in the local oscillators used to mix the signal to I and Q baseband components will impair the receiver's ability to discriminate between components above and below the IF centre frequency. To achieve 40 dB of IQ discrimination requires these local oscillators to be orthogonal to within 0.5o, including all drift from ageing, temperature and manufacturing tolerances. This phase accuracy must then be maintained throughout the pair of analogue paths up to and including the A-D conversion function. Similarly, the amplitude response of the two analogue paths, including any gain mismatch between the two converters, must be well matched to preserve the IQ discrimination of the receiver. Again, to obtain discrimination of 40 dB it is necessary to match the amplitude response of the two paths to better than 0.1 dB. Such tolerances are possible and may be exceeded by using a calibration routine; however, obtaining this tolerance in a pair of digital paths is routine and provides some motivation for encoding an IF signal directly.

Further problems exist in this architecture. Any DC offset present at the input to the converters will, after demodulation, be indistinguishable from a tone at the centre of the IF, and may mask a real signal. Similarly, in a sensitive very-low-noise receiver, such as is required for radar, low frequency 1/f noise will be present and will resemble a region of high noise in the centre of the channel. Both of these effects can be removed by performing the conversion process at an IF.

Fig. 3 shows the receiver with the A-D converters replaced by a single bandpass device operating on the second IF. The sampling frequency is conveniently chosen to be four times the IF because this simplifies the subsequent mix to baseband by allowing the orthogonal local oscillators to be the repeated data sequences [1 0 -1 0] and [0 1 0 -1].

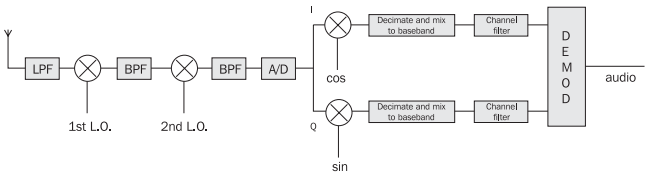

LPF=Low-pass filter

BPF=Band-pass filter

Fig. 3 Heterodyne receiver with second IF A-D conversion

The mixing process may most easily be implemented by time demultiplexing the encoded IF signal into the I and Q paths and then incorporating the sign inversions into the taps of the decimation filter. Because the mixing process and the baseband paths are implemented digitally, perfect orthogonality and phase-and-amplitude matching may be achieved routinely, and IQ mismatch is extensively removed. Similarly, DC offsets at the input to the A-D converter are mixed harmlessly out of the passband, as are any low-frequency noise components. The dynamic range required of the converter is generally slightly greater than would have been required of the baseband converters because the degree of adjacent channel rejection is slightly less without the contribution of the baseband anti-aliasing filters.

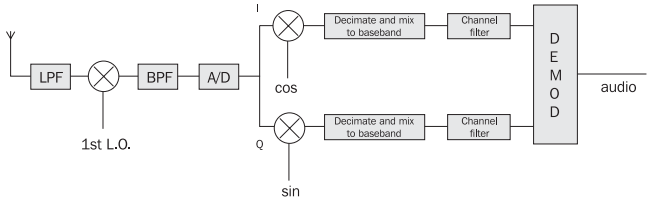

LPF=Low-pass filter

BPF=Band-pass filter

Fig. 4. Heterodyne receiver with first IF A/D conversion

Finally, fig. 4 shows the receiver configured with a bandpass converter in the first IF. The motivation here is to remove the second analogue IF stage, including both the filters and the mixer, and thus to simplify the overall design and to alleviate the problems associated with the second mixer. This is an ambitious conversion task with the IF at 70 MHz and a sampling frequency of 280 MHz. Though both conversion and decimation are possible at these frequencies, the task is made yet more difficult by the lack of adjacent channel rejection at the first IF. Many channels will now be encoded simultaneously, and the dynamic range required of the converter is typically some 20 dB greater than that required at the lower IF. Hence the compound requirements of both higher conversion rate and increased dynamic range make this converter rather more difficult (and consequently more expensive) to make. Moreover, the decimation circuitry will consume more power.

Moving the conversion process from baseband to the second IF has thus been shown to improve the IQ matching within the receiver, allowing greater discrimination between components above and below the IF centre frequency, and to negate the effects of DC offsets and low-frequency noise at the input to the converter. Moving the converter to the first IF requires far greater performance from the converter and the decimator, with a consequent increase in both power consumption and the cost of the technology. Whether this is negated by the convenience of removing the second analogue IF stage is specific to any particular receiver.

SIGMA DELTA FUNDAMENTALS

Sigma delta coders are noise-shaping modulators that reduce the number of quantizing levels required to represent a signal by greatly oversampling the signal. Accuracy in the amplitude domain is traded for accuracy in the time domain. In data conversion circuits this trade-off is generally favourable, and this is exploited in the design of sigma delta A-D converters. Because of the high oversampling frequencies required, they are particularly suited to low-frequency applications such as audio, and the large markets in this field have helped to drive the technology.

An example of a second-order baseband sigma delta modulator is given in fig. 5.

The basic elements are:

- the input port,
- the summing node which compares the input signal with a delayed version of the modulator's output and generates the difference,
- a loop filter in the feedforward path,
- an oversampled quantizer which in this case is single-bit and the output of which forms the output of the modulator, and
- a feedback path containing a delay equal to one sample period. The delay could equally be positioned in the loop filter, but must appear somewhere in the closed loop to make it realizable.

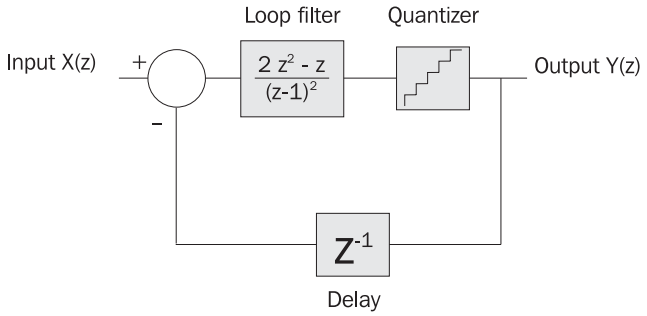

Fig. 5. Example of a second-order baseband sigma delta modulator

A simplified model of the modulator is shown in fig. 6. The loop filter is generalized to the filtering function A(z) and the quantizer is modelled as a unity gain stage with the addition of a white noise source n(z), representing the addition of quantizing noise. The input to the modulator is x(z) and the output is y(z).

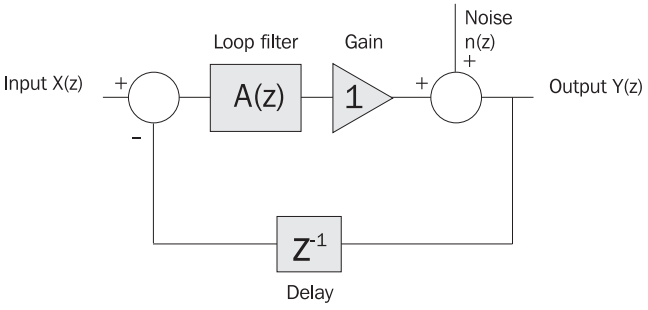

Fig. 6 Generalized model of a sigma delta modulator in which the quantizer is modelled by a unity gain stage with the addition of quantizing noise.

Transfer functions, describing the way in which the input signal and the quantizing noise are filtered at the output of the modulator, may be established using the linearized model of the quantizer. These are referred to as the Signal Transfer Function (STF) and the Noise Transfer Function (NTF) respectively, and are related to the inputs and outputs of the model by the equation

$$y(z) = x(z)\ STF(z) + n(z)\ NTF(z) \qquad (1)$$

The STF(z) is derived from the model in fig. 6 as

$$STF(z) = \frac{A(z)}{1+z^{-1}A(z)} \quad (2)$$

and the NTF(Z) is similarly derived as

$$NTF(z) = \frac{1}{1+z^{-1}A(z)} \quad (3)$$

It is seen that for a large value of A(z) the magnitude of the STF approximates unity, whilst the magnitude of the NTF tends towards 1/A(z). Thus the STF adopts a flat response and the NTF promotes attenuation of quantizing noise within the passband. A(z) is generally selected as a network of digital integrators, giving noise suppression at low frequencies. It is not possible to select a wideband function for A(z) with the intention of suppressing noise over a substantial portion of the operating frequency because such a filter would result in an unstable closed-loop response, and the denominators of both the STF and the NTF would contain poles outside the unit circle.

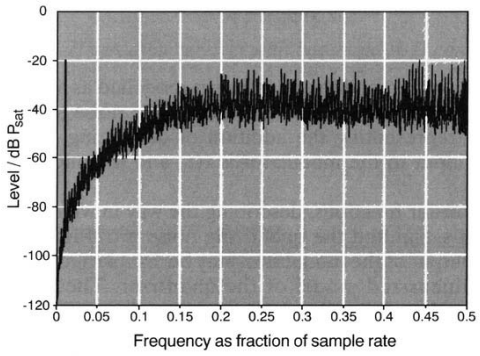

Fig. 7. Typical output spectrum of a second-order baseband sigma delta modulator

Fig. 7 shows a typical output spectrum generated using a simulation of the modulator depicted in fig. 5. The single tone on the left hand side of the display represents the input signal, whilst the rest of the display represents the single-bit quantizing noise shaped by the modulation process. The noise is seen to be suppressed at the low frequencies that constitute the passband of the modulator. Typically a passband of only 1/128 to 1/16 of the sampling frequency would be used, with the narrower passbands offering the higher dynamic ranges because of the enhanced noise suppression.

The order of a single-loop modulator is determined by the number of integrators in the loop filter. In the modulator shown in fig. 5 the loop filter contains a double pole at DC and thus the modulator is second order. The order of the modulator has a strong effect on the way it behaves, and this is now considered.

The noise power density within the passband is attenuated by the presence of zeros in the NTF, and by considering equation (3) this is seen to be determined by the poles of the loop filter A(z). A high density of zeros in the NTF results in a high degree of noise suppression, and generally it may be said that higher order coders result in lower inband noise power. However, this extra suppression is obtained at a price. First, the rate at which the noise rises away from the passband increases at 6 dB per octave with each increase in the order of the loop filter, and the requirements on the decimation filter become more severe. Stability also becomes more of an issue with higher order solutions. There was a time when single-bit converters were thought to be unstable above second order, although fifth-order coders are now routinely designed. The true penalty of higher order solutions is now considered to be a reduction in the overload point, because the onset of unstable behaviour occurs when large signals are applied. Hence, in an A-D converter, increasing the order of the modulation initially brings benefits because the quantization noise is reduced, but ultimately the dominant noise source becomes the circuitry itself, rather than the quantization effects; any further increase of the order serves only to lower the overload point and hence to reduce the dynamic range.

In general it is preferable to use a converter of order greater than two because the noise and linearity advantages of the higher order systems more than compensate for the additional complexity of their designs. The stability issue must thus be addressed, and in the next section some examples of the techniques that have been developed to stabilize higher order coders will be examined. All the techniques are applicable to baseband and bandpass coders alike.

ADVANCED SIGMA DELTA ARCHITECTURES

Three sigma delta architectures that represent the main methods employed to achieve stability in higher order systems will now be presented.

The first is commonly referred to as a MASH coder and employs a cascade of first or second-order loops in a noise-cancellation scheme. The error across the quantizer in the first loop is encoded by the second loop, and this is then subsequently subtracted from the digital output of the first loop to achieve cancellation. The technique may be used iteratively to provide further reduction in the noise simply by adding extra loops. The example given is by Uchimura et al.(1) and is restricted here to second order to simplify the explanation.

The second architecture represents the method used in many integrated devices, namely to provide overload detect circuitry which activates a hard reset of the integrators in the loop filter. The technique is rather crude, but if the input signal is limited prior to application to the converter then the occurrence of overloads is prevented, and the mechanism serves merely to reset the coder on power-up. The presence of the limiter however does limit the dynamic range slightly. For radio IF applications the presence of the limiter is undesirable, because ideal limiters are difficult to build, and the reduction of linearity that accompanies the addition of the limiter may degrade the overall system performance. Similarly, the use of overload detect and reset circuitry is intolerant to momentary overloads and noise bursts or rapid fading conditions could result in the converter frequently resetting, with a corresponding drop in performance. An example is given by Lee and Sodini(2).

The third class of coder employs nonlinear filtering techniques to enhance stability at overload. The basic principle is that, when

large signals are present in the filter, the net output of the higher order filter stages is clipped by a limiter and the converter reverts to a fundamentally lower order mode of operation. If carefully designed, the clipping noise can be shaped out of the passband by the same mechanism as that which shapes the quantizing noise. The general form of these coders has many independent variables and obtaining an optimum set of coefficients can be time-consuming. A simplified architecture which has considerably fewer independent variables, thus easing the optimization process(3), is presented here.

A MASH coder is presented by Uchimura et al.(1) and is shown in fig. 8. It consists of two first-order sigma delta coders where the input to the second coder is the error signal generated by the quantizer in the first coder. The output of the second coder is differentially encoded to compensate its shaped spectrum, and then it is added to the output of the first coder to achieve cancellation of the quantizing noise. The residual noise is that produced by the quantizer in the second loop, though this is heavily attenuated within the passband by the action of the differentiator, and thus a net improvement is obtained.

The scheme has the advantage that, because the input stage is first order, the overload point is maintained at a high level and hence the dynamic range is not compromized as in single-loop high order systems. The limitations of the technique stem from the fact that it employs cancellation and is thus sensitive to component tolerances. If for example there is a 1% error in the generation of the error signal across the quantizer in the first loop, then only 40 dB of noise suppression can be obtained, regardless of how many subsequent stages are used.

Other authors have considered the use of a second-order coder for the first stage, simply because the noise and distortion levels are very much lower, and if component mismatch does occur in the first stage its visibility at the final output is smaller(4). Additionally, because the output of the coder is derived from several sources, it is multi-bit, and consequently the decimation filter will have to employ multipliers. This contrasts with the single-bit coders, where only adders are required in the decimator, assuming a finite impulse response (FIR) filter is used.

Single-bit solutions can be achieved using a single higher order loop and addressing the stability issue accordingly. Possibly the first published account of a higher order loop came from Lee and Sodini(2) and is shown in fig. 9. A cascade of integrators is provided in the feedforward path and taps of suitable weight are combined from each stage and applied to the input of the single bit quantizer. It is predominantly the ratio in which these taps are combined that determines the stability, although the amplitude of the input signal is also largely influential, particularly near overload. A set of feedback taps is also included. These taps are used to move the poles of the integrators away from DC to form a low-pass filter with less ripple across the passband. These poles become the zeros of the NTF and thus distribute the in-band noise more evenly across the passband whilst giving an overall improvement in the coder's SNR.

This architecture is very flexible because each integrator's output has a separate tap weight before addition at the quantizer, and so the higher order outputs may be set at a sufficiently low level that stability is retained for a useful range of input signal levels. If, however, the input signal is taken too high, then an overload condition will occur and the coder will not automatically recover. This problem is generally approached by detecting the output of the integrators and, in the case of a large signal being detected, the integrator is reset. Such a radical attack on the information contained within the filters generates some considerable noise and it is found that the output of the coders is generally unusable under

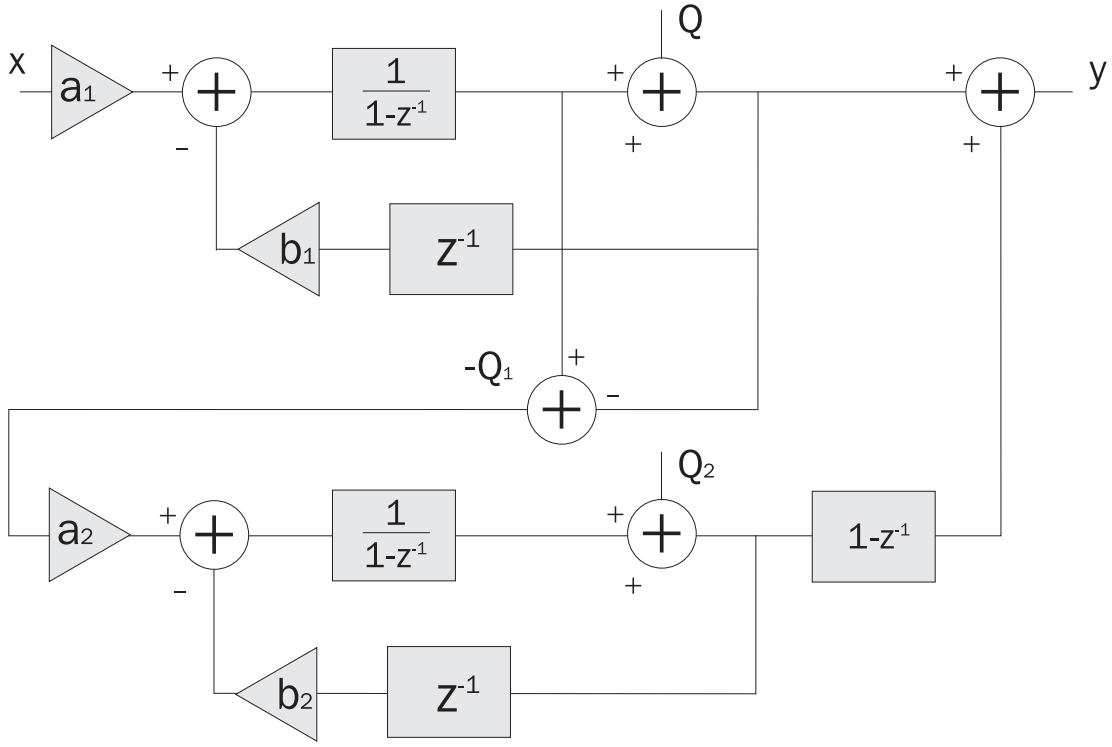

Fig 8. Multi-loop cascaded modulator by Uchimura et al. [1]

such conditions. The input signal level associated with this overload detect threshold is thus deemed the overload point of the coder. For radio applications, this is an unacceptably sensitive failure mechanism because the compound effects of multipath and other fading mechanisms, together with noise burst effects, could frequently overload the system. A limiter placed at the input to the converter could prevent this, but, for a wideband signal, any distortion introduced by this limiter could result in the raising of the system noise floor and possibly result in a blocking condition and the loss of communication.

The third scheme considered uses nonlinear filtering techniques to automatically clip large signals within the coder wherever they occur. By this means, the lower-order filtered components attain precedence in the overall coding scheme at times when the linearity is most threatened and the stability is thus enhanced. Fig. 10 shows a highly versatile architecture by Dunn(5) in which every integrator output is equipped with a limiter, allowing complete control of the gain and dynamic range of each filtered component at the input to the comparator. Using this arrangement it is possible to defer the point of overload of the converter and to enhance the performance above overload. Because of the sheer number of variables involved, optimization of such a coder is intensive. A rather simpler architecture containing only two variables is offered by the author in reference (3), and shown here in fig. 11. Its operation can best be understood by considering it as a nested pair of second-order modulators; because it is possible to make second-order modulators both stable for normal signal levels and also capable of unassisted overload recovery after the removal of an overloading input, the resulting fourth-order coder can be shown to have similar behaviour.

The fundamental elements of a second order coder are the second-order filter and the quantizer in the feedforward path, and the single delay element in the feedback path. The outer loop that bypasses the limiter represents a standard second-order sigma delta coder, and is thus known to be stable. When modulating under normal conditions, the quantizer can be tentatively modelled as a unity gain stage with the addition of quantizing noise. Using this model, the transfer function between points v and u can be shown to be a one-sample-period delay. The input u is effectively the delta modulation input to the sigma delta loop, and the delta coding is cancelled by the integrators of the outer loop before the signal is retrieved at output v. Thus the filter in the inner loop that encompasses the limiter is, to a first approximation, unaware of the existence of the filters in the outer loop because their filtering function is effectively negated by the action of the closed loop. Thus, provided that the outer loop is operating in a non-overload condition, the inner loop will exhibit second-order stability and overload recovery properties. The limiter is present so that after the removal of an overloading input the outer loop becomes the dominant signal source at the input to the quantizer, and the outer loop is guaranteed to recover. After this, normal second-order modulation ensues, the transfer function between v and u returns to a single delay, and the inner loop also subsequently recovers from the overload.

This particular method of overload recovery has advantages: first, because the output of the limiter is applied directly to the input of the quantizer, any noise introduced by the clipping effect is subject to the full noise-shaping properties of the modulation scheme, and thus is is comparatively silent compared to other techniques - particularly the overload detect/reset schemes. As a result of this, it is permissible that the limiter be frequently activated, and consequently the gain in the outer loop may be set at a high level that would result in unstable behaviour in the absence of the limiter. The additional noise introduced by the occasional

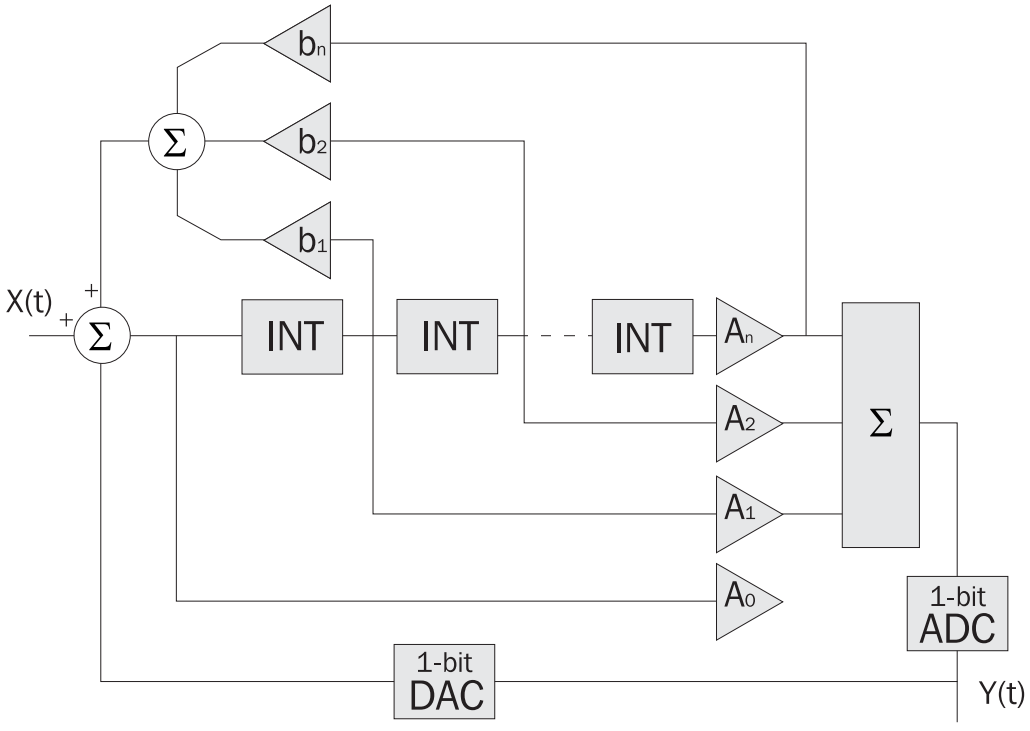

Fig. 9. Nth-order modulator by Lee and Sodini[2)]

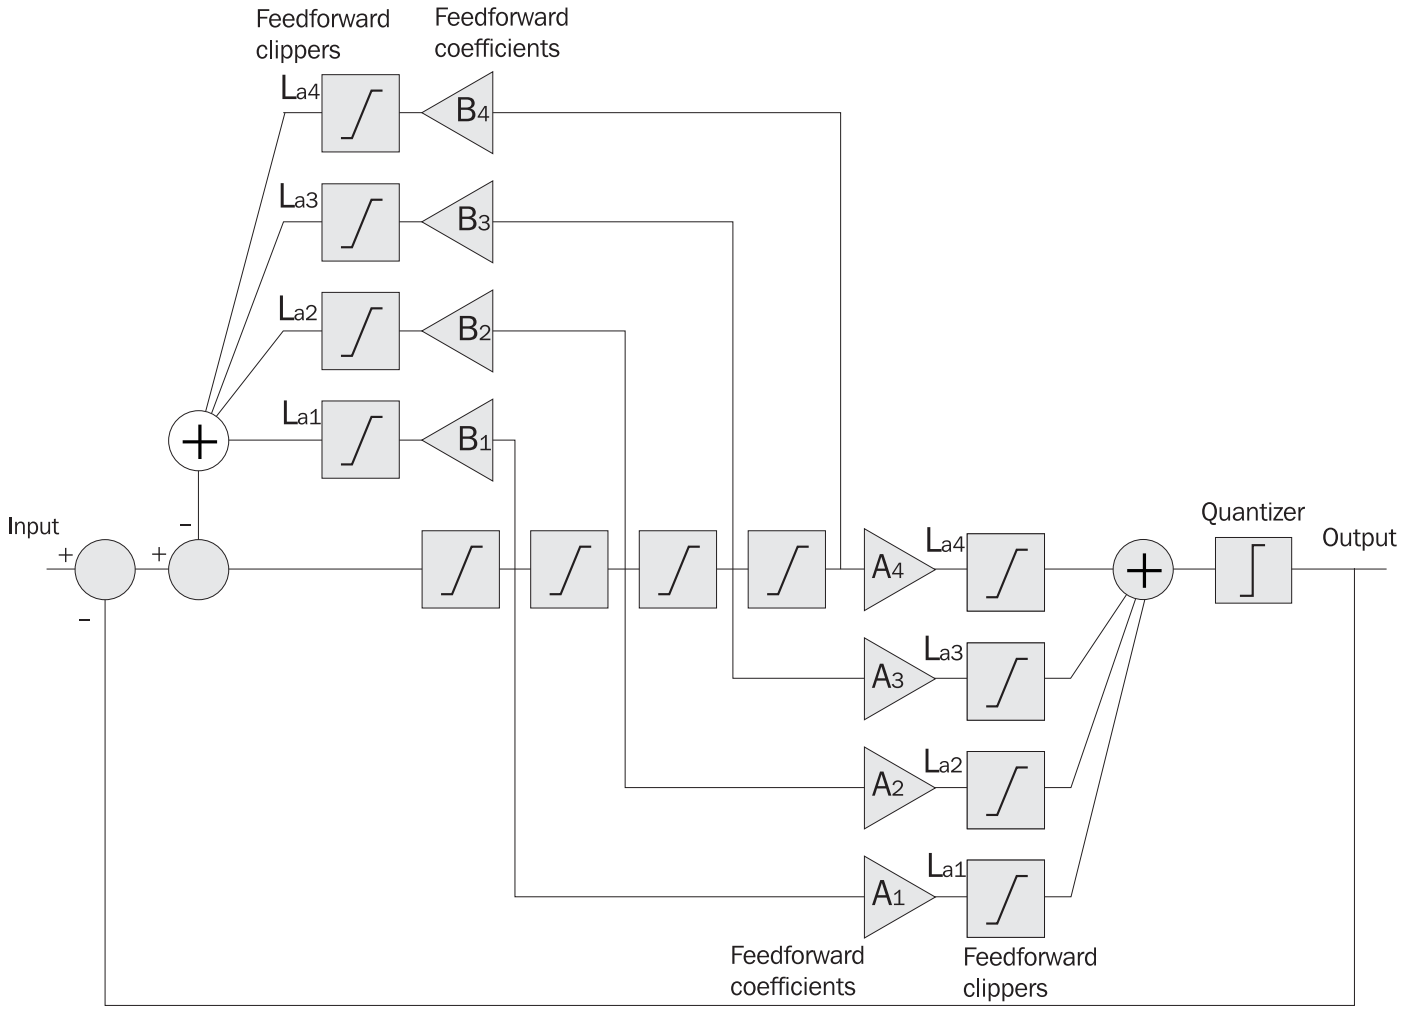

Fig. 10. Adaptive sigma delta modulator architecture by Dunn and Sandler[5]

clipping of the signal is low, and the extra attenuation of quantizing noise obtained by raising the gain of the higher order filter taps allows a few extra decibels of performance to be obtained. The final benefit occurs at high signal levels, where it is found that the overload characteristics of the coder are rather more gradual than the catastrophic failure associated with overload detect/reset coders, and extended operational range is achieved, albeit with reduced performance. The coder is found to be extremely tolerant to frequent low-level overload and is thus well-suited to use in a receiver application. A fuller account of the workings of this coder can be found in reference (6).

TRANSFORMATION FROM BASEBAND TO BANDPASS CONVERTERS

All the sigma delta converters considered so far have been baseband systems in which noise shaping is implemented typically in the audio frequency band by means of integrators in the sigma delta converter's loop filter. Noise-shaping is not constrained to occur at baseband and may, in fact, be applied at any fraction of the sampling frequency(7). One frequency is of particular interest, namely the one-quarter sampling frequency, because a converter that samples an IF signal at four times the IF produces an innately useful output. The encoded IF signal requires subsequent digital mixing to baseband in-phase and quadrature (I and Q) components with digital local oscillators. Provided that the sampling frequency is four times the IF, then the in-phase local oscillator is simply the repeated digital sequence [1 0 -1 0] whilst the quadrature local oscillator is the repeated digital sequence [0 1 0 -1]. These mixing functions can be included in the decimation filters and a considerable simplification of the digital filtering hardware can be achieved.

A simple transformation exists to convert a baseband sigma delta modulator into its quarter sampling frequency bandpass equivalent. The conversion can easily be envisaged by considering the pole zero diagram of the NTF of the modulator shown in fig. 5. By substituting -z for z in the modulator's architecture, the pole zero diagram is rotated by 180° about the origin of the Argand plane, thus translating the position of the noise-shaping zeros from baseband to the half sampling frequency (-1, 0j). A further substitution exists that translates the noise-shaping zeros from the half sampling frequency to the one quarter and three quarter sampling frequencies, namely the substitution of z2 for z. By means of this substitution the noise-shaping zeros are translated to the square roots of the original positions in the Argand plane, namely the one quarter (0, j) and three quarter (0, -j) sampling frequencies. The overall substitution is thus -z2 for z in the baseband modulator. A physical interpretation of this is that each delay element is

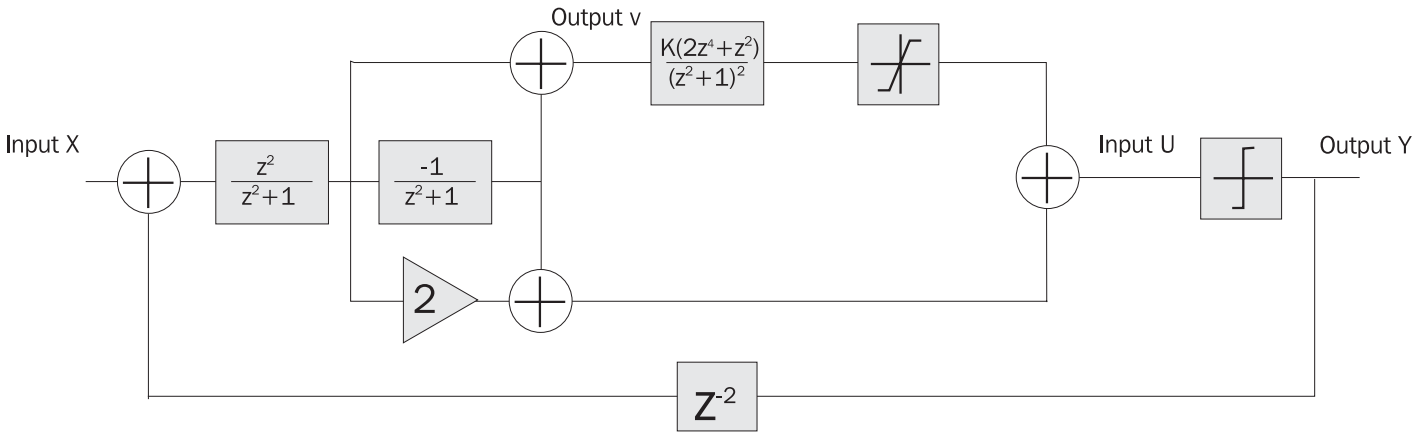

Fig. 11. Fourth-order bandpass sigma delta modulator with limited first, second, third and fourth-order components and outer second-order loop

replaced by a two-stage delay and an inversion. A practical example of this transformation can be found by making the above substitution in the baseband modulator of fig. 5, giving the new bandpass modulator shown in fig. 12.

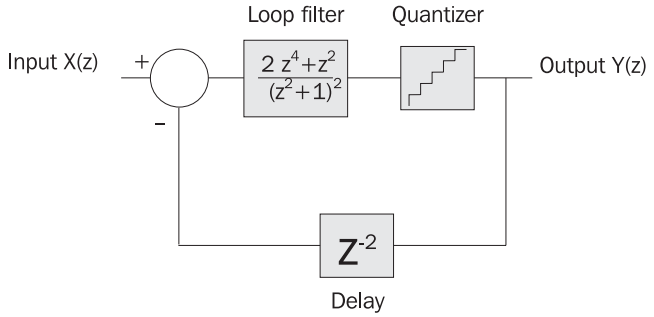

Fig. 12. Second-order bandpass sigma delta modulator obtained bysubstituting -z² for z in the baseband architecture

One aspect of the new modulator that often causes some initial confusion is the summation at the input node, rather than the subtraction normally found in a feedback loop. This is simply explained as a result of the IF operation, the sample rate at four times the IF and the two sample period delay in the feedback path. This delay term corresponds to a 180° phase shift at the IF, and hence cancellation is still achieved.

IMPLEMENTATION

The usual means of implementing sigma delta A-D converters has been with switched capacitor filters. In the case of an IF converter for a radio receiver, these have a number of disadvantages. First the noise figure associated with the switched capacitors is generally very high and this requires that the IF signal is amplified before application to the converter; this increases the power consumption and may impair the linearity. Second, because the IF signal is effectively sampled at the input to the converter, any noise introduced by the sampling process is indistinguishable from noise on the IF signal and cannot be shaped out of band by the sigma delta process. Finally the required linearity and speed required of the filters at higher IFs is currently not simultaneously obtainable with switched capacitor filters. Consequently, discrete component A-D converters based on standard RF circuit principles have been developed, where high frequencies and linearity are easier to obtain.

Discrete component realizations of the A-D converters using parallel LC filters have been produced encoding IFs up to 70.0 MHz and sampling up to 280.0 MHz, although generally lower IFs have been encoded because the usable dynamic range is not compromised by the distortion introduced by the very high speed D-A converters. The reaction that greets the idea of an A-D converter produced from discrete components is generally one of surprise. In practice, however, very high performances have been achieved, because - unlike multi-bit coders - the simple structure of sigma delta loops does not require large arrays of transistors or resistive dividing networks and the circuitry remains quite simple.

As an example, with an intermediate frequency of 2.5 MHz a dynamic range of 99 dB has been achieved over a passband of 50 kHz using purely discrete construction, with in excess of 100 dB spurious-free performance, when the noise is measured in a narrow bandwidth(8). Additionally, the use of continuous time filters allows very low noise figures compared with discrete time filters, and the converter just described is designed with an 18 dB noise figure, compared with the 39 dB noise figure of an example 16-bit converter, sampling at 100 kHz and overloading for an input of 2 V pk-pk. This has advantages in both power consumption and the preceding gain requirements. The small signal swings will be particularly useful in any system pushing towards the 3.3 V power supplies commonly sought in mobile communications applications. The disadvantage of the LC filters is that complete integration is not possible because of the inductors.

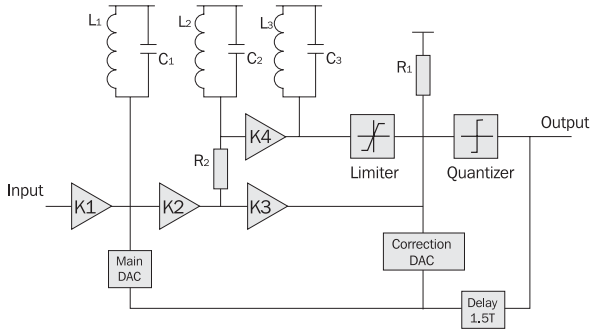

Fig. 13. Third-order bandpass sigma delta A-D converter

Fig 13 shows a typical discrete implementation of a third-order converter. Each LC filter is separated by a transconductance amplifier such that the voltage generated across a preceding filter stage is converted into a current input for the following stage. The quantizer is a sampled comparator, the output of which is the output of the converter itself. The output is delayed by one and one half sampling periods before application to the main digital-to-analogue converter (DAC) as full-width data pulses. The combination of the delay and the width of the resulting pulse from the DAC delays the centre of the DAC pulse by two sample periods after the sampling instant. A correction DAC is included that bypasses the LC filter stages. The purpose of this DAC is to supplement the output of the filter at each sampling instant to account for the fact that the finite-width pulse from the main DAC has imparted only one half of one quantum of charge into the filter.

Complete integration of the A-D converter becomes possible if the LC filters are replaced by state-variable filters. These have been integrated with IFs of up to 50 MHz and with quality factors up to approximately 60(9), although both of these figures may soon be improved. Typically a Q-factor between 100 and 200 is sought, although the upper limit on the IF is far in excess of typical second IFs and would probably be sufficient. The IF and Q-factor may be controlled via lines connected to pins on the integrated device. Whether this type of filter could be integrated to give the required linearity for use in a sigma delta application has not been investigated.

CONCLUSIONS

The advantages of encoding the IF signal in a heterodyne radio receiver have been considered. It has been argued that this approach will offer improved discrimination between upper and lower sidebands, reduced spurious carrier component and the removal of flicker noise in the demodulated signal. The concept of sigma delta conversion has been explained at a fundamental level and examples of higher order advanced architectures have been presented. Because most previous work has been carried out on baseband A-D conversion, a mapping between baseband and digital modulator architectures has been shown, allowing bandpass converters to be designed from their baseband equivalents. Finally, a discrete component implementation of a converter was presented with an overview of the performance obtained in practice. The converters, implemented with discrete components and LC filter stages, have been successfully developed at the GEC-Marconi Research Centre and are currently being designed into a variety of receivers by a number of the GEC product companies.

REFERENCES

1 Uchimura, K., Hayashi, T., Kimura, T. and Iwata, A., `VLSI A to D and D to A converters with multi-stage noise-shaping modulators', ICASSP 86, Tokyo, p. I545-1548, 1986.

2 Lee, W.L. and Sodini, C.G., `A topology for higher order interpolative coders', Proc. IEEE ISCAS 87, p. 459-462, May 1987.

3 Thurston, A.M. and Hawksford, M.J., `Dynamic overload recovery mechanism for sigma delta modulators', IEE 2nd International Conference on Advanced A-D and D-A Conversion Techniques and their Applications, Conference Publication No.393, p. 124-129, July 1994.

4 Ribner, D.B., `A comparison of modulator networks for higher-order over-sampled sigma delta analog-to-digital converters', IEEE Trans. Circuits and Systems, 38, 2, p. I45-159, Feb. 1991.

5 Dunn, C. and Sandler, M., `Use of clipping in sigma delta modulators', IEE Colloquium on Oversampling Techniques and Sigma Delta Modulation, Digest No. 1994/083, p. 8/1-8/9, 30th March 1994.

6 Thurston, A.M., `Bandpass Sigma Delta A-D Conversion for Application in Digital Receivers', Ph.D. thesis, University of Essex, 1994.

7 Schreier, R. and Snelgrove, M., `Bandpass sigma delta modulation', Electronics Letters, 25, 23, p. 1560-1561, November 1989.

8 Pearce, T.H., `Cascading bandpass sigma delta A-D converters', IEE 2nd International Conference on Advanced A-D and D-A Conversion Techniques and their Applications, Conference Publication No. 393, p. 1-6, July 1994.

9 Visocchi, P., Taylor, J., Mason, R., Betts, A. and Haigh, D., `Design and evaluation of a high precision, fully tunable OTA-C bandpass filter implemented in GaAs Mesfet technology', Submitted to the Journal of Solid State Circuits.

CHINA Beijing
Tel: [+86] (10) 6467 2716
Fax: [+86] (10) 6467 2821

CHINA Shanghai
Tel: [+86] (21) 6282 8001
Fax: [+86] (21) 62828 8002

FINLAND
Tel: [+358] (9) 2709 5541
Fax: [+358] (9) 804 2441

FRANCE
Tel: [+33] 1 60 79 96 00
Fax: [+33] 1 60 77 69 22

GERMANY
Tel: [+49] 8131 2926-0
Fax: [+49] 8131 2926-130

HONG KONG
Tel: [+852] 2832 7988
Fax: [+852] 2834 5364

INDIA
Tel: [+91] 80 5115 4501
Fax: [+91] 80 5115 4502

KOREA
Tel: [+82] (2) 3424 2719
Fax: [+82] (2) 3424 8620

SCANDINAVIA
Tel: [+45] 9614 0045
Fax: [+45] 9614 0047

SPAIN
Tel: [+34] (91) 640 11 34
Fax: [+34] (91) 640 06 40

UK Burnham
Tel: [+44] (0) 1628 604455
Fax: [+44] (0) 1628 662017

UK Stevenage
Tel: [+44] (0) 1438 742200
Fax: [+44] (0) 1438 727601
Freephone: 0800 282388

USA
Tel: [+1] (316) 522 4981
Fax: [+1] (316) 522 1360
Toll Free: 800 835 2352

AEROFLEX
A passion for performance.

www.aeroflex.com
info-test@aeroflex.com

Our passion for performance is defined by three attributes represented by these three icons: solution-minded, performance-driven and customer-focused.

Part No. 46891/867, Issue 2, 06/04