Algorithmic $\Delta\Sigma$-Modulated FIR Filter

Ashkan Olyaei
Department of Electrical and Computer Engineering
University of Toronto
Toronto, ON M5S 3G4, Canada
Email: olyaei@eecg.toronto.edu

Roman Genov
Department of Electrical and Computer Engineering
University of Toronto
Toronto, ON M5S 3G4, Canada
Email: roman@eecg.toronto.edu

***Abstract*— We present an algorithmic $\Delta\Sigma$-modulated FIR filter which computes digital convolution of a continuous-time analog input signal with a programmable digital impulse response. Selective sampling of the input signal controlled by unary-encoded FIR coefficients yields bit-serial analog-digital multiplication. A $\Delta\Sigma$-modulated analog-to-digital converter samples a time-varying input at multiple instances in time generating a quantized version of the average of all weighted samples. Computational throughput of an arbitrary FIR filter is maximized by algorithmic resampling of the modulation residue to obtain higher resolution bits. This yields a bit resolution linear in the number of conversion cycles. A 1.9mm $\times$ 1.3mm 128-channel FIR filter integrated prototype was fabricated in a 0.35 μm CMOS technology. It yields a computational throughput of up to 3.8 GMACS, with computational quantization time, power dissipation, and integration area comparable to those in a conventional oversampling analog-to-digital converter.**

I. Introduction

Various techniques for realizing finite-length impulse response (FIR) filters in sensory systems have been developed. Dedicated digital signal processors (DSP) use high-throughput architectures to compute programmable FIR filter convolution sums, but require an analog-to-digital converter (ADC) to convert the analog sensory input into the digital domain prior to digital signal processing. In addition, the low data rate serial input of such processors limits their sustained throughput [1], particularly in multi-sensor applications. FIR filters can also be implemented in the analog domain. Current mode implementations use zero-latency addition but deploy multiple matched current mirrors limiting the number of filter coefficients [2]. Capacitor bank implementations use charge sharing to compute average sum and difference [3] but may have limited scalability. Analog domain implementations of FIR filters also require an ADC to generate digital outputs after the signal processing operation.

As compared to the existing FIR design approaches, the mixed-signal FIR filter presented here combines both signal processing and quantization in a single compact low-power architecture as shown in Figure 1. It directly computes digital convolution of a continuous-time analog input signal with a programmable digital impulse response. Our approach combines selective sampling, accumulation, and concurrent quantization in a single $\Delta\Sigma$-modulated analog-to-digital conversion cycle. Signed multiplication and averaging oversampling quantization are performed with computational quantization time, power dissipation, and integration area comparable to those in a conventional oversampling analog-to-digital converter. The rest of this paper is organized as follows. Section II gives an overview of the $\Delta\Sigma$-modulated multiplying ADC architecture. Section III extends the $\Delta\Sigma$-modulated multiplying ADC functionality to weighted averaging. Section IV presents the $\Delta\Sigma$-modulated FIR filter implementation. Section V discusses an algorithmic quantization scheme, comprised of multiple oversampling data conversion cycles, which optimizes the throughput of the architecture for an arbitrary impulse response. Section VI presents experimentally measured results obtained from a 0.35 micron CMOS prototype of a 128-channel bank of algorithmic $\Delta\Sigma$-modulated FIR filters.

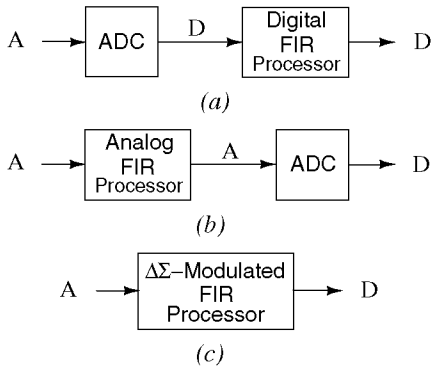

Fig. 1. FIR filter architectures with *(a)* a digital processor, *(b)* an analog processor, and *(c)* a $\Delta\Sigma$-modulated processor.

II. $\Delta\Sigma$-Modulated Multiplying ADC

A first-order incremental oversampling ADC is depicted in Figure 2*(a)*. It converts a sequence of analog samples into a digital word representing a quantized version of the average of all samples. It is comprised of a sample-and-hold (S/H) circuit, an integrator, a comparator and a decimating counter. As shown in Figure 2*(b)*, this architecture can be extended to perform both quantization and signed multiplication of the input with a digital word

$$C = sign(C)\ |C|, \tag{1}$$

with

$$|C| = \sum_{i=0}^{N-1} c[i], \tag{2}$$

where $c[i]$ are unsigned unary coefficients of C.

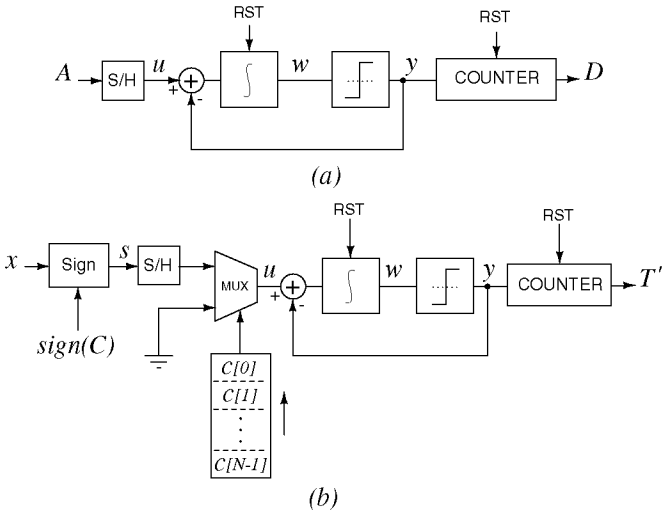

Fig. 2. *(a)* First-order $\Delta\Sigma$ incremental A/D converter. *(b)* $\Delta\Sigma$-modulated multiplying ADC.

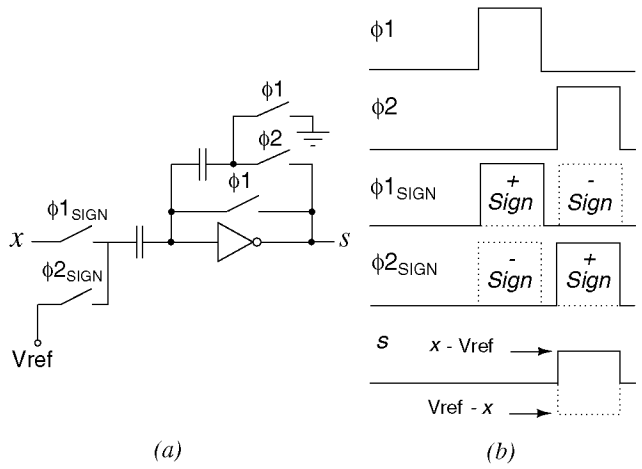

Fig. 3. *(a)* The sign unit circuit in the $\Delta\Sigma$-modulated FIR filter. *(b)* Sign circuit timing diagram; for a continuous range of signed outputs, Vref = min{x}. The circuit also performs the sample-and-hold operation. (S/H is explicitly shown in Figure 2*(b)* for clarity.)

Selective sampling of the sign-transformed input s, controlled by the bit-serial sequence $c[i]$, yields an analog sequence $u[i] = sc[i]$. The first-order modulator converts the sequence $sc[i]$ into a bitstream $y[i]$ in N cycles, using a 'resetable' (RST) analog integrator:

$$w[0] = 0, \tag{3}$$

$$w[i+1] = w[i] + \alpha\ (sc[i] - y[i]), \quad i = 0, \ldots N-1, \tag{4}$$

$$w[N+1] = w[N] - \alpha\ y[N], \tag{5}$$

and a single-bit quantizer:

$$y[0] = -1, \tag{6}$$

$$y[i] = \text{sign}(w[i]), \quad i = 1, \ldots N, \tag{7}$$

where α is the intrinsic gain of the integrator.

A binary counter accumulates the bits $y[i]$ to produce a decimated output. The rectangular decimation window, and initial reset of the integrator, avoid tones in the quantization noise spectrum at DC input that are characteristic of a conventional first-order $\Delta\Sigma$ modulator with lowpass decimation filter [4]. The quantization error (conversion residue) is directly given by the final integrator value $\frac{1}{\alpha}w[N+1]$, as verified by summing (4) and (5) over i:

$$\sum_{i=0}^{N} y[i] = \sum_{i=0}^{N-1} sc[i] - \frac{1}{\alpha}w[N+1], \tag{8}$$

where $\sum_{i=0}^{N} y[i] = T'$ is the digital output. This operation yields multiplication of the sign-transformed analog input s with the digital unary-coded coefficients $c[i]$, while a digital output resolution of $\log_2(N)$ bits is warranted:

$$T' = |C|\ s + q', \tag{9}$$

where

$$|q'| = |\frac{1}{\alpha}w[N+1]| \tag{10}$$

is the multiplication quantization noise. Higher resolution at lower oversampling ratio N can be obtained using higher-order incremental conversion [5].

The double sampling sign unit circuit shown in Figure 3*(a)* applies the sign of the coefficient to the input by selecting a switched-capacitor sampling sequence order as illustrated in the timing diagram in Figure 3*(b)*:

$$s = sign(C)\ x, \tag{11}$$

which in combination with (1) and (9) yields signed multiplication:

$$T' = C\ x + q'. \tag{12}$$

III. $\Delta\Sigma$-Modulated Weighted Averaging ADC

The multiplying ADC architecture in Section II can be extended to perform weighted averaging. Accumulation is achieved by sampling a time-varying input, $s_1, s_2 \ldots s_K$, at multiple instances in time without resetting the integrator or the binary counter.

The architecture of the oversampling weighted averaging ADC is depicted in Figure 4. The bit stream $y[ki]$ is now generated for K inputs each sampled N times:

$$w[0] = 0, \tag{13}$$

$$w[k(i+1)] = w[ki] + \alpha\ (s[k]c[i,k] - y[ki]), \quad i = 0, \ldots N-1, \quad k = 1, \ldots K, \tag{14}$$

$$w[K(N+1)] = w[KN] - \alpha\ y[KN], \tag{15}$$

using a single-bit quantizer:

$$y[0] = -1, \tag{16}$$

$$y[ki] = \text{sign}(w[ki]), \quad i = 1, \ldots N, \quad k = 1, \ldots K. \tag{17}$$

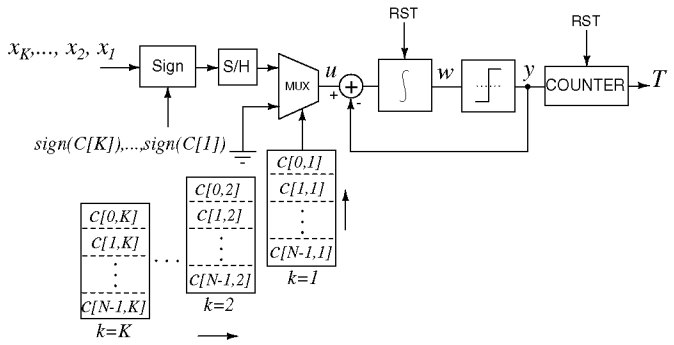

Fig. 4. $\Delta\Sigma$-modulated weighted averaging ADC.

The quantization error, $\frac{1}{\alpha}w[K(N+1)]$, is obtained similarly:

$$\sum_{k=1}^{K}\sum_{i=0}^{N} y[ki] = \sum_{k=1}^{K}\sum_{i=0}^{N-1} s[k]c[i,k] - \frac{1}{\alpha}w[K(N+1)], \quad (18)$$

where $\sum_{k=1}^{K}\sum_{i=0}^{N} y[ki] = T$ and the notation is consistent with that of (3)–(8).

This realizes an inner product of input $s[k]$ with the digital unary-coded coefficients $c[i,k]$ with an output resolution of $\log_2(KN)$ bits:

$$T = \sum_{k=1}^{K} |C[k]|\ s[k] + q, \quad (19)$$

which in combination with (1) and (11) yields:

$$T = \sum_{k=1}^{K} C[k]\ x[k] + \ q, \quad (20)$$

where

$$|q| = |\frac{1}{\alpha}w[K(N+1)]| \quad (21)$$

is the weighted averaging quantization error.

IV. $\Delta\Sigma$-Modulated FIR ADC

The weighted averaging architecture in Section III can be further extended to perform FIR filtering. A moving window selects the corresponding input samples, which in combination with sign, multiplication and accumulation operations shown in (20) results in computation of the FIR filter convolution sum:

$$T[n] = \sum_{k=1}^{K} C[k]\ x[K+n-k] + q,$$

where q is the weighted averaging quantization error and K is the number of FIR filter taps. A ring shift register of length K and width N applies the coefficients to the corresponding input samples in a cyclic fashion, as shown in Figure 5. An analog memory stores K consecutive samples of the input sequence. The memory requirement is not an additional overhead in spatial sensor arrays, such as CMOS imagers with in-pixel memory [6].

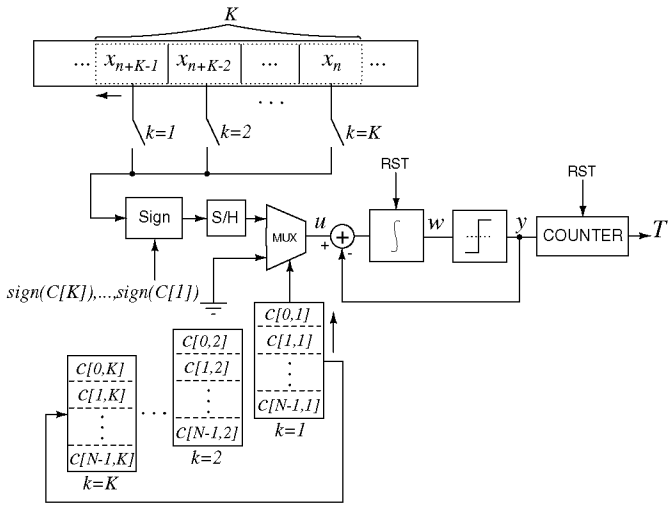

Fig. 5. $\Delta\Sigma$-modulated FIR filter.

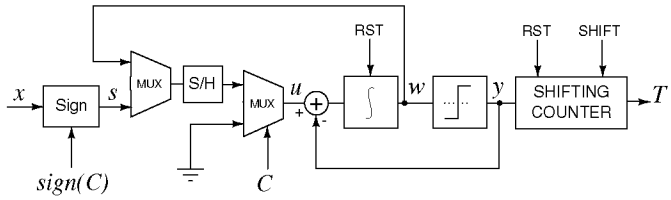

Fig. 6. Algorithmic $\Delta\Sigma$-modulated FIR ADC, with residue resampling across the integrator, and shifting counter for the decimator.

V. Algorithmic $\Delta\Sigma$-Modulated FIR ADC

An algorithmic quantization scheme [7] is employed to further increase the sustained throughput of the FIR processor. When $log_2(KN)$ is less than the desired digital output resolution, the conversion residue $\frac{1}{\alpha}w[K(N+1)]$ in (18) is quantized further to obtain higher resolution bits. Ratio-insensitive matching is achieved by resampling the residue through the same signal path as used for accumulation (14), and employing the same modulator to convert the residue [7] as shown in Figure 6.

At the end of the oversampling conversion, the residue $w[K(N+1)]$ is fed back to the input for subsequent incremental conversion over N' additional cycles, $u'[i] = \beta w[K(N+1)]$, $i = 0, \ldots N'-1$, where β represents the residue resampling gain. Thus

$$\sum_{i=0}^{N'} y'[i] = N'\beta\ w[K(N+1)] - \frac{1}{\alpha}w'[N'+1], \quad (22)$$

which under the matching condition $\alpha\beta \equiv 1$ combines with (18) to cancel the first residue $w[K(N+1)]$:

$$N'\sum_{k=1}^{K}\sum_{i=0}^{N} y[ki] + \sum_{i=0}^{N'} y'[i] = N'\sum_{k=1}^{K}\sum_{i=0}^{N-1} s[k]c[i,k] - \frac{1}{\alpha}w'[N'+1]. \quad (23)$$

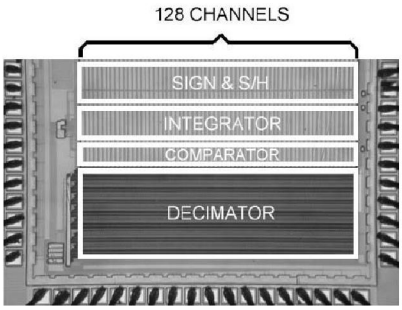

Fig. 7. Micrograph of the 128-channel algorithmic $\Delta\Sigma$-modulated FIR filter. To demonstrate the advantages of the approach to large-scale sensory signal processing, the FIR filters are integrated with a 128 $\times$ 128-pixel array, not shown, as part of a mixed-signal focal-plane processor.

The left-hand side of (23) is readily available in digital form, and the right-hand side terms conform to (18) with effective KNN'-level resolution. Therefore, the residue conversion (22) enhances the resolution of (18) by an extra $\log_2(N')$ bits [7].

Optimization of the number of oversampling and algorithmic conversion cycles based on the FIR coefficients can enhance the computational throughput of the architecture. When the sum of absolute values of coefficients of an FIR filter (scaled to all integers), N_{FIR}, is less than N, the oversampling computational cycle is stopped on the N_{FIR}-th sample. Higher resolution bits are obtained by subsequent algorithmic residue resampling and extended counting on the residue.

VI. Experimental Results

The architecture has been experimentally validated on a 0.35 micron CMOS VLSI prototype containing a bank of 128 algorithmic $\Delta\Sigma$-modulated analog-to-digital converters (ADCs) simultaneously computing FIR of 128 analog input sequences. The micrograph of this multi-channel FIR filter is shown in Figure 7. We perform spectral analysis of the input waveform using a 15-tap integrated Morlet wavelet filter shown in Figure 8*(a)*, with sum of absolute values of all FIR coefficients equal to 32. The algorithmic $\Delta\Sigma$-modulated ADC computes the response of this filter to a sinusoidal input signal, with frequency increasing linearly in time, as shown in Figure 8*(b)*. The filter response computed by a floating-point processor is illustrated in Figure 8*(c)*. Using only oversampling conversion, the filter can be realized in 256 cycles (with 4-bit FIR filter coefficients) yielding 8-bit digital outputs. Taking advantage of the algorithmic residue resampling scheme, the same filter response is computed in 32 oversampling cycles plus another 8 oversampling cycles to obtain the three less significant bits for a total of 40 conversion cycles yielding 8-bit outputs. Figure 8*(d)* shows this FIR filter response, experimentally computed and measured from a single channel of the VLSI $\Delta\Sigma$-modulated FIR filter array. Computing this FIR on the 128-channel filter yields a computational throughput of 3.8 GMACS based on a conservative quantizer sampling rate of 62.5 ksps.

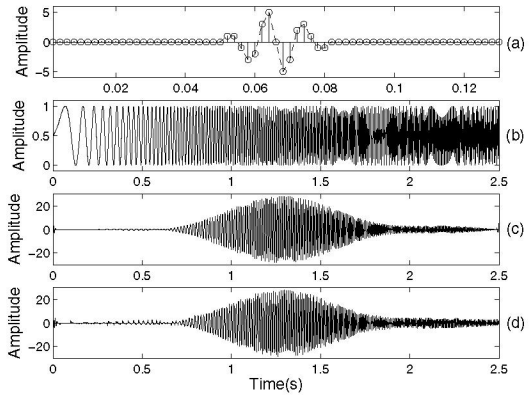

Fig. 8. *(a)* 15-tap integrated Morlet wavelet. *(b)* Sinusoidal input signal with frequency increasing linearly in time. *(c)* FIR Filter response computed by a floating-point processor. *(d)* FIR filter response, experimentally computed and measured from a single channel of the VLSI $\Delta\Sigma$-modulated FIR filter array.

VII. Conclusion

We present a mixed-signal VLSI architecture of an algorithmic $\Delta\Sigma$-modulated FIR filter which computes digital convolution sum of a programmable impulse response with a continuous-time analog input signal. Signed multiplication, accumulation, and oversampling quantization are combined in a single $\Delta\Sigma$-modulated analog-to-digital conversion cycle. Selective sampling yielding multiplication, and multiple input sampling yielding accumulation are performed with computational quantization time, power dissipation, and integration area comparable to those in a conventional oversampling analog-to-digital converter. Algorithmic post-conversion is used to achieve a higher bit resolution by iteratively resampling the modulation residue through the same signal path. Optimization of the number of oversampling and algorithmic conversion cycles based on the FIR coefficients further enhances the computational throughput of the architecture. Experimental results obtained from a 0.35 micron CMOS prototype with a bank of 128 $\Delta\Sigma$-modulated FIR filters validate the utility of the design for large-scale parallel microsystems with electrical, chemical or photonic on-sensory-plane signal processing.

References

[1] J. Park, W. Jeong, H. Mahmoodi-Meimand, Y. Wang, H. Choo, and K. Roy, "Computation Sharing Programmable FIR Filter for Low-Power and High-Performance Applications," *IEEE J. Solid-State Circuits,* vol **39** (2), pp 348-357, 2004.

[2] F.A. Farag, C. Galup-Montoro, and M.C. Schneider, " Digitally Programmable Switched-Current FIR Filter for Low-Voltage Applications," *IEEE J. Solid-State Circuits,* vol **35** (4), pp 637-641, 2000.

[3] L. Lentola, G.M. Corelazzo, E. Malavasi, and A. Baschirotto, " Design of SC Filters for Video Applications," *IEEE T. Circuits and Systems for Video Technology,* vol **10** (1), pp 14-22, 2000.

[4] J. Robert, G.C. Temes, V. Valencic, R. Dessoulavy, and P. Deval, "A 16-Bit Low-Voltage CMOS A/D Converter," *IEEE J. Solid-State Circuits,* vol. **SC-22**, pp 157-163, 1987.

[5] O.J.A.P. Nys and E. Dijkstra, "On Configurable Oversampled A/D Converters," *IEEE J. Solid-State Circuits,* vol. **28** (7), pp 736-742, 1993.

[6] A. Olyaei, and R. Genov, "Oversampling Focal-Plane Spatio-Temporal Processor," *IEEE Trans. Circuits and Systems (TCAS'06)*, Subm.

[7] G. Mulliken, F. Adil, G. Cauwenberghs and R. Genov, " Delta-sigma algorithmic analog-to-digital conversion," *IEEE Int. Symp. Circuits and Systems (ISCAS'02)*, vol **4**, pp IV.687-IV.690, 2002.