

Area-Efficient Multichannel Oversampled PCM Voice-Band Coder

BOSCO H. LEUNG, MEMBER, IEEE, ROBERT NEFF, PAUL R. GRAY, FELLOW, IEEE, AND
ROBERT W. BRODERSEN, FELLOW, IEEE

Abstract—This paper describes the design of a four-channel oversampled A/D converter with transmit filter for voice-band application. The decimation filter is time-shared between the four channels and the architecture of the sigma-delta coder is selected on the basis of minimizing the chip area. The analog front-end loop is fully differential to minimize the channel-to-channel crosstalk. The key issues in designing multichannel oversampled ADC for area efficiency will be addressed. The proper choice of the coder as well as the filter architecture are discussed. It is concluded that for the multichannel telephony voice-band application implemented in CMOS technology, a first-order loop is most area efficient. The performance of the coder has been evaluated and it has a dynamic range of 79 dB, occupies a total active area of 33 000 mils² or 8250 mils² per channel, and meets the D3 specifications for the transmit filter. It runs on a 5-V supply and consumes 50 mW per channel.

I. INTRODUCTION

BECAUSE of their compatibility with scaled digital processes, oversampled techniques are particularly attractive for the implementation of A/D interfaces for voice-band telecommunication applications. In comparing the oversampled techniques with the conventional analog approach in this application, the key issue is that of minimizing the total chip area. Oversampled coders require a large percentage of the circuit area to implement the digital filter. These digital circuits can easily take advantage of technology scaling and, perhaps more importantly, they can be shared without crosstalk. Considerations of die area create a strong motivation to use a multichannel time-shared filter on a single chip to minimize the per-channel area. Crosstalk between analog circuits from different channels can be reduced by using separate analog front ends, which adds little extra area because the analog circuits make up only a small percentage of the total chip

Manuscript received July 11, 1988; revised August 29, 1988. This work was supported by DARPA, the California MICRO Program, Bellcore, Fairchild, AMD, Rockwell International, and Honeywell.

B. H. Leung was with the Electronics Research Laboratory, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720. He is now with the Department of Electrical Engineering, University of Waterloo, Waterloo, Ont., N2L 3G1, Canada.

R. Neff was with the Electronics Research Laboratory, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720. He is now with IBM, San Jose, CA 95193.

P. R. Gray and R. W. Brodersen are with the Electronics Research Laboratory, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720.

IEEE Log Number 8824594.

area. In the past, considerations of crosstalk and die area have precluded a multichannel single-chip approach, and previous implementations of both sigma-delta and successive-approximation coders for this application have used a single-channel architecture [1]. In this paper, an experimental four-channel oversampled voice-band coder in a 3- μ m, single-poly, double-metal CMOS technology is described. This prototype achieves a dynamic range of 79 dB, idle channel noise of 12 dBnC0, channel-to-channel crosstalk of -83 dB, power supply rejection ratio of 44 dB at 1 kHz and more than 60 dB at 20 kHz, and has a frequency response that meets the D3 specifications in a silicon area of 8250 mils² per channel. The total power consumption is 200 mW, or 50 mW per channel. This power consumption comes mainly from the separate analog loops, and can be powered down if the specific channel is not in use.

This paper consists of three additional parts. In Section II, architectural considerations unique to the design of a multichannel coder are addressed. Section III goes into descriptions of the various circuits used in the prototype. Finally the experimental results will be discussed in Section IV.

II. CHOICE OF CODER IMPLEMENTATION FOR MULTICHANNEL OVERSAMPLED PCM CODERS

For voice-band telephony applications, first- and second-order sigma-delta coders are of the most practical interest. These coders are shown schematically in Fig. 1. A sigma-delta loop consists of one or two integrators and a 1-bit A/D and 1-bit D/A. The loop generates a 1-bit code whose density is proportional to the input level. A digital low-pass filter averages this signal and generates a digital output. As shown in Fig. 2 both loops have an output quantization noise spectrum that has a high-pass characteristic.

In Fig. 2, f_s is the input sampling frequency, f_D is the downsampling frequency, and f_o is the highest frequency in the passband. The oversampling ratio D is defined as the ratio f_s/f_D .

In per-line components for voice-band telephony, per-line component cost is a critically important parameter.

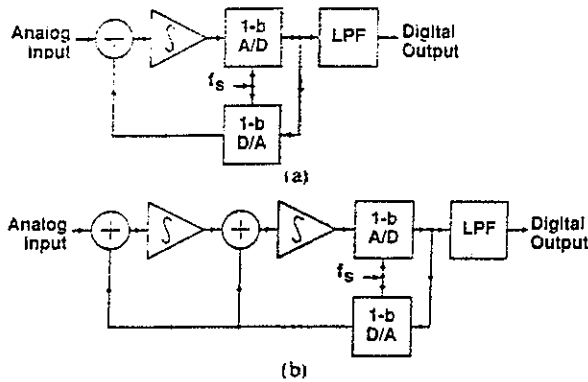


Fig. 1. (a) First-order coder. (b) Second-order coder.




Fig. 2. Noise spectra of sigma-delta coders.

An important objective of the work described here was to reduce the silicon area per channel in the coder function to the lowest possible value. While second-order sigma-delta coders are most frequently used in single-channel applications because a lower sample rate can be used, the realization of multichannel PCM interfaces in a single chip brings about a different set of trade-offs in silicon area minimization which tend to make a first-order implementation more attractive than would be the case for a single-channel coder. As part of this work an area comparison was made to determine the relative area consumptions of first-order and second-order coders in a multichannel implementation.

Since the decimation filter takes up the vast majority of the area of the overall coder, the comparison of the filter area required in the two types of coders is of central importance. The decimation filter has three functions:

- 1) out-of-band quantization noise suppression;
- 2) out-of-band signal anti-aliasing; and
- 3) maintaining the passband ripple to be less than the value dictated by transmission requirements, usually 0.25 dB for typical telephony codecs.

Since the band-edge frequency of the voice channel is only 3 kHz, the digital filtering required to maintain the passband shape can be most efficiently carried out at low sample rate after decimation. As a result this circuitry can be time-shared over the multiple channels and has an area contribution that is small on a per-channel basis. The requirement for signal anti-aliasing can be considerably simplified by incorporating a multistage filter. The second-stage filter can use decimation so that it can then be time multiplexed. Consequently the area of the circuitry




Fig. 3. Out-of-band quantization noise suppression.

required to do signal anti-aliasing can be made relatively small.

The out-of-band quantization noise suppression must be performed at the input sampling frequency, and as a result the hardware required to achieve this suppression cannot be easily time multiplexed in the technology used for this work. This in turn implies that a critical issue for the resulting area per channel in the final implementation is the means of realizing this initial decimation function. As shown in Fig. 3, the noise density from the sigma-delta loop increases with frequency. This is filtered by the first-stage filter but after decimation the unfiltered out-of-band noise is aliased into the baseband. The filter requirement is to reduce this noise to a level that is small compared to the quantization noise that is in-band before decimation. The area of the filter to achieve this is strongly influenced by the choice of the order of the coder.

A realistic comparison of the area required for the decimation filter required for first- and second-order coders is quite complex and must necessarily be made under a set of assumptions that might be violated in some innovative future design. In the work described here, such a comparison was made in order to compare the areas of such filters as they would be implemented in the prototype described later. These assumptions are as follows.

a) In the telephony codec application, the first-order coder implementation utilizes a sinc^2 filter with an L/D of 2, a decimation ratio of 512, and a sampling rate of 4 MHz. The second-order coder utilizes a sinc^3 filter with an L/D of 3, a decimation ratio of 128, and a sampling rate of 1 MHz [2]. These filter configurations and sample rates meet the dynamic range and idle channel requirements of PCM telephony with comparable margins. These particular FIR filter configurations have been previously shown to reduce out-of-band aliased noise to a value that is small compared to in-band noise for the two coder cases. Other decimation filters that meet the telephony transmission requirements are certainly possible. For example, use of a second-order loop operating at a relatively high sampling rate with a sinc^2 filter gives a coder whose in-band noise is dominated by aliased out-of-band noise which is capable of meeting the noise requirements. Such configurations do not appear to have great promise for having minimum area and were not investigated further.




Fig. 5. Filter structure.

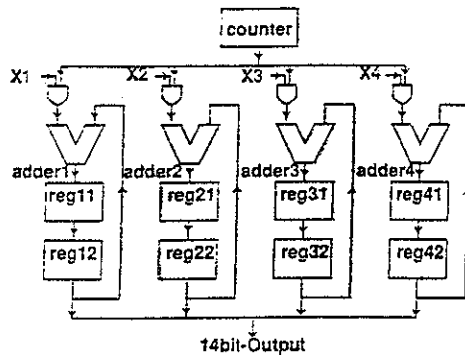


Fig. 6. Block diagram of FIR1 filter.

passband ripple of less than 0.25 dB. As discussed in Section II, an L/D of 2 is required to ensure that the out-of-band quantization noise will be sufficiently suppressed. To satisfy the other two requirements the filter would need a higher L/D ratio. Since the FIR filter computes at the high input sampling frequency f_s , its complexity should be minimized. By breaking the filter into multiple stages, the requirements can be met while keeping the L/D ratio of the first-stage FIR at 2. The overall filter structure is shown in Fig. 5 where an FIR filter with a triangle impulse response takes a 4-MHz 1-bit input data. The 14-bit output is decimated to 32 kHz and fed into a second-stage FIR filter. The output is further decimated to 16 kHz and fed into the IIR filter where it is decimated to 8 kHz. FIR1 is implemented in a custom structure. FIR2 and IIR are implemented by a time-shared microprogrammed structure.

Next the prototype custom FIR1 design will be discussed. Fig. 6 shows the block diagram of the FIR1 filter. For this four-channel filter one counter provides all the coefficients. Although an FIR decimation filter that does not require an impulse response generator exists [4], [5], this approach does not have a particular advantage in area because the generator is shared over different channels. Furthermore, this approach does not lend itself easily to reduction in area by time multiplexing the data path without introducing extra registers and multiplexers. As shown in the figure, X_1 – X_4 are the 1-bit output code from the four analog front ends. The distinct feature of this approach is that the adder in each channel is multiplexed. As will be shown later, careful design of the control structure allows this to be implemented in a simple fashion. The detailed implementation of the 1-bit slice of the adder and the registers is shown in Fig. 7. Here PC_0 and PC_8 are the LSB and MSB from the program counter, which is an up-counter. PC_i is the i bit of the counter. The




Fig. 7. One-bit slice of FIR1 filter.




Fig. 8. FIR2 and IIR filter implementation.

PC_8 signal is used to invert the counter output after it counts D samples. It converts the up-counter into an up-down counter. As discussed in Section II, since L/D equals 2, two sets of coefficients are needed. If two accumulators are used these two sets of coefficients have to be provided simultaneously. Since only one adder is used, they will have to be provided in consecutive clock cycles. Furthermore, since the impulse response is triangular, the two sets of coefficients are just the inverse of one another, leading to further simplification. In the circuit the PC_0 signal is used to invert the counter output between every clock cycle to generate these two sets of coefficients. There are altogether four registers. Two of the registers are used to hold state variables for the addition. The other two registers are used to hold the decimated outputs. These outputs are then latched and transferred to the next stage filter. The control signals are generated by ANDing a latch signal and the two-phase clock. The latch signal happens once every D samples and it latches the new output as well as pushing the old output down like a FIFO. The reset signal is used to reset the FIR filter after the output has been latched.

The IIR filter and the FIR2 filter are both implemented using a microprogrammed architecture. The FIR2 filter is used to provide anti-aliasing around 16 kHz and the IIR filter is used to maintain the passband ripple as well as the stopband attenuation. The IIR filter has to compensate for the droop in the passband due to FIR1 and FIR2 filters while maintaining a ripple of less than 0.25 dB. In the stopband the suppression must be more than 33 dB at 4.6 kHz and beyond. The microprogram structure as shown in Fig. 8 is composed of the program counter (PC), the ROM for the program, the RAM for the state variables, the arithmetic unit and I/O unit (AUIO) for computations,

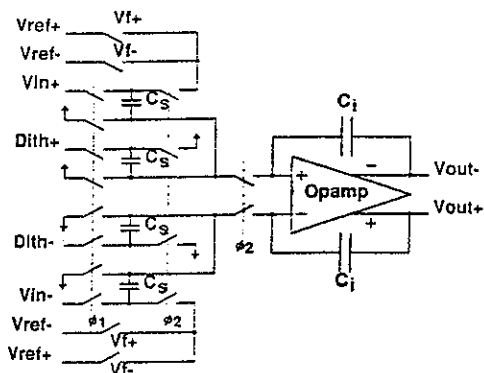


Fig. 9. Prototype switched-capacitor integrator.

and the address arithmetic unit (AAU) for calculating addresses for state variables that belong to different channels. The basic hardware is generated by the Lager system [6], a silicon compiler that accepts an assembler program for the filter as the input and generates the layout of the filter. The hardware is heavily pipelined to increase throughput. To save area no hardware multiplier is used, and multiplications are achieved by shifting and adding. In order to reduce the number of shift and add cycles, thus improving the speed of the IIR filter and allowing for more multiplexing, the IIR filter coefficients are optimized by a program Candi [7] that generates the canonical signed digit (CSD) code with the fewest number of bits in the filter coefficients. The program uses a data path of 20 bits which ensures that the idle channel noise is low enough even in the presence of limit cycles. The ROM is 50 words by 26 bits and the RAM is 40 words by 20 bits. The microprogrammed structure runs at 4 MHz. In addition special modifications are introduced to achieve area efficiency for the present application. In particular the AAU is replaced by simple logic to calculate the addresses of the state variables for the different channels.

The switched-capacitor first-order analog front-end integrator is shown in Fig. 9. To minimize the crosstalk, four separate analog front ends are used. With the digital filter being integrated on the chip the power supply rejection is also a primary consideration. A fully differential architecture is used to reduce the crosstalk and improve the power supply rejection. Since the process has only one layer of poly, a metal2-metal1-poly capacitor was used. $Dith+$, $Dith-$ are inputs for dithering, and $Vf+$, $Vf-$ are outputs from the comparators. Furthermore the output digital pad buffers are switched only at the downsampled frequency to minimize switching transients.

Next the op amp and comparator design will be described. The op-amp gain requirement in a sigma-delta loop is determined by its effect on the quantization noise suppression at low frequency. The effect of the finite gain of the op amp is to modify the frequency response of the integrator from having a pole at ∞ to having a pole at some finite frequency. Moreover, the dc gain of an integrator is no longer infinite, but equals A , the op-amp gain. The effect of changing the frequency response of the




Fig. 10. Chip photograph.

integrator is to modify the noise transfer function. This modification results in more quantization noise in the passband. From simulation it is determined that when the op-amp gain is less than the oversampling ratio, the increase in quantization noise becomes significant [8]. For the prototype the oversampling ratio is 512. Based on the above result it is determined that a gain of about 1000 is needed.

A fully differential folded cascode configuration with dynamic common-mode feedback was chosen. The common-mode feedback [9] is connected directly to the output node to improve the speed at the expense of having a lower gain. From simulations the amplifier has a typical gain of 1000 and settles to 0.01 percent in 75 ns with a 1.2-V differential step into a 1.2-pF load. To reduce the harmonic distortion due to signal-dependent charge injection the sampling switches are turned off sequentially, with the switches connecting to ground being turned off first.

Since the comparator offset does not affect the performance of the converter, a simple regenerative latch is used to achieve fast switching. In simulation the comparator latches correctly in 20 ns with an overdrive of 10 mV. A chip photograph of the core of a prototype is shown in Fig. 10.

IV. EXPERIMENTAL RESULTS

The measured signal-to-noise ratio versus input amplitude curve is shown in Fig. 11. The signal-to-noise ratio of the converter was evaluated by putting a sine wave of 1.024 kHz into the converter, uploading the digital output to a computer, and then running a 1024-point FFT to compute the ratio of the signal power to the noise power.




Fig. 11. Measured signal-to-noise ratio versus input amplitude.

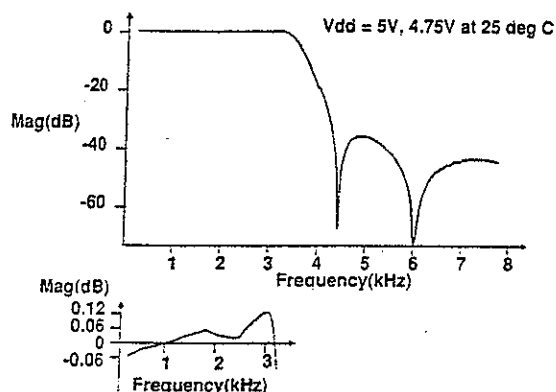


Fig. 12. Measured frequency response.




Fig. 13. Measured gain tracking.

The measured result meets the D3 specifications. In Fig. 12 the frequency response of the overall coder is shown for power supplies of 5 and 4.75 V, which show no appreciable difference. The response has more than 33-dB suppression at 4.6 kHz and beyond. The passband ripple is less than 0.25 dB. In Fig. 13 the gain tracking characteristic of the coder is shown. The measured result meets the D3 specifications. The observed performance of the prototype is summarized in Table I. The crosstalk was measured by putting a full-scale sine wave at 1 kHz at the input of channel one, a zero input to channel two, and then measuring the 1-kHz component at the output of channel two. The crosstalk was measured to be -83 dB. Power supply rejection ratio at 1 kHz was measured by applying a 20-mV peak-to-peak sine wave at 1 kHz at the V_{dd} of the chip, zero input to the chip, and measuring the 1-kHz component at the output. Measurements were repeated for different frequencies from 500 Hz to 20 kHz. The PSRR at 10 kHz is better than that at 1 kHz because the power noise is suppressed by the digital filter. The idle channel noise was measured by putting zero input at the channel input, summing up the noise at the output, and taking the

TABLE I
DATA SUMMARY
Typical performance: 5-V power supply and 25°C

Area	8250 sq. mils/ch
Dynamic range	79dB
Idle channel noise	12dB _{BrnC0}
Crosstalk	-83dB
PSRR	43dB (1 kHz) 60dB (10 kHz)
Harmonic distortion	-76dB (2nd) -82dB (3rd)
Power	50mW/ch

ratio between the measured noise and the input signal level at the overload point. The measured value was 12 dB_{BrnC0} C-message weighted. In addition the variation in idle channel noise as a function of dc offset has been evaluated. A dc offset in increments of 5 mV and a range from -100 to $+100$ mV was applied to the coder and the measured idle channel noise was found to vary less than 1 dB. The anti-alias requirement was also tested at higher frequencies of 13 and 29 kHz to measure how effective the FIR1 and FIR2 filters are in removing frequency components around 16 and 32 kHz, respectively. At 13 kHz the signal is attenuated by 42 dB and at 29 kHz the signal is attenuated by 39 dB, both attenuations larger than the 33-dB requirement.

V. SUMMARY

This paper reports on a prototype four-channel over-sampled CMOS PCM voice-band coder that meets the D3 specifications in an area of 8250 mils² per channel. In summary, the prototype demonstrates that when compared to the conventional single-channel sigma-delta approach, a multichannel first-order sigma-delta coder in voice-band application can substantially reduce the chip area per channel in a scaled technology.

ACKNOWLEDGMENT

The authors wish to thank Y.-M. Lin for his valuable discussions, particularly on the implementation of the FIR filter. They would also like to thank MOSIS for fabricating the chip.

REFERENCES

- [1] P. Defraeye *et al.*, "A 3- μ m CMOS digital codec with programmable echo cancellation," *IEEE J. Solid-State Circuits*, vol. SC-20, pp. 679-688, June 1985.
- [2] J. C. Candy, "Decimation for sigma delta modulation," *IEEE Trans. Commun.*, vol. COM-34, pp. 72-76, Jan. 1986.
- [3] J. E. Iwersen, "Calculated quantizing noise of single-integration delta-modulation coders," *Bell Syst. Tech. J.*, vol. 48, pp. 2359-2389, Sept. 1969.

- [4] E. B. Hogenauer. "An economical class of digital filters for decimation and interpolation." *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-29, pp. 155-162, Apr. 1981.
- [5] A. N. Netravali. "Optimum digital filters for interpolative A/D converters." *Bell Syst. Tech. J.*, vol. 56, pp. 1629-1641, Nov. 1977.
- [6] J. Raebaeys, S. P. Pope, and R. W. Brodersen. "An integrated automated layout generation system for DSP circuits." *IEEE Trans. Computer-Aided Des.*, vol. CAD-4, pp. 285-296, July 1985.
- [7] R. Jain *et al.* "Custom design of a VLSI PCM-FDM transmultiplier." *IEEE J. Solid-State Circuits*, vol. SC-21, pp. 73-86, Feb. 1986.
- [8] M. W. Hauser and R. W. Brodersen. "Circuit and technology considerations for MOS delta-sigma A/D converters." in *ISCAS Conf. Proc.*, May 1986, pp. 1310-1315.
- [9] R. Castello and P. R. Gray. "A high-performance micropower switched-capacitor filter." *IEEE J. Solid-State Circuits*, vol. SC-20, pp. 1122-1132, Dec. 1985.



Paul R. Gray (S'65-M'69-SM'76-F'81), for photograph and biography please see this issue, p. 1323.



Bosco H. Leung (S'84-M'87) received the B.Sc. degree in electrical engineering from the Rensselaer Polytechnic Institute, Troy, NY, in 1979, and the M.Sc. degree in electrical engineering from the California Institute of Technology, Pasadena, in 1981. In 1983 he was awarded the Bell Northern Research Scholarship for his Ph.D. studies at the University of California, Berkeley, where he received the Ph.D. degree in electrical engineering and computer science in 1987.

From 1980 to 1982 he was working in the Digital Switching Division of Northern Telecom, Calgary, Canada, and from 1982 to 1983 he was with the Business Product Division of Northern Telecom, Calgary, Canada, where he was involved in the design of the speakerphone. Since 1988 he has been an Assistant Professor at the University of Waterloo, Waterloo, Ont., Canada. His main research interest is in integrated circuits for analog-to-digital conversion and digital signal processing.



Robert Neff received the B.S. degree in engineering from Swarthmore College, Swarthmore, PA, in 1985 and the M.S. degree in electrical engineering from the University of California, Berkeley, CA, in 1987.

He joined the IBM Corporation in 1987 at the General Products Division in San Jose, CA, in the VLSI design area. Currently he is working on data acquisition circuits for the servo function in advanced storage products.

Robert W. Brodersen (M'76-SM'81-F'82) received B.S. degrees in electrical engineering and mathematics from the California State Polytechnic University, Pomona, in 1966. He received the M.S. and Engineering degrees in 1968 and the Ph.D. degree in 1972 from the Massachusetts Institute of Technology, Cambridge.

From 1972 to 1976 he was with the Central Research Laboratory at Texas Instruments Incorporated, Dallas, where he was engaged in the study of operation and applications of charge-coupled devices. In 1976 he joined the Faculty of the University of California at Berkeley, where he is now a Professor. In addition to teaching, he is investigating the use of MOS technology for signal processing applications.

Dr. Brodersen has received the Morris Libermann Award of the IEEE (1982), the W. G. Baker Award for the outstanding paper in *IEEE JOURNALS AND TRANSACTIONS* (1979), and the award for the best paper in the *IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN* (1986). He has also won best paper awards at *EASCON* (1973), the *International Solid State Circuit Conference* (1975), and the *European Solid-State Circuit Conference* (1978). He has been named the *Outstanding Engineering Alumnus* of the California State Polytechnic University (1978).