Efficient Implementation of Ternary SDM Filters using State-of-the-Art FPGA

TAYABUDDIN MEMON*, PAUL BECKETT**, AND AMIN Z SADIK***

RECEIVED ON 27.01.2010 ACCEPTED ON 10.03.2010

ABSTRACT

We present the analysis of a ternary FIR filter at varying OSR (Over Sampling Ratios). The sigma delta modulated ternary filter impulse responses obtained using Matlab at varying OSRs show that each doubling of OSR results in an increase of 10dB in the stopband attenuation. BT-FIR (Balanced Ternary FIR Filters) at varying OSRs have been implemented in VHDL using an efficient adder tree organization to gather the partial products. Filters in both pipelined and non-pipelined modes were synthesized on a small number of representative commercial FPGA (Field Programmable Gate Arrays) devices. Both the filter taps and binary inputs use 2's complement format. The synthesis results show the tradeoffs between hardware area and performance at varying OSRs. In pipelined mode, a 6MHz video stream can easily be handled at an OSR of 64, while occupying less than 8% of a Stratix-III device.

Key Words: **Ternary FIR filter, FPGA, Sigma Delta Modulation, VHDL Implementation.**

1. INTRODUCTION

While SWL (Short Word Length) systems have been in use over the last 40-50 years for analog-digital signal processing, their commercial use has increased since the advent of VLSI (Very Large Scale Integration) circuits. They are now found in numerous applications in both wired and wireless communication systems. For example, in [1] the conflicting requirements of high sampling rate, large dynamic range, and removal of the power interference before the amplification are overcome by the use of sigma delta modulation ADCs. One very attractive aspect of single-bit systems is their intrinsic simplicity of operation, low power consumption, stability and efficient hardware implementation.

So far, the applications of single-bit sigma delta modulated ($\Sigma\Delta M$) systems in mobile communications have tended to be restricted to the analog-to-digital conversion domain because it has proved difficult to perform complicated DSP tasks using 1-bit processing. However, in last two decades a new generation of SWL systems (1 or 2 bits) were developed that can perform general-purpose DSP functions, including classical filtering and adaptive LMS filtering. The application of such systems will lead to substantial reduction in hardware size and execution time. Short word-length systems do not require complex integer multiplication hardware [2-4], these being replaced by

* Ph.D. Student School of Electrical and Computer Engineering, RMIT, Melbourne, Australia.
** Senior Lecturer, School of Electrical & Computer Engineering, RMIT, Melbourne, Australia.
*** School of Engineering and Physical Sciences Heriot-Watt University Dubai Campus, UAE.

simple multiplexers. This is attractive for hardware implementation using and especially ASIC, as reducing the number of general-purpose digital multipliers in the chip is a major challenge in both these domains [5-7]. Technology scaling has allowed sigma delta modulation to operate above 10MHz with dynamic range specifications beyond 70dB, making the technique applicable to mobile handset receivers and transmitters [8].

Single bit systems become more efficient when their input values and the transverse FIR impulse response are all represented in binary or ternary format [2]. BT-TFF exhibit a basic architecture that is similar to the FIR transversal filter with coefficients in ternary format i.e. {-1,0 and +1}. However, the design of an efficient BT-TFF requires a number of coefficient taps proportional to the OSR and this may be a very large number. Thus, in common with all filters of this type, using a larger OSR will result in better filter performance at the expense of an increase in hardware along with a corresponding decrease in its maximum operating frequency.

In this paper we show the results of Matlab® simulations that illustrate the effect of OSR on the structure and performance of the BT-TFF. We also show the synthesis results for the filter targeting a small number of representative commercial FPGA devices using both pipelined and non-pipelined topologies to determine the area/performance tradeoffs implicit in these filters.

The remainder of this paper proceeds as follows. In Section 2, the design of FIR filters with BT coefficients and their Matlab simulations are discussed briefly. Section 3 describes the design, simulation and comparative results of the BT-FIR in VHDL. Finally, we summarize and conclude the paper and point to future work.

2. THE TERNARY FIR FILTER

The overall structure of the BT-FIR considered in this paper (Fig. 1) is essentially identical to its equivalent multi-bit filter and is formed from a tapped delay line followed by a coefficient multiplication stage and finally the addition of the partial products [1]. The detail of the filter hardware will vary with the application. For example, although the filter taps will always be ternary {-1, 0 +1}, the input data values, x(k), may either be single-bit (binary) or ternary. In either case, the multiplication stage reduces to either a simple bit-wise AND/OR logic block or a small LUT (Look-Up-Table).

As shown in Fig. 1, the FIR filter output y(k) is given by the convolution of the ternary taps h_i and the input signal x(k) as follows:

$$y(k) = \sum_{i=0}^{M} h_i, x_{k-i} \text{ with } h_i \in \{1,0,-1\} \quad (1)$$

where M is the order of the filter (=number of taps). The filter taps can be derived using the Sigma-Delta

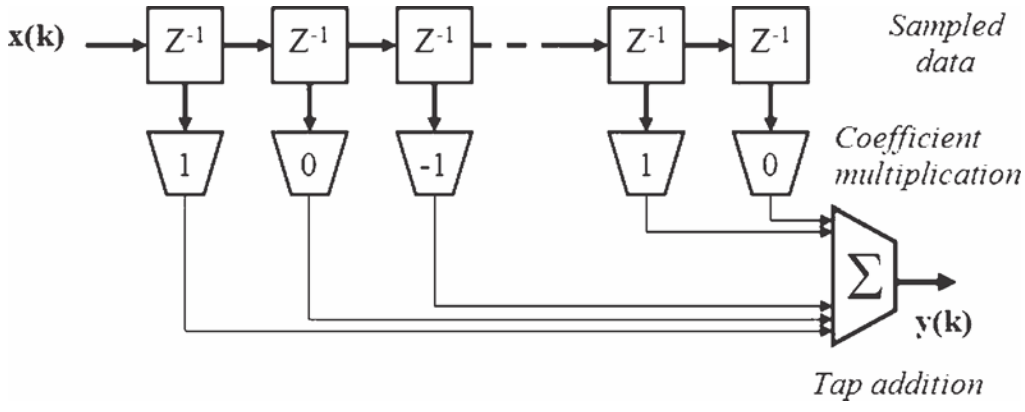

FIG. 1. BLOCK DIAGRAM OF TERNARY FIR FILTER (ADAPTED FROM)

Modulation (ΣΔM) of target impulse response. As their frequency response is directly related to the over-sampling ratio and therefore to the number of taps, these filter types typically require a very large number of coefficient taps (M). The ternary taps generated by ΣΔM must fulfil two mandatory requirements: ternary quantization and flat frequency response over the bandwidth of interest. A typical structure of a second order ΣΔM is shown in Fig. 2. The z-domain transfer function of this second order modulator is given by:

$$H(z) = G(z)z^{-1} + E(z)(1-2z^{-1}+z^{-2}) \quad (2)$$

where G(z) represents the target impulse response and E(z) represents the quantization noise transfer functions. The noise shaping effect of the ΣΔM is evident from the presence of the filtering term, $(1-2z^{-1}+z^{2})$ acting on the noise term, E(z). The frequency response of the above ΣΔM is given by:

$$H_{\Sigma\Delta T}(e^{j\Omega}) = G(e^{j\Omega})\, e^{-j\Omega} + E(e^{j\Omega})(1-2e^{-j\Omega}+e^{-2j\Omega}) \quad (3)$$

where $\Omega=2\pi f/f_s$ is the normalized frequency (radians).

In this work, the second order sigma delta modulator structure of Fig. 2 has been used for the generation of ternary taps. The process commences with the generation of a target impulse response by the Parks-McClellan best-design Remez exchange algorithm (Fig. 3). The passband and stopband edges of the target filter were set at 0.2π and 0.3π respectively with 60dB stopband attenuation. As the filter oversamples the input spectrum, it must be scaled before encoding into the ternary format so that the maximum input to ΣΔM operates within its maximum SQNR (Signal to Quantization Noise Ratio). Of many potential scaling and interpolation techniques, we have used the FFT method. After scaling (i.e. oversampling) the taps are ternary encoded.

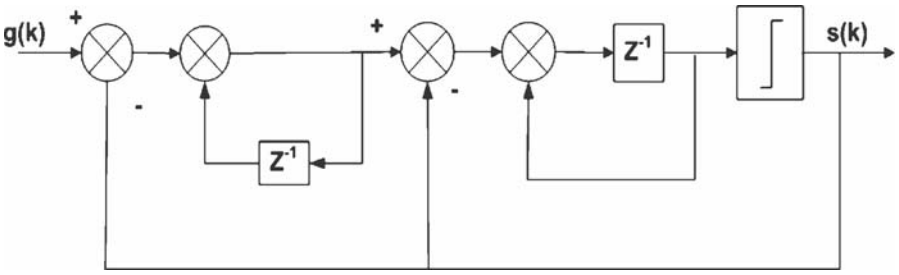

FIG. 2. SECOND ORDER M ARCHITECTURE USED FOR THE GENERATION OF SIGMA DELTA MODULATION

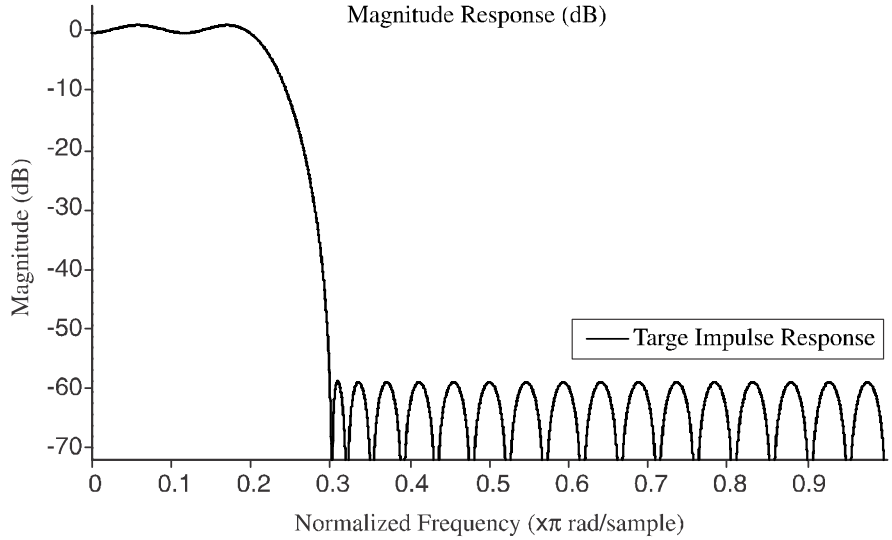

FIG. 3. TARGET IMPULSE RESPONSE BY REMEZ EXCHANGE ALGORITHM

The ternary filter generated via ΣΔM was simulated at varying OSRs. The corresponding impulse response was found to closely match the target impulse response of the ternary filter for each OSR. From Fig. 4, it can be seen that each doubling of the OSR results in an increase of about 10dB in the stopband attenuation, starting with 20dB at 32 OSR and reaching about 50dB at an OSR of 256. Note that the frequency axis in Figure 4 has been normalized in order to make the comparison clearer. This is necessary because varying the OSR implies a different absolute sampling rate for each (i.e. the Nyquist rate x OSR) so that although the relative characteristics of the filter will not change, the actual edge frequencies will be correspondingly different. A down sampling filter is used at the very end of the process to readjust the overall filter frequency responses to its original (Nyquist) rate.

3. TERNARY FIR FILTER DESIGN IN VHDL

As shown previously in Section 2, the operation of this ternary filter can be divided into two sections (Fig. 1) the multiplication of input data with the coefficients followed by the addition of the partial products. In the speech filter of [1], the number of taps (M) can be large: 1024 for an OSR of 32, or 2048 for an OSR of 64. Our implementation divides this up into N coefficient "multiply" blocks (i.e. by ±1,0) followed by an adder tree with $\log_2 N$ levels to perform the summation. Thus, in a binary implementation the addition requires 10 or 11 stages.

To perform signed arithmetic in FPGAs, 2's complement is a reasonable choice as the ternary taps {-1/0/+1} and binary data {-1/+1} can easily be represented as two bit numbers in two's complement format. Summing over N=1024 implies ten levels and a final multi-bit result of ±N. As the two's complement binary number range is offset around zero, 12 bits are required to completely express the full output range of ±1024. As mentioned above, SWL systems reduces the complex multiplication required for standard multi-bit processing to a simple AND-OR logic equation. In a FPGA context, the multiplication of two 2-bit numbers will map to a small LUT, potentially saving significant power and area compared to fully parallel implementations.

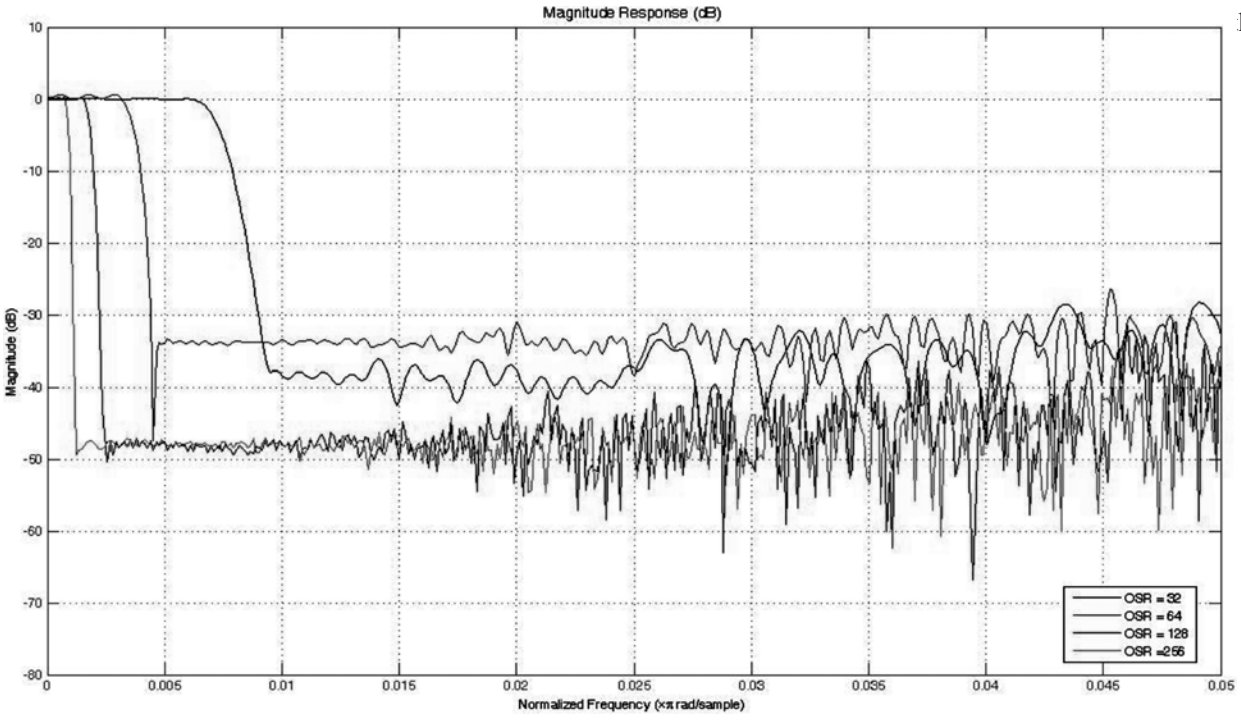

FIG. 4. TERNARY FILTER IMPULSE RESPONSE AT OSR=32, 64,128 AND 256

latency as well as improve the efficiency of this type of adder circuit. For example, a traditional Wallace adder tree is used in [9] and the VHDL code is generated by Matlab. In our work, we have focussed on techniques that map efficiently onto FPGA organizations but that are also suitable for ASIC implementation.

In the general overview of the filter structure shown in Fig. 5, there are 1024 data/coefficient pairs so the adder tree comprises ten levels. At each stage of the adder tree, the number of addition blocks halves while their length increases by one bit, culminating in the final 12-bit output. As a typical FPGA LUT structure takes a small number of inputs (in the range 6-8), the first two adder levels will be mapped to individual LUT blocks in the FPGA architecture that will operate in parallel. The remainder will comprise small ripple-carry blocks up to twelve bits long. Note that it would be equally possible to use optimised IP blocks created specifically for this purpose. In this paper, we have taken a more general approach, so that our implementation results can be considered to be worse-case.

3.1 VHDL Implementation

Both filter design alternatives were coded in VHDL and compiled, simulated and synthesized using Quartus-II 9.0 and Model Sim 5.8 targeting a small number of Altera™ Cyclone and Stratix FPGAs (Table 1), chosen as representative of a two "classes" of FPGA components that exhibit slightly different underlying architectures and performance vs. cost. The adder tree was synthesized in both pipelined and non-pipelined configurations. In the pipelined mode, registers are used in between each adder stage as illustrated in general terms in Fig. 5. As in all pipelined systems, the tradeoffs involve a small increase in hardware area and greater latency in exchange for increased throughput. It can be noted that the latency of the adder stage ($\log_2 N+1$) is small compared with the overall data latency across the filter shift registers, but its impact will be entirely application dependent.

In all of these tests, the operating frequency (F_{MAX}) target for the synthesis process was set to a value higher than achievable for the given technology in order to force the tools to generate a final routing that was comparable across devices. The approximate area values listed are those reported by the flow summary while F_{MAX} has its usual definition as the maximum clock rate (minimum period) achievable with zero slack on the worse-case critical path, as reported by the post-place and route timing analyser.

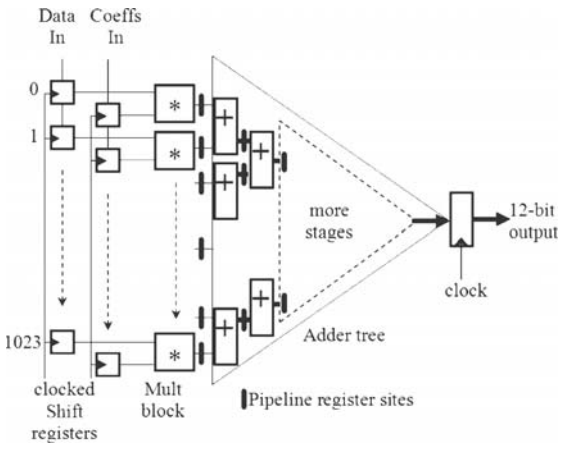

FIG. 5. BT-TFF HARDWARE ARCHITECTURE DESIGNED IN VHDL

TABLE 1. VHDL IMPLEMENTATION RESULTS

Device	OSR	Area: Non-Pipelined			Area: Pipelined		
		No. of LUTs	No. of Registers	FMAX (MHz)	No. of LUTs	No. of Registers	FMAX (MHz)
Cyclone-II EP3C120F48C7	32	7739 (6%)	4108 (3%)	60.1	7581 (10%)	10226 (9%)	202.1
	64	15521 (13%)	8205 (7%)	55.0	15,152 (13%)	20,465 (17%)	191.6
	128	31,113 (26%)	32,783 (28%)	04.9	30,423 (30%)	40,944 (34%)	169.7
	256	62,263 (52%)	32,783 (28%)	40.9	60,714 (51%)	81,903 (69%)	163.2
Statix-III EO3SK34F1517C2	32	6359 (2%)	4117(2%)	115.3	6,807 (3%)	10,226 (4%)	400.0*
	64	12,732 (5%)	8,221 (3%)	100.5	13,846 (5%)	20,441 (8%)	400.0*
	128	25,389 (9%)	16,459 (6%)	85.4	27,284 (10%)	40,944 (15%)	400.0*
	256	50,730 (19%)	32,874 (12%)	72.0	54,593 (20%)	81903 (30%)	356.6
* I/O constrained Clock Rate							

Although the implementations are directly comparable, the results may vary in a real application as no account was taken of pin capacitance and no specific optimizations were made, such as forcing the use of I/O registers. In a few cases as noted below, F_{MAX} exceeds the highest rate achievable at the I/O pins (typically limited to less than 400MHz).

4. CONCLUSIONS

In this paper we have presented an analysis of a novel balanced ternary FIR filter structure at varying OSR. Ternary impulse responses were obtained at varying OSRs using Sigma Delta Modulated FIR filter calculations in Matlab showing that each doubling of the OSR results in an increase in the stopband attenuation of around 10dB. The absolute stopband attenuation was around 20dB at an OSR of 32 and reached about 50dB at an OSR of 256.

The various filters were implemented in VHDL using a hierarchical adder tree organization to gather the partial products. Both pipelined and non-pipelined modes were investigated to show the tradeoffs between hardware area and performance at varying OSRs. Using a high performance device (e.g. Stratix III), the filter can operate at clock frequencies in excess of 400MHz (actually beyond the capability of the chip package). We note that a 400MHz sample rate will easily handle a 6MHz video stream at an OSR of 64, while occupying less than 8% of the available FPGA resources. Even using a low-cost device family (Cyclone III), worse case operating speeds over 160MHz are possible in pipelined mode, although a large filter (e.g. OSR=256) may occupy most (up to 70%) of the FPGA resources.

This ?M -based balanced ternary FIR filter can achieve high performance without requiring the use of built-in DSP components such as fast parallel multipliers. The technique is thus well suited to ASIC implementation as well as to low-cost FPGA architectures that do not offer these elements as part of their architecture. To date, this work has focussed on the design and simulation of the ternary FIR filter structure using fixed coefficients. In future work, we will "close the loop" on this organization by adaptively adjusting the coefficients using information extracted from the data stream.

ACKOWLEDGEMENT

The authors would like to express their thanks to Prof. Peter O'Shea, and Mr. Zahir Hussain, Associate Professor, who provided assistance in single-bit domain DSP applications.

REFERENCES

[1] Thompson, A.C., O'Shea, P., Hussain, Z.M., and Steele, B.R., "Efficient Single-Bit Ternary Digital Filtering Using Sigma-Delta Modulator", IEEE Letters on Signal Processing, Volume 11, pp. 164-166, February, 2004.

[2] Thompson, A.C., Hussain, Z.M., and O'Shea, P., "Efficient Digital Single-Bit Resonator", IEEE Electronic Letters, Volume 40, pp. 1396-1397, 2004.

[3] Thompson, A.C., Hussain, Z.M., and O'Shea, P., "Performance of a New Single-Bit Ternary Filtering System", Australian Telecommunications, Networks and Applications Conference, Melbourne, 2003.

[4] Sadik, A.Z., Hussain, Z.M., and O'Shea, P., "An Adaptive Algorithm for Ternary Filtering", IEE Electronics Letters, Volume 42, pp. 420-420, March, 2006.

[5] Rusu, A., Rodriguez de Llera Gonzalez, D., and Ismail, M., "Reconfigurable ADCs Enable Smart Radios for 4G Wireless Connectivity", IEEE Circuits and Devices Magazine, Volume 22, pp. 6-11, 2006.

[6] Memon, T.D., Beckett, P., and Hussain, Z.M., "Design and Implementation of Ternary FIR Filter Using Sigma Delta Modulation", ISCCC'09, 2009.

[7] Memon, T.D., Beckett, P., and Sadik, A.Z., "Performance-Area Tradeoffs in the Design of Short Word Length FIR Filter", 5th IEEE International Conference on MEMS, Nano and Smart Systems, Dubai, UAE, 2009.

[8] Galton, I., "Delta-Sigma Data Conversion in Wireless Transceivers", IEEE Transaction on Microwave Theory and Techniques, Volume 50, pp. 302-315, January, 2002.

[9] Chen, D., "VHDL Implementation of Fast Adder Tree", Department of Electrical Engineering, Linkoping University, 2005.