COMBO II™ CODEC/ FILTER—Who Says Analog Can't Shrink?

National Semiconductor
Application Note 503
Chris Stacey, Roy Batruni and
Jim Weiser
December 1987

IS ASP STILL ALIVE?

Many circuit designers in the Telecommunications industry may have heard the term ''Digital Signal Processing'' (DSP) so many times in recent years that perhaps they now assume that ''Analog Signal Processing'' (ASP) is a dying art. COMBO II, a second generation programmable PCM Codec/Filter combination from National Semiconductor, should set the record straight on that score, however. Built on an advanced 2 micron* double-metal CMOS process, COMBO II employs all the same analog techniques used so successfully in the TP3054/7 series of COMBO™ devices. Analog circuit blocks such as switched-capacitor filtering and charge re-distribution D/A converters have been scaled down to take advantage of the smaller geometries provided on this process. This has enabled a number of additional line card functions, previously implemented with external passive components, to be economically included with the PCM Filters and Codec on the device.

COMBO II = COMBO + . . .

Foremost among these additional functions is a programmable hybrid balance filter. Hybrid balancing is Telecom jargon for the cancellation of echo from the 4-wire receive path towards the 4-wire transmit path introduced by the 2-to-4 wire conversion, either via a transformer or a monolithic SLIC (Subscribers Line Interface Circuit). So, too, the gain in both the transmit and receive directions can be programmed, over a range of 25.4 dB in 0.1 dB steps. Several digital functions on COMBO II are also programmable, including time-slot and port assignment and 6 general-purpose I/O latches. In all, there are 10 write-able and read-back-able program registers, accessed through a standard serial control interface. *Figure 1* shows a block diagram of the device.

ASP OR DSP? FOR AND AGAINST

When the design of COMBO II was embarked upon, the pro-DSP ball was already rolling, and beginning to gather momentum. The writing seemed to be on the wall for an Analog Signal Processing (ASP) approach. After all, weren't 2 and 1.5 micron dual-metal CMOS processes developed to optimize VLSI digital circuits? And doesn't DSP reduce the analog complexity to that of a simple over-sampled delta-modulation type of A/D and D/A converter at the front-end of the device? Certainly there is some substance in these statements, and there is little doubt that when a filtering function is very complex, or requires an adaptive transfer function, DSP is generally the technique of choice.

Not that DSP is without its drawbacks, however. For example, a DSP codec requires building an A/D converter, sampling in the 0.5–1 MHz range, with the necessary dynamic range to meet PCM specifications for noise and linearity, which is a considerable challenge in itself. Achieving a good

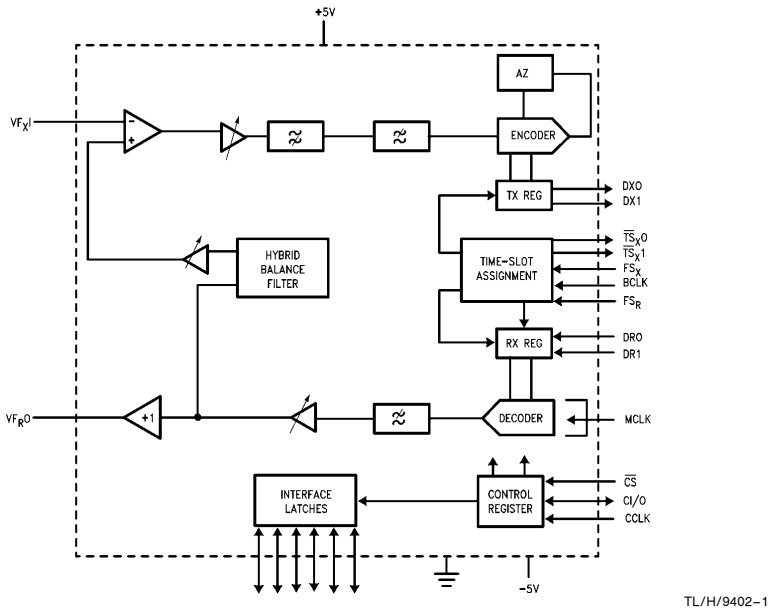

FIGURE 1. COMBO II = COMBO I + Functions in Bold Lines

*2 microns is the minimum ''drawn'' dimension; it is roughly equivalent to what is commonly called a 1.5 micron (effective channel length) process.

RRD-B30M115/Printed in U. S. A.

Power Supply Rejection Ratio (PSRR) with a wideband A/D is also difficult. Furthermore, studies revealed that digital decimation and PCM channel filters with the necessary number of significant bits would actually occupy considerably more die area, and consume more power, than a switched-capacitor solution.

Thus, the challenge was clear: could switched-capacitor filters and capacitive A/D and D/A techniques be scaled down to the advanced process while retaining the low noise and high PSRR of the original COMBO designs? The goal would be both better performance and less cost than a DSP solution.

SURPRISE! YOU CAN HAVE YOUR CAKE AND EAT IT TOO!

In many respects it is the advanced high-density processing technology which, rather than making life difficult for an analog design, actually comes to the rescue. For example, one of the key essentials for good transmission performance on COMBO II has been the development of a high-quality capacitor on the 2-micron process. The requirements for coefficients of voltage and temperature are very stringent, and the influence of parasitic elements on the analog switches must be carefully controlled. Devices such as the original COMBO, typically built on processes in the 3 to 5 micron range, utilized a double-polysilicon process with inter-poly oxide capacitors to implement accurate and consistent filters. Now, a single implant has been added to the standard M^2CMOS process to produce a high quality capacitor between poly and diffusion. This capacitor, with twice the capacitance per unit area as in the older double poly process, allows the filter to be designed in only one half the die area. Scaling also improves capacitor matching, since the poly is a critical process layer which made it difficult to control the ratio of capacitors.

In order to yield high density VLSI circuits, wafer processing must be meticulously clean, a factor which is a boon to analog designers, since it produces lower 1/F noise in MOS transistors. Process accuracy also produces better matched transistors and thus low offset op-amps. Voltage offset and noise of the amplifiers is further minimized by advanced circuit techniques. And there's more yet: other process advantages include smaller parasitics due to scaled interconnect, lower charge injection due to smaller switches, and better performance op-amps. Reduced parasitics and lower charge injection go a long way to ensuring good power supply rejection from both supplies.

Thus, with direct-step-on-wafer exposure, and the ultra-clean conditions so essential for good yields on a high-density process, the performance of these analog elements, as measured on actual COMBO II devices, is proving to surpass that of the current generation of codec designs. That certainly dispels one of the myths about ASP! Die-hard DSP advocates would have you believe that there is some fundamental law of physics which renders this impossible to achieve. The truth of the matter is that there is no such physical limitation. Taking advantage of the fact that the PCM channel filters and hybrid balance filter are low-order fixed, rather than adaptive, functions, and applying these ASP techniques, COMBO II has actually resulted in an economical combination of more functions than current designs without compromising performance.

COMBO II: HOW IT'S DONE

Let's take a closer look at the architecture of COMBO II, beginning at the transmit input. Analog voice signals enter at the V_{FXI} pin to a summing function, at which echo cancellation can be accomplished when using the hybrid balance filter circuit. Following this is a digitally programmable gain block, which provides from 0 dB to a maximum of 25.4 dB of gain in 0.1 dB steps. This is implemented by using polysilicon resistors as the gain-setting elements in a three stage configuration, illustrated in *Figure 2*. These polysilicon resistors are tapped logarithmically, and are well controlled due to the stringent etching control in a high-density CMOS process. Such a tapped resistor approach is well suited to logarithmic steps, while capacitor arrays are generally better suited to binary-weighted steps. A 3-stage configuration was chosen to minimize the number of op-amps, while still achieving the desired accuracy with the least number of resistor taps. Infinite input resistance CMOS op-amps allow the use of MOS transfer gates as tap selection elements with no effect on performance. Following the gain block the signal passes through an 8th order switched-capacitor bandpass filter which meets CCITT and LSSGR PCM specifications. Again, thanks to the 2 micron process, the result is a scaled down, high-performance filter with low noise and offset in a surprisingly small area.

NOW FOR THE HARD PART!

Probably the most critical section in a COMBO, and without doubt the heart of it, is the A/D encoder which, to meet CCITT and other specifications, must convert the output of the transmit filter at an 8 kilosample/sec rate, with 8 bits/sample. Because of the need to preserve signal integrity for even small amplitude samples of the signal, however, a compression of the output codes must be included, wtih an equivalent resolution of 13 bits for small amplitude samples. Two standardized coding characteristics are specified by CCITT, known as μ-255 or A-law, and COMBO-II allows for the selection of either law by programming control bits.

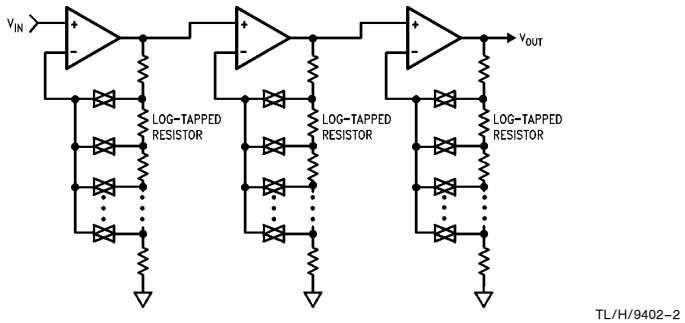

TL/H/9402–2

FIGURE 2. Programmable Gain Block

Meeting CCITT performance specifications demands at least a 78 dB dynamic range in the A/D converter. This range, coupled with the 25 dB programmable gain required for different applications, needs a total system dynamic range of 103 dB! For a DSP approach this would require a front end A/D with 17-bit resolution. That's quite a tall order for any A/D, especially if it needs to be monolithic and cost-effective. In practice, DSP codecs have far less resolution and linearity than this, both in their A/D converters and the arithmetic circuits, with a consequent penalty of degraded noise and distortion performance when gain is programmed in. But a low-noise programmable gain amplifier preceeding a 78 dB successive-approximation A/D is readily achievable, and allows gain to be added without the dB-for-dB performance degradation of a DSP codec, as illustrated in *Figure 3*.

Just as in the original COMBO, COMBO II uses a charge-redistribution capacitor array approach in the A/D, gaining the same benefits as the filter from the high quality of the capacitors on this process. Appropriate weighting of the capacitor values also provides a convenient and low cost method of μ or A-law compressing in the A/D, which, in a purely DSP device needs to be implemented by a separate digital processing look-up table. Completing the A/D is a strobed comparator, which senses at high speeds with modest power thanks to the high speed process. Again transistor matching and small switch transistors help to keep offset and noise low. Of course, all control logic and the digital auto-zero integrator take full advantage of the high density of the 2 micron process.

The output of the A/D encoder is an 8-bit word at the standard 8 kHz sample rate, which is then parallel loaded into a shift register and serially shifted out at a bit rate from 64 kb/s to 4.096 Mb/s. Transmit data can be shifted out at the instance of a controlling frame sync signal or assigned to any one of up to 64 8-bit time slots in a frame, depending on program bytes sent to the device. Two-port switching is also available, since the digital output can be switched to either of two output driver ports, D1 or D0.

Likewise, serial PCM data in the receive direction can be switched from the system into the device through one of two input ports, DR1 or DR0, again selectable by a programmable bit. Time-slot assignment, in which the data is selected from any one of up to 64 time-slots in a frame, can be programmed independently from the transmit side by a similar programming byte. PCM data is then parallel loaded to a latch to be used by the expanding D/A converter.

This D/A expansion in the decoder follows a complementary law to that in the encoder. The output of the D/A is a voltage appearing on the top-plate of its capacitor array which is buffered by a unity gain non-inverting amplifier. A voltage reference, which is used by both the decoder and encoder, is generated from a bandgap type circuit. N-well CMOS processes produce a vertical PNP transistor which is reasonably well controlled. Since the collector is actually the device substrate, and must therefore be connected to the most negative potential, the reference is isolated from the positive supply, which is generally the noisiest in any mixed digital/analog system. The resistors which set the current densities of the PNPs are again polysilicon, to provide high accuracy matching and an order of magnitude lower temperature coefficient than diffused resistors. A low offset op-amp is used which allows sufficient reference accuracy with minimal trim. With this bandgap voltage reference, the D/A produces an accurate stepped analog version of the digital words it receives as input.

Next, this output from the decoder is fed to a 5th-order switched-capacitor lowpass filter to recover only the audio part of the spectrum, and correct for the sinX/X attenuation characteristic (high-frequency ''droop'') of the sample-hold decoder output. This receive filter gains the same advantages as the transmit filter from process scaling, making it an efficient analog signal processor.

Following the receive filter is a digitally programmable gain block with an attenuation range of 0 dB to -25.4 dB in 0.1 dB steps. As in the transmit side, the combination of D/A and 25 dB attenuator range demands a 103 dB dynamic range, which would need to be provided by the reconstruction D/A alone in a DSP approach, a formidable objective!

After this attenuator the signal splits to drive two blocks, the hybrid balance filter and a class AB power driver, capable of driving up to 2.45 Vrms into a 300Ω load with < -60 dB THD. High density processing helps the power driver design

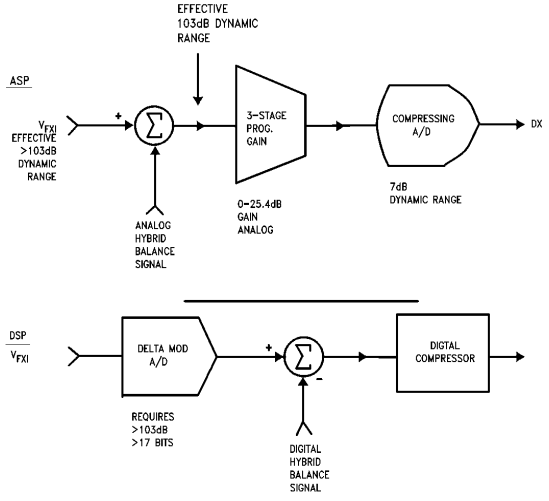

TL/H/9402–3

FIGURE 3

by allowing shorter channel output drive transistors, which thus produce more drive per unit area than was possible on earlier processes.

A HYBRID BALANCING ACT

Undoubtedly the most complex new feature on COMBO II, the hybrid balance function has been designed to meet the demanding International specifications not only of local loops, but also Special Service circuits and Trunk Interfaces. Transformer and monolithic SLIC (Subscriber Line Interface Circuit) are equally well provided for, with no need for any external hybrid balance components such as the R-C networks usually used in current designs. Moreover, the performance of this filter enables considerably better hybrid balancing to be achieved than is generally obtainable with simple R-C designs.

To understand the reasons why hybrid cancellation is so critical to good speech performance in the telephone network, *Figure 4* illustrates a typical single-link connection. Network switching and transmission require the conversion of the signal path from a 2-wire subscriber loop to a 4-wire path at the central office, thereby separating the directions of transmission. The conversion junction, commonly known as a hybrid, can consist of a transformer or an electronic SLIC. Practical hybrids, however, are imperfect signal separators, so that a portion of the signal from subscriber 1, intended for subscriber 2, couples across this junction back

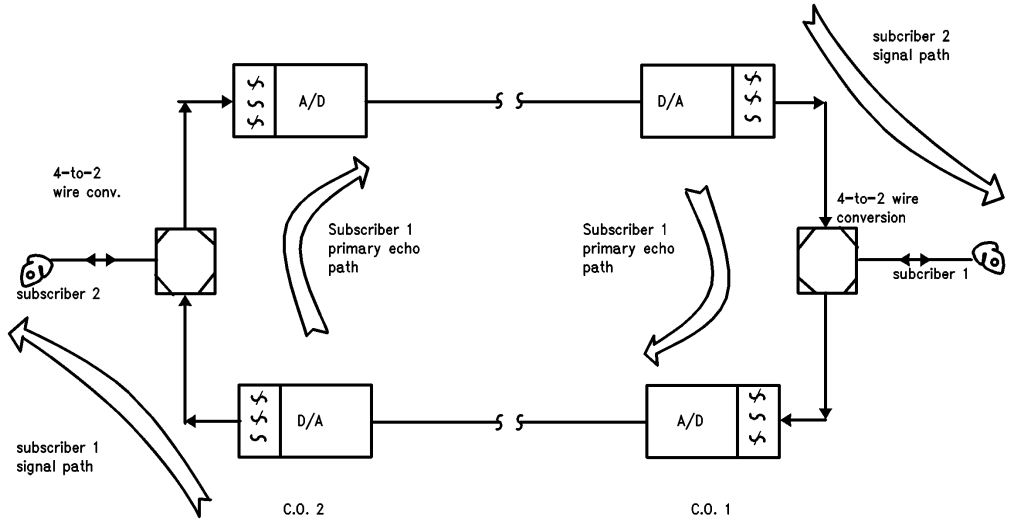

TL/H/9402–4

FIGURE 4. Subscriber Loop Signal Paths

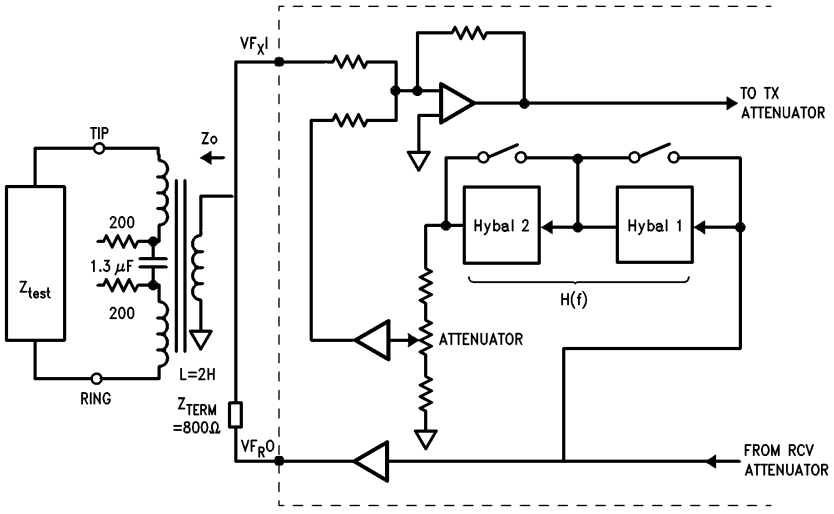

TL/H/9402–5

FIGURE 5. Example of a Typical Transformer Line Interface Implementation. Z_{test} is Shown in *Figure 6*.

to subscriber 1, its originating source. This far-end echo not only degrades the quality of speech but may also result in instability of the overall connection, since the echo at one junction may couple back at the other junction again, and unless there is sufficient attenuation in the 4-wire paths it can prevent a conversion from being intelligible. ''Balancing'' the hybrid refers to the minimization of this echo signal by adding a special circuit which improves the 4-wire signal separation.

To gain a better appreciation of the demands on the balancing circuit, consider the 4-to-2 wire conversion junction shown in *Figure 5*, where a transformer line interface is used. This is a realistic implementation of a transformer SLIC, which satisfies the 2-wire termination impedance specification.

Letting the input impedance of the subscriber loop in *Figure 5* be Zo(f), the transfer function between the unwanted echo V_E and the input V_R is given by the simple voltage-dividing action of the 2 impedances:

$$R(f) = \frac{Zo(F)}{Zo(f) + Zt(f)} \qquad \text{Equation 1}$$

In a typical line card application, the termination impedance Zt(f) is fixed by the Telephone Administration or PTT to ensure good sidetone and transmission level performance in the telephone set. Zo(f) is a function of the transformer size and construction and the a.c. and d.c. termination impedances, as well as the characteristics of the particular subscriber loop, such as length, gauge constituency and presence or absence of loading coils. As a consequence of the line impedance's wide variation across subscriber loops, telephone administrations ordinarily substitute a test network for the real subscriber loop when testing the degree of echo cancellation. Somewhat curiously, these test networks are generally different for each country, and the networks for the USA are shown in *Figure 6*, along with balance specifications attached to each network. The requirements regarding hybrid balance appearing in BellCore specification LSSGR-7 require the ratio 20 $\log(V_X/V_E)$ to be higher than that shown in *Figure 6*. This ratio simply measures the energy of the residual (uncancelled) echo as a fraction of the energy in the source signal in dB. A high ratio indicates better cancellation or balancing.

In Equation 1, the function R(f) is called the echo path frequency response, and it determines the strength of the echo signal V_E. The echo path therefore produces an undesirable signal V_E due to a signal V_R. If one attempts to synthesize R(f) with external R-C networks, practical difficulties arise. Typically the number of poles and zeros in R(f) are a function of the reactive components in the SLIC and test network. In *Figure 5*, for example, three poles and three zeros are present in the echo path. Tolerances of the reactive and resistive components, however, accompanied by the variability of the transformer inductance as a function of the d.c. feed current, cause movement in the expected values of the poles and zeros. This puts stringent requirements on balancing the SLIC with an external R-C network in terms of the number of discrete components required and their tolerance.

ASP WINS OUT AGAIN

On COMBO II, hybrid balancing is accomplished not by an impedance, but by means of a filter H(f), whose input is the same as the echo path R(f). If the filter H(f) is made equal to R(f) at all frequencies then their outputs are equal, so that when they are subtracted the resultant signal V_X is zero, i.e. the echo is cancelled. The advantage in this approach is gained by use of high precision switched capacitor filters in implementing H(f), and the low cost of making it of a high order and programmable. A considerable degree of programmability is provided in this filter, enabling users to tailor

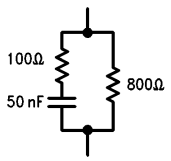

Non-loaded lines test network, Z_{test1}.

TL/H/9402–6

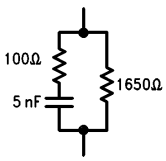

Loaded lines test network, Z_{test2}.

TL/H/9402–7

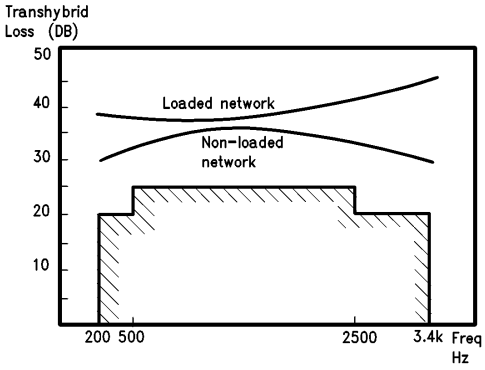

TL/H/9402–8

FIGURE 6. Balance Test Networks (U.S.) with Balance Filter Performance when Z_{test1} and Z_{test2} are used in *Figure 5*.

the response to accurately match that of their SLIC and termination circuits. Both transformer and electronic SLICs can be accurately balanced thanks to this programmability.

To make this complex problem more manageable, the balance transfer function H(f) on COMBO II is implemented as a switched-capacitor filter consisting of 3 sections, which can be factored as:

$$H(f) = G \times H1(f) \times H2(f) \qquad \text{Equation 2}$$

Here, the term G is the gain stage multiplicative factor selected by the user, which can take any of 25 values covering a 6 dB range in 0.25 dB increments. Its function is to perform signal level matching between H(f) and R(f), whereas H1(f) and H2(f) are programmed to provide frequency response matching. The H1(f) term is given by:

$$H1(f) = \frac{f\,(f - z1)}{(f - p1)\,(f - p1)} \qquad \text{Equation 3}$$

which is a biquad filter section with a fixed zero at 0 Hz, a programmable low frequency zero, and a pair of complex conjugate low frequency poles. Altogether there are 256 possible configurations for H1(f) under user control. Referred to as Hybal 1, this section is particularly suited to balancing transformer SLICs and other terminations having a dominant low frequency pole, typically lower than 500 Hz.

At these frequencies the problem of balancing the echo path is a common one, the reason being the considerable effect of the SLIC components whether a monolithic or transformer SLIC is used on the echo response R(f). Consequently, on COMBO II special attention has been paid to the frequency range of 200 Hz to 500 Hz, and the balancing problem separated into a low and a high frequency (greater than 500 Hz) problem. Yet another advantage of ASP over DSP becomes apparent here and it is this: the low frequency poles in the echo path imply an excessively long time domain impulse response due to the long time constant and slow decay they generate. In a digital filter it becomes impractical to provide enough coefficients to meet such a long impulse response, and external R-C components are usually necessary to improve low frequency hybrid balancing.

This is entirely avoided on COMBO II by the ability to program in actual poles, rather than trying to synthesize them with a large number of zeros, as is typically done in a digital filter.

The third section of the hybrid balance filter response H(f) is H2(f):

$$H2(f) = \frac{1 - (f/Z)}{1 - (f/P)} \qquad \text{Equation 4}$$

Here the approximate range of Z is 1 kHz to 80 kHz, and P is between 1 kHz and 10 kHz, and again there are around 256 possible configurations for H2(f). Referred to as Hybal 2, this filter is optimized to give fine control of the H(f) response at frequencies above about 500 Hz. For the most part this filter section balances the effects of the test network itself, although the large selection of pole/zero values was chosen to give it balancing ability that is equivalent to or exceeds what is possible with a digital filter approach. One example of a feature not possible with a digital filter is the ability to program poles and zeros with values above the voice-band (4 kHz), which have a critical influence within the voice-band. These are quite often needed for optimum balance in practice, and since a digital filter uses a PCM signal sampled at 8 kHz it cannot implement any pole or zero outside 4 kHz by virtue of the Nyquist sampling theorem.

Implementing the hybrid balance filter in an analog signal processing architecture gains two other advantages over DSP approaches. Firstly is the avoidance of potential input overload problems resulting from the unwanted echo being superimposed on the transmit input signal. In COMBO II, cancellation takes place at an input summing junction, so that the signal dynamic range is reduced before it passes through any active circuitry. In contrast, the front-end A/D in a DSP codec is further strained by the need to cope with this additional dynamic range, which is nominally 6 dB. As a consequence of this the second big advantage is attained, which is that the echo replica signal is not susceptible to any A/D non-linearities. Finite quantization steps in the A/D, and inevitable deviations from the nominal input/output characteristic introduce non-linear distortion that a linear echo canceller filter cannot handle. Since the hybrid balance filter on COMBO II cancels the echo before it is quantized by the A/D, the problem is avoided.

All this is not to say that an ASP design is a piece of cake! At the circuit level, the performance goals for these filters impose more severe requirements on amplifiers and other components than do the other on-chip filters. Placing a low frequency pole as low as 28 Hz in the bi-quad requires large capacitor ratios which need to match to better than 0.5%. Also, because of the very high effective resistances emulated by the switched capacitors (in the region of 350 MΩs!), leakage and charge injection from the switches must be kept low to avoid large offsets. Once again, the advanced process comes to the rescue, with its low leakage needed to produce micro-power digital logic, and very low gate overlap capacitances needed for density and speed. Both of these characteristics help to keep the offsets low, a vital factor in such a low frequency switched-capacitor filter. In the high frequency pole-zero stage, a wide-band, low-noise amplifier is required, which the process can also provide. Accuracy is very critical for good hybrid balance, and again the high performance capacitors pull through. Since all signals pass through capacitors in these filters, it is relatively straightforward to add transfer gates where necessary to program capacitors in or out, thereby changing the response, or bypassing the filters altogether.

To show just how well COMBO II's hybrid balance filter performs, the echo-cancellation of the interface from *Figure 5* was entered into the COMBO II Hybrid Balance Design and Optimization Program, which runs on a personal computer and performs joint optimization of Hybal 1 and Hybal 2 over the entire voiceband. With 256 Hybal 1 and 256 Hybal 2 filter functions to choose from, the optimization program will pick one of the 65,536 filter transfer functions available on COMBO II to give the best possible echo cancellation. Plotted in dotted lines in *Figure 6* are the balance filter's matching response of each echo frequency response, showing that margins of 5 dB to 10 dB are readily achievable even at the most difficult frequencies. In *Figure 7*, the echo frequency response with these networks is plotted in a polar (complex variable) format which captures magnitude and phase information in one plot. The magnitude of the response at a certain frequency is the distance from the origin to the point on the plot representing the response at that frequency; the phase is the angle between the horizontal axis and the line connecting the origin and the response point.

VIVE LA DIFFERENCE!

Designing COMBO II, with its stringent demands on analog performance, began as a considerable challenge. By successfully meeting this challenge with the application of Analog Signal Processing techniques on the most advanced high-volume CMOS process, COMBO II provides Telecom

designers with the next logical step in reducing both the cost and space of their line interfaces. It is also undoubtedly breathing a new lease of life into the world of ASP. Certainly DSP codecs will continue to be given a good run for their money!

ABOUT THE AUTHORS

Chris Stacey comes from London, England, and holds the B.Sc. degree with honors in Electronic Engineering from the University of Southampton in 1969 and the M.Sc. degree in Telecommunications Technology from the University of Aston in 1972. From 1969 to 1978 he worked with the Plessey Company on the design of PCM channel banks, multiplexers and digital line transmission systems. In 1978 he joined National Semiconductor Corporation in Santa Clara, California, where he has specialized in system-level definition of advanced telecommunication integrated circuits. He is currently Telecom Applications Manager.

Roy Batruni comes from Lebanon, and holds the B.S. degree with honors in Electronic, Computer and Systems Engineering from Loughborough University, England, in 1979. He also holds the M.Eng. degree from Cornell University, which he received in 1980. Since then, he has been with National Semiconductor's Telecom Group specializing in the field of communications using adaptive and optimal filtering techniques. He is the author of several pending patents.

James Wieser comes from Detroit, MI., receiving the B.S. degree and M.S. degree in Electrical Engineering from the University of Michigan, Ann Arbor in 1977 and 1978 respectively. In 1978 he joined the Telecom Group at National Semiconductor Corporation in Santa Clara, CA., where he is now an Engineering Section Manager. He specializes in CMOS analog circuit design and is the holder of several patents in this field.

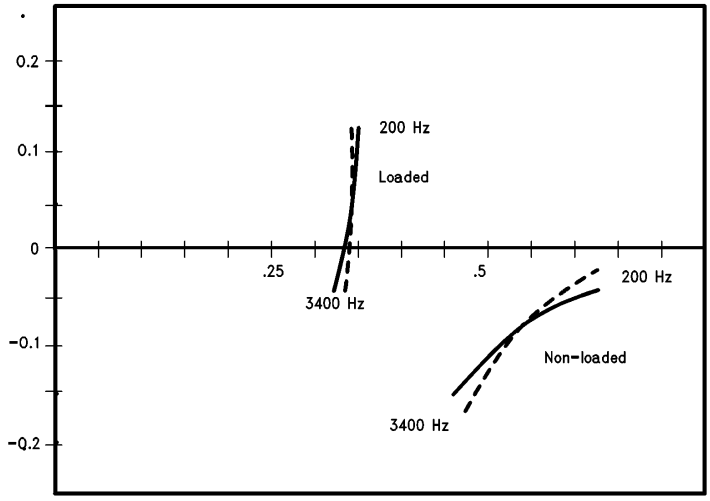

TL/H/9402–9

— Echo Path Response.

. . . Balance Network Response.

FIGURE 7. Frequency Responses of Echo Path and Balance Network for Non-Loaded Test Network and Loaded Test Network in *Figure 6.*

LIFE SUPPORT POLICY

NATIONAL'S PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES OR SYSTEMS WITHOUT THE EXPRESS WRITTEN APPROVAL OF THE PRESIDENT OF NATIONAL SEMICONDUCTOR CORPORATION. As used herein:

1. Life support devices or systems are devices or systems which, (a) are intended for surgical implant into the body, or (b) support or sustain life, and whose failure to perform, when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.
2. A critical component is any component of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.

National Semiconductor Corporation
1111 West Bardin Road
Arlington, TX 76017
Tel: 1(800) 272-9959
Fax: 1(800) 737-7018

National Semiconductor Europe
Fax: (+49) 0-180-530 85 86
Email: cnjwge@tevm2.nsc.com
Deutsch Tel: (+49) 0-180-530 85 85
English Tel: (+49) 0-180-532 78 32
Français Tel: (+49) 0-180-532 93 58
Italiano Tel: (+49) 0-180-534 16 80

National Semiconductor Hong Kong Ltd.
13th Floor, Straight Block,
Ocean Centre, 5 Canton Rd.
Tsimshatsui, Kowloon
Hong Kong
Tel: (852) 2737-1600
Fax: (852) 2736-9960

National Semiconductor Japan Ltd.
Tel: 81-043-299-2309
Fax: 81-043-299-2408