

VLSI Implementation of a Reconfigurable Mixed-Signal Finite Impulse Response Filter

Erhan Özalevli, Walter Huang, Paul Hasler, and David V. Anderson
 School of Electrical and Computer Engineering
 Georgia Institute of Technology, Atlanta, Georgia 30332-0250

Abstract— We present an implementation of a reconfigurable 16-tap finite impulse response filter for post-processing applications. This filter exploits the distributed arithmetic technique for signal processing and floating-gate voltage references for setting tunable analog coefficients. The filter is fabricated in 0.5 μm CMOS process, and its order can be increased at the cost of 0.011 mm^2 of die area and 0.02 mW of power per tap. Measurement results for low-pass and band-pass filters at 50 kHz sampling frequency are presented.

I. INTRODUCTION

Low-power circuits are crucial for the implementation of signal processing applications in portable electronics that require extended battery life. Decreasing the power consumption per functional block in these devices is a design challenge that has to be dealt with to accommodate more functionality. Among the most versatile and widely used functional blocks are finite impulse response (FIR) filters. Therefore, the design of a low-power FIR filter can readily ease the shrinking power budget of portable electronics.

Although reconfigurable FIR filters are usually implemented in the digital domain to take advantage of full programmability and wide dynamic range, their analog counterparts have significant potential for decreasing the power consumption in certain signal processing applications. The difficulty of obtaining reconfigurability in the analog domain has been overcome by employing switched-capacitor or switched-current design techniques [1], [2]. Switched-capacitor techniques are very suitable for FIR filter implementations since they enable a precise control over their filter coefficients [1], [3]. Although the power and speed trade-off and the error accumulation in these filters are alleviated by using a transposed FIR filter structure [4] or a rotating switch matrix technique [5], these implementations still suffer from high area overhead and high power consumption. Similar to switched-capacitor techniques, switched-current techniques allow for the integration of the digital coefficients with analog filter implementations through the use of the current division technique [2], circular buffer architecture [6], or multiplying digital-to-analog converters (MDAC) [7]. Recently, a switched-current FIR filter based on distributed arithmetic (DA) has also been suggested to decrease the hardware complexity [8]. These techniques, although designed for pre-processing applications, can be also utilized for post-processing applications such as MP3 decoding by using them after a digital-to-analog converter (DAC). However, the use of an additional high-resolution

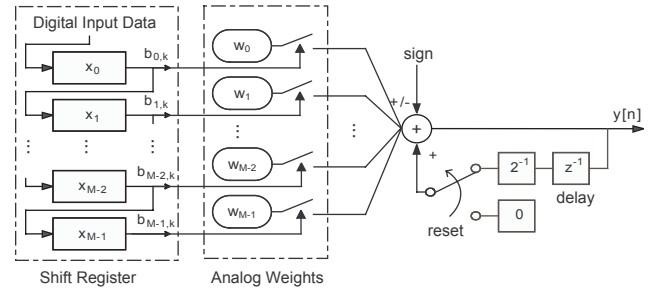


Fig. 1. Proposed FIR filter architecture using distributed arithmetic (DA). Digital input data is processed in the analog domain.

and/or high-speed DAC causes an increase in area and power consumption. In addition, these design techniques still pose problems when implementing compact high-order filters with low-power consumption. This paper proposes a reconfigurable FIR filter architecture that allows for the implementation of a low-power and compact high-order filter by employing the DA technique for signal processing and by utilizing the analog storage capabilities of floating-gate (FG) transistors to obtain programmable analog coefficients for reconfigurability.

II. DA COMPUTATION

The DA concept was first introduced by Croisier et al. [9]. It is an efficient computational method for computing the inner product of two vectors in a bit-serial fashion [10]. The operation of DA can be derived from the inner product equation, $y[n] = \sum_{m=0}^{M-1} x[n-m]w[m]$. In the case of FIR filtering, x is the input vector and w is the weight vector. Using a K -bit 2's-complement representation for x , the above equation can be written as

$$y[n] = - \sum_{i=0}^{M-1} w_i b_{i0} + \sum_{j=1}^{K-1} 2^{-j} \sum_{i=0}^{M-1} w_i b_{ij} \quad (1)$$

where b_{i0} is the sign bit, b_{ij} is the j^{th} bit of the i^{th} element in the vector x , and $b_{i(K-1)}$ is the least significant bit.

In digital implementations, the summation, $\sum_{i=0}^{m-1} w_i b_{ij}$, is pre-computed and stored in a memory for a multiplier-less operation and reduced hardware complexity. Relative to digital implementations, addition in the analog domain is much more power and area efficient. Therefore, the high memory usage of digital implementations can be eliminated by processing the digital input data in the analog domain as illustrated in

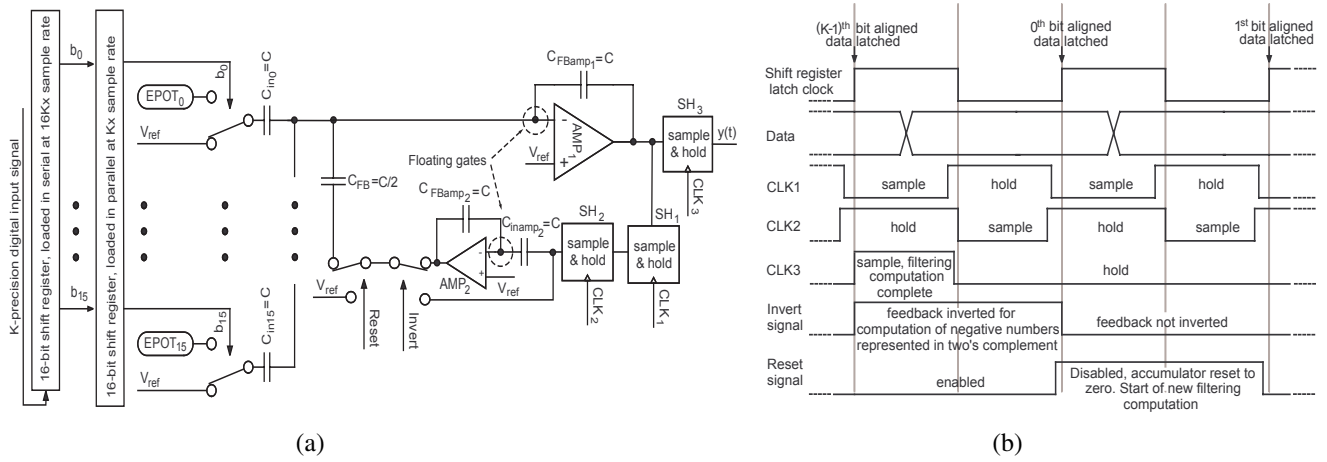


Fig. 2. (a) Implementation of the 16-tap hybrid FIR filter. b_i is the input bit for j^{th} cycle of operation and $y(t)$ is the output. Epots store the analog weights. Sample-and-holds, SHs, are used to obtain the delay and hold the computed output voltage. (b) Timing diagram of the digital input data and control bits.

Figure 1. For an individual analog weight, data is processed in a similar way as it is achieved by serial DACs, where the conversion is performed sequentially.

III. FIR FILTER ARCHITECTURE

The proposed hybrid FIR filter architecture, shown in Figure 2a, consists of four components, 16-bit shift registers, an array of tunable FG voltage references (epot) [11], inverting amplifiers (AMP), and sample-and-hold (SH) circuits.

The timing of the digital data and control bits governs the filtering operation and is illustrated in Figure 2b. Digital inputs are introduced to the system by using the serial shift register. These digital input words represent the digital bits, b_{ij} in (1), which selects the appropriate epot voltages to form the appropriate sum of weights necessary for DA at the j^{th} bit. The clock frequency of this shift register is dependent on the input data precision, K , and the length of the filter, M , and is equal to $M \cdot K$ times the sampling frequency. The analog weights of the FIR filter are stored by the epots. When selected, these weights are added by employing a charge amplifier structure, which is composed of same size capacitors and a two-stage amplifier, AMP_1 . For the first operation cycle, the result of the addition stage represents the summation, $\sum_{i=0}^{m-1} w_i b_{i(K-1)}$, in (1), which is the addition of weights for the $LSBs$ of the digital input data.

The feedback path of the system is used to obtain the delay, to achieve a 2's complement compatible computation, and to implement the divide-by-two operation for DA computation. For that purpose, sample-and-hold circuits, SH_1 and SH_2 , and inverting amplifiers, AMP_1 and AMP_2 , are employed. SH circuits store the amplifier output to feed it back into the system for the next cycle of the operation. The non-overlapping clocks, CLK_1 and CLK_2 , have a frequency of K times the sampling frequency. The stored data is then inverted relative to the reference voltage by using the second inverting amplifier, AMP_2 , to obtain the same sign as the summed epot voltages. AMP_2 is identical to AMP_1 , and has the same

size input/feedback capacitors. After obtaining the delay and the sign correction, the stored analog data is fed back to the addition stage as delayed analog data. During the addition, it is also divided by two by using $C_{FB} = C_{FBamp1}/2 = C/2$, which gives a gain of one-half when it is added to the new sum. This operation is repeated $K - 1$ times. During the last cycle of operation, as shown in Figure 2b, the 2's-complement operation is achieved by disabling the inverting amplifier in the path by enabling the *Invert* signal. Finally, when the computation is finished, it is sampled by SH_3 using CLK_3 , which is enabled once every K cycles. SH_3 holds the computed voltage till the next analog output voltage is ready. The new computation starts first by enabling the *Reset* signal to zero out the effect of the previous computation.

A. Computational Error Sources

The main sources of error in this architecture are the gain and offset errors in the feedback path. Due to the serial nature of the computation, as the digital input precision increases, the offset error diminishes, but the gain error monotonically increases. Therefore, the gain error becomes dominant. The worst case gain error occurs when the sum of the weights is largest. For $\max\left(\left|\sum_{i=0}^{M-1} w_i b_{ij}\right|\right) = \alpha$, this error becomes

$$\varepsilon = \alpha \left(\frac{1 - (0.5 - \Delta)^K}{0.5 + \Delta} - \frac{1 - 0.5^K}{0.5} \right) \quad (2)$$

where Δ is the gain non-ideality and K is the input precision. This error, depicted in Figure 3a, is normalized by α and given as a function of K for a range of Δ values. The plot illustrates the fact that the intersection of the quantization error and ε determines the maximum Δ tolerable by the system for accurate computation. Therefore, as long as the desired ε for a specific Δ value is less than the gain error at the intersection point, the system is limited by the digital input precision. For example, when the gain error is 2^{-9} and $K = 8$, Δ has to be smaller than 2^{-11} to meet the accuracy requirements.

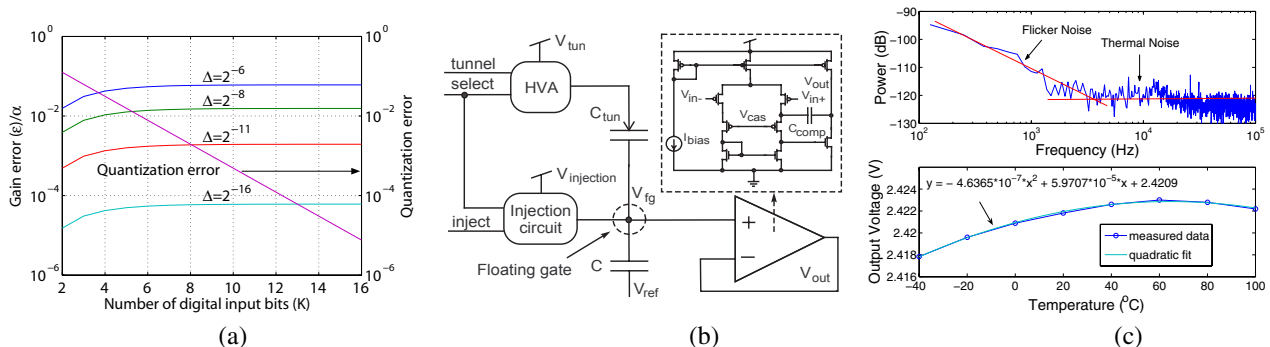


Fig. 3. (a) Computational error of the system due to quantization error and gain non-ideality. (b) Modified epot schematic. Epot is built by using a low-noise amplifier and FG programming circuit. C_{tun} is the tunnelling junction and HVA is a high-voltage-amplifier used for tunnelling. (c) Output noise and temperature sweep of the epot. The epot exhibits second-order temperature dependence when programmed to 2.422V.

B. Circuit description of the computational blocks

To achieve an accurate computation using DA, the circuit components are designed to minimize the gain and offset errors in the signal path. In this architecture, those components are epots, inverting amplifiers, and sample-and-holds.

The epots and inverting amplifiers use FG transistors to exploit their analog storage and capacitive coupling properties. The epot, shown in Figure 3b, is modified from its original version [11] to obtain a low-noise voltage reference that has a low temperature coefficient. It is a dynamically reprogrammable, on-chip voltage reference that uses a low-noise amplifier integrated with FG transistors and programming circuitry to tune the stored analog voltage. An array of epots is used for storing the weights of the filter and individual epots are controlled and read by employing a decoder. A tuning of the stored voltage on the FG node is achieved by utilizing the hot-electron injection and Fowler-Nordheim tunnelling mechanisms. Tunnelling is used for coarse programming of the epot voltage. It decreases the amount of electrons by using a 14V tunnelling voltage. Precise programming is performed by using the hot-electron injection, which increases the amount of electrons on the FG node using a 6.5V injection voltage.

Unlike switched-capacitor amplifiers, the addition and inversion are achieved without resetting the inverting node of the amplifiers. This is because the FGs at the inverting-node of the amplifiers allow for continuous-time operation. Inverting amplifiers are implemented by using a two-stage amplifier structure to obtain high gain and large output swing. Similar to epots, the charge at the FG node of these amplifiers is precisely programmed by monitoring the amplifier output while the system operates in the reset mode. In this mode, the shift registers are cleared and *Reset* signal is enabled. By using this technique, the offset at the amplifier output is reduced to less than 1mV.

Lastly, SH circuits are designed to simultaneously achieve high sampling speed and high sampling precision for meeting the requirements of the bit-serial DA computation. Therefore, these circuits are implemented by utilizing the sample-and-hold technique using Miller hold capacitance [12]. This compact circuit minimizes the signal dependent error, while

maintaining the sampling speed and precision. Also, a gain-boosting technique [13] is incorporated into the SH amplifiers to achieve a high gain and fast settling. In this FIR filter implementation, two SH circuits with non-overlapping clocks are used in the feedback path to obtain the fixed delay for the sampled analog voltage. In addition, the third SH is utilized to sample and hold the computed final output once every K cycles. This SH uses a negative-feedback output stage [14] to buffer the output voltage off-chip. Due to the performance requirements of the system, these SH circuits consume more power than the rest of the system.

IV. MEASUREMENT RESULTS

In this section, we present the experimental results for the proposed FIR filter that was fabricated in a 0.5 μm CMOS process. The precision of the digital input data is set to 8 bits, and the data is loaded into the shift register at 6.4MHz for a 50kHz sampling frequency.

The epots are characterized to observe their noise level and temperature dependence. As shown in Figure 3c, the epot exhibits a thermal noise of -120db and a temperature coefficient of 14.8ppm/ $^{\circ}C$. The noise corner is measured to be around 4kHz. For an array of 16 epots programmed to different voltages, the mean is found as 16ppm/ $^{\circ}C$ with a maximum variation of 20.8ppm/ $^{\circ}C$. This coefficient is mainly determined by the temperature coefficients of the input capacitor and of the amplifier offset. The epots are programmed relative to the reference voltage, which is set to 2.5V. The error of the epot voltages are kept below 1mV to minimize the effect of weight errors on the filter characteristics.

To demonstrate the reconfigurability, the filter is configured as low-pass and band-pass. An 858Hz sinusoidal output of the low-pass filter at a 50kHz sampling rate is illustrated in Figure 4a. The spurious-free-dynamic-range (SFDR) of the output signal is measured to be 43dB. Although the input precision was set to 8 bits, the gain error in the system as well as noise in the experimental set-up limits the maximum achievable SFDR. The magnitude and phase responses of the filters, shown in Figures 4b and 4c, are measured by sweeping the frequency of the input sine wave from DC to

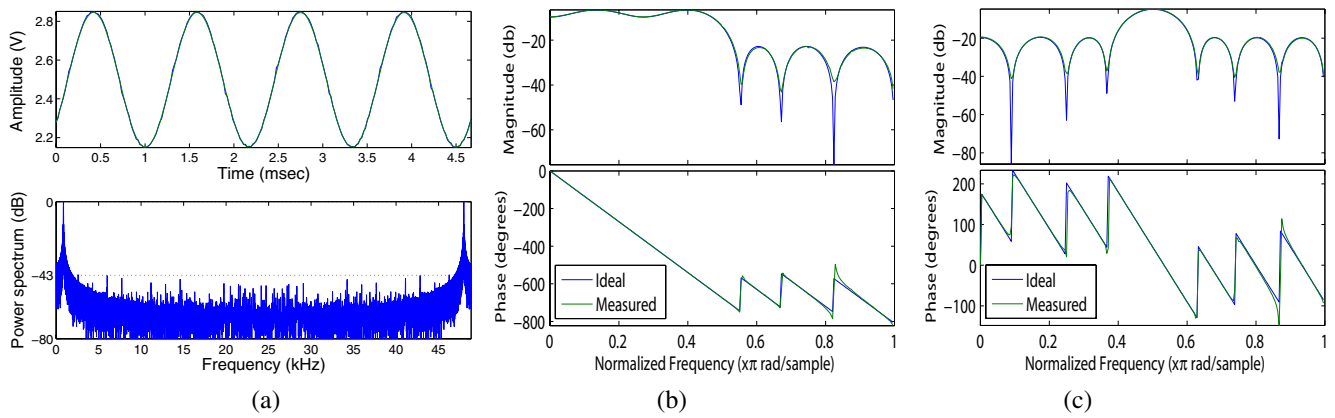


Fig. 4. (a) Low-pass filter transient response and its power spectrum. The sampling frequency is set to 50kHz . The sinusoidal transient signal has a frequency of 858Hz . (b) Low-pass filter magnitude and phase responses. (c) Band-pass filter magnitude and phase responses.

25kHz . For this experiment, 256 data points are collected to accurately measure the frequency response of these filters. The magnitude and phase responses follow the ideal responses closely. However, as the output signal amplitude becomes very low, the experimental set-up limits the resolvable magnitude and phase. As expected from symmetrical FIR filter, the measured phase responses of low-pass and band-pass filters are linear. The static power consumption of the fabricated chip is 16mW . Most of the power is consumed by the SH and inverting amplifier circuits. The die photo of the designed chip is illustrated in Figure 5. The system occupies around half of the $1.5 \times 1.5\text{mm}^2$ die area. Each additional filter tap contributes 0.011mm^2 to the total area and 0.02mW to the total power consumption, which readily allows for the implementation of high-order filters. Lastly, the performance of the filter is summarized in Table I.

V. CONCLUSION

We have presented an implementation of a reconfigurable FIR filter exploiting the distributed arithmetic to minimize hardware complexity and floating-gate transistors to obtain programmable analog coefficients. The proposed architecture is well suited for implementing high-order FIR filters due to its low-area and power requirements for each additional FIR filter tap. Also, the programmable analog coefficients of this

TABLE I

PERFORMANCE AND DESIGN PARAMETERS OF THE FIR FILTER.

Process	$0.5\mu\text{m}$ CMOS
Power supply	5V
Programming Mechanisms	Hot-Electron Injection and Electron Tunneling
Sampling frequency	50kHz
Input data precision	8
Number of filter taps	16
SFDR	43dB
Increase in the power per tap	0.02mW
Increase in the area per tap	0.011mm^2
Total static power consumption	16mW
Unit capacitor size	300fF
Used chip area	$\sim 1.125\text{mm}^2$




Fig. 5. Die photo of the FIR filter chip.

filter will enable the implementation of adaptive systems that can be used in applications such as noise cancellation.

REFERENCES

- [1] G. Fischer, "Analog FIR filters by switched-capacitor techniques," vol. 37, pp. 808–814, 1990.
- [2] K. Bult and G.J.G.M. Geelen, "An inherently linear and compact MOST-only current division technique," vol. 27, pp. 1730–1735, 1992.
- [3] Z. Ciota, A. Napieralski, and J.L. Noullet, "Analogue realisation of integrated FIR filters," vol. 143, pp. 274–281, 1996.
- [4] B.C. Rothenberg, S.H. Lewis, and P.J. Hurst, "A 20-Msample/s switched-capacitor finite-impulse-response filter using a transposed structure," vol. 30, pp. 1350–1356, 1995.
- [5] Y.S. Lee and K.W. Martin, "A switched-capacitor realization of multiple FIR filters on a single chip," pp. 536–542, 1988.
- [6] S. Lyle, G. Worstell, and R. Spencer, "An analog discrete-time transversal filter in $2.0\mu\text{m}$ CMOS," vol. 2, pp. 970–974, 1992.
- [7] A.M. Chiang, "Low-power adaptive filter," pp. 90–91, 1994.
- [8] P. Sirisuk, A. Worapishet, S. Chanyavilas, and K. Dejhan, "Implementation of switched-current FIR filter using distributed arithmetic technique: exploitation of digital concept in analogue domain," vol. 1, pp. 143–148, 2004.
- [9] A. Croisier, D.J. Esteban, M.E. Levilion, and V. Rizo, "Digital filter for PCM encoded signals," 1973.
- [10] S. A. White, "Applications of distributed arithmetic to digital signal processing: A tutorial review," vol. 6, pp. 4–19, 1989.
- [11] R.R. Harrison, J.A. Bragg, P. Hasler, B.A. Minch, and S. Deweerth, "A CMOS Programmable Analog Memory Cell Array using Floating-Gate Circuits," *IEEE Trans. on Circuit and Systems*, 2001.
- [12] P.J. Lim and B.A. Wooley, "A high-speed sample-and-hold technique using a Miller hold capacitance," vol. 26, pp. 643–651, 1991.
- [13] K. Bult and G.J.G.M. Geelen, "A fast-settling CMOS opamp for SC circuits with 90-dB DC-gain," vol. 25, pp. 1379–1384, 1990.
- [14] K.E. Brehmer and J.B. Wieser, "Large swing CMOS power amplifier," vol. SC-18, pp. 624–629, 1983.