

UNIVERSITY OF WESTMINSTER



WestminsterResearch

<http://www.wmin.ac.uk/westminsterresearch>

Reconfigurable implementation of recursive DCT kernels for reduced quantization noise.

Suleyman Demirsoy
Robert Beck
Andrew Dempster
Izzet Kale

Cavendish School of Computer Science

Copyright © [2003] IEEE. Reprinted from Proceedings of the 2003 International Symposium on Circuits and Systems (ISCAS '03), pp. 289-292.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Westminster's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

The WestminsterResearch online digital archive at the University of Westminster aims to make the research output of the University available to a wider audience. Copyright and Moral Rights remain with the authors and/or copyright owners. Users are permitted to download and/or print one copy for non-commercial private study or research. Further distribution and any use of material from within this archive for profit-making enterprises or for commercial gain is strictly forbidden.

Whilst further distribution of specific materials from within this archive is forbidden, you may freely distribute the URL of WestminsterResearch. (<http://www.wmin.ac.uk/westminsterresearch>).

In case of abuse or copyright appearing without permission e-mail wattsn@wmin.ac.uk.

RECONFIGURABLE IMPLEMENTATION OF RECURSIVE DCT KERNELS FOR REDUCED QUANTIZATION NOISE

Süleyman Sirri Demirsoy, Robert Beck, Andrew G. Dempster and Izzet Kale

Applied DSP and VLSI Research Group, Department of Electronic Systems
University of Westminster, 115 New Cavendish St, London, W1W 6UW, UK

Tel: +44 20 7911 5000 - 3611(ext) e-mail: demirss@cmsa.wmin.ac.uk

ABSTRACT

Time multiplexed implementations of the recursive DCT processors are widely used in many multimedia and compression applications. Recently proposed three Goertzel kernels offer significant improvement (up to 90 %) in the noise performance of the time-multiplexed architecture to allow word-length specifications get reduced. In this paper, a highly optimized reconfigurable DCT architecture is proposed that can perform the function of three different kernels (Type A, B and C) on Virtex FPGA.

1. INTRODUCTION

Recursive DCT implementations are attractive due to their regular structure and reduced computational complexity. The transfer function of the DCT given in (1) can be implemented with a second order IIR filter. The kernel employed in most of the designs in the literature [6], [7] is a resonator of Type B configuration as shown in Figure 1(a). In [2], two alternative forms, Type A and Type C were proposed (Figure 1b and 1c).

$$H(z) = \frac{P_k(1 - z^{-1})}{1 - 2\beta_k z^{-1} + z^{-2}} \quad (1)$$

$$P_k = (4L_k / N) \cos(k\pi / N), \quad k: \text{frequency bin index}$$

$$L_k = 1/\sqrt{2} \text{ for } k=1, L_k = 1 \text{ for all other } k$$

$$2\beta_k = 2\cos(k\pi / N) \quad N: \text{transform length}$$

The benefits of these structures are investigated in [1] and substantial area gains in a fully parallel implementation were demonstrated. In that design, Type A, Type B and Type C structures were all employed for different frequency bins to have the optimum coefficient magnitude. Figure 2 shows how these magnitudes vary for different frequencies. By using Type A for the first one-third of the filter bins, Type B for the next one-third and Type C for the last one-third of the filter bins, the optimum coefficient values are used.




Figure 1 Recursive Goertzel filters, a) Type B, b) Type A, c) Type C

For a time-multiplexed system, where only one recursive structure is used for all frequency bins by multiplexing it in time, Type B is the natural choice to employ for having the average coefficient magnitude among the three types. However, if a reconfigurable structure that can be transformed into three different types for different frequencies is developed with *minimum reconfigurability overhead*, a *significant improvement* on internal quantization noise figures can be achieved.

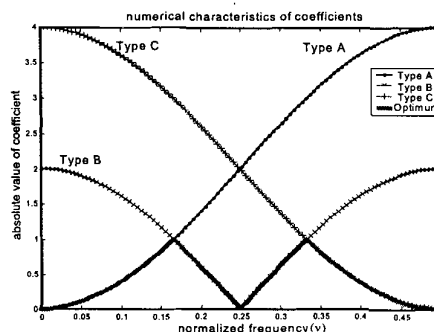


Figure 2 Coefficient magnitude behaviour of Type A, B and C structures vs. frequency. The optimum behaviour is also marked.

FPGAs are attractive platforms for easy and fast implementation and offer high performance when the available resources are used efficiently. The effective use of the Look-Up Tables (LUT) available on-board in Xilinx Virtex FPGAs is investigated in [3].

In this paper, a reconfigurable recursive kernel that is highly optimized for the Xilinx Virtex FPGA series is given and investigated for its noise performance. Section 2 describes the details of the Virtex slice architecture and the optimization steps of the recursive kernels. Section 3 gives the noise performance analysis and Section 4 concludes the paper.

2. OPTIMIZATION

The Virtex FPGA series comprise Configurable Logic Blocks (CLB), which are the main programmable functional blocks. Each CLB contains two slices each of which has two of the structures given in Figure 3a. Fast carry logic is implemented in each half-slice by MUXCY. The Look-up Table (LUT) is a 16-bit SRAM with 4 inputs. It can be programmed to realise any logic function having up to four inputs. Figure 1b and 1c shows two examples for a 1-bit adder subtractor implementation on each half-slice, and a subtractor with a multiplexer choosing one of its two inputs for subtraction from A. The half-slice also has a D-type flip-flop that can function like an output register.




Figure 3 (a) Half of the Virtex Slice diagram (simplified), (b) A combination of XOR gates implemented in the LUT to make an adder/subtractor, (c) a multiplexer choosing the input to subtract from A

Any circuit optimization for a given device with limited resources would be a job of facilitating the available resources as much as possible, bearing in mind that the required operation can be satisfied by a modification on the structure that allows fitting the resources in a better way while maintaining the functionality. For the reconfigurable DCT kernel of this paper given in Figure 4a, there are several multiplexers (MX1 to MX5) which allow us to use the same components with different inputs. The letters A, B and C indicate the Type of the kernel a line is required for. The signs on top of the letters are the functionality that is required in the adder/subtractors (A1 to A4) for that input when the specified kernel is configured. The loop multiplier (M1) takes the coefficient of all kernels. The optimization is performed as follows:

The adder/subtractor A1 and the delay element D2 can fit into a half-slice if the multiplexer MX1 does not exist. Moreover, the adder/subtractor A4 consumes half a slice by itself without making use of the delay element. By replicating D2 as D2b and storing the output of A4 in D2b, it is possible to get rid of MX1. (See Figure 4b)

MX3 requires an extra LUT. However by simply feeding A3 with '0' from MX5, the output for Type A can be achieved via A3 from the same output as Type B and C. This implies we have three different outputs from MX5 (Figure 4c).

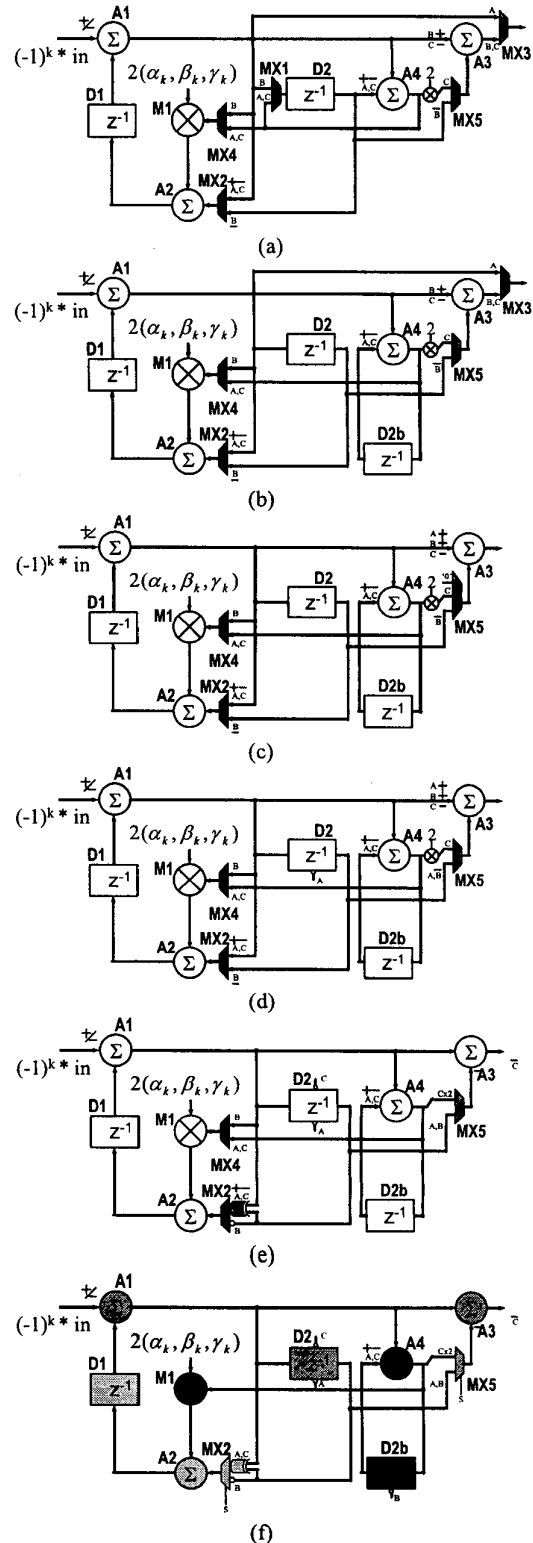


Figure 4 Optimization steps of the reconfigurable kernel

It is possible to generate the '0' signal by clearing D2. This reduces the number of necessary inputs to MX5 down to two. (Figure 4d)

The input to A5 coming out from A1 requires negation for type C and the other input needs to be negated for Type B too. This functionality of negation on both inputs for different kernels is not possible to implement in a half-slice. However by performing negation for Type C at the same input as Type B, it is possible to fit MX5 and A3 to the same half-slice. Hence the output of A3 for the Type C kernel needs to be negated. It is done by using the negated coefficients in the Pk multiplier outside the recursive loop shown in Figure 1.

The multiplication by 2 is performed by hardwiring the output of A4 one bit shifted toward the MSB.

The output of the MX2 needs to be negated for Type B and C. In its current form, it is not possible to fit MX2 with A2 to the same half-slice. It is known that for Type A, D2 generates a '0' signal for MX5. By generating a '1' signal on D2, the output of D2 can be used as a select signal for the operation on the other input to MX2. This modification allows fitting MX2, A2 and D1 to the same half-slice. (See Figure 4e)

MX4 requires an extra LUT. It is possible to get rid of this multiplexer by using the output of A4 for Type B kernel as well. A '0' signal generated by D2b would allow having the A1 signal at the output of A4, after adding a small amount of delay. (Figure 4f)

After all the modifications are performed, the reconfigurable DCT kernel occupies four half-slices where Type B kernel uses three half slices. This overhead is easily compensated by the reduction in the coefficient word-length of multiplier M1. Hence, if a general multiplier is used for M1, the reconfigurable structure needs the same number of half-slice as Type B. For the choice of fixed or Reconfigurable Multiplier Blocks (ReMB) [8], the necessary coefficients are smaller in magnitude and are constructed using less adder stages. An example to this fact is given in [8] for a coefficient word-length of 12-bits. The loop multiplier M1 for Type B kernel is constructed with four basic structures, whereas the loop multiplier for the reconfigurable kernel required only three basic structures.

3. NOISE PERFORMANCE

The quantization noise exists in a circuit due to the fixed word-length of the intermediate and output signals. The performance of the system is affected by the choice of the signal word-lengths. Several restrictions on the maximum quantization noise are defined by the standards to avoid performance degrade below a limit. For example, MPEG-4 standard defines the maximum allowed MSE, mean error and magnitude error that occurs in an 8 by 8 block of pixels in a video data for Two Dimensional Inverse DCT (2D-IDCT) [5].

The coefficient values required by the Type B kernel for the frequencies close to DC and half-Nyquist are higher than the values required by Type A near DC and Type C near half-Nyquist (see Figure 2). An extra bit for the integer part of the coefficient is required if just the Type B coefficients are used. Therefore by using the Type A and C kernels for these frequency

bins, we save one bit for the same fractional precision. On top of that, the analysis of the integer parts of the intermediate signals shows that Type B kernel requires more bits for the integer parts of the signals than the Type ABC kernel does. Table 1 shows how many bits are required to represent the integer parts of the signals for a transform length of N=16. The reconfigurable kernel, Type ABC, would spare more bits for the representation of the fractional parts at the output of A1, A2, and M1, hence would lead to better noise performance.

Type	A1	A2	A3	M1	Loop	Pk	out	A4
B	15	15	16	16	2	0	11	-
ABC	13	13	16	13	1	0	11	15

TABLE 1 Number of bits required for the integer parts of the signals at the output of the specified components for N=16. Loop and Pk are the coefficient values for the loop multiplier (M1) and Pk multiplier (see Figure 1)

A set of experiments were performed to demonstrate the enhanced performance. For various word-length specifications, both Type B and Type ABC design were stimulated with a thousand uniformly distributed random numbers between -300 and 300. A reference design with floating point precision was used to find the quantization noise generated by the two designs. Our error measure for these experiments was the Mean Square Error (MSE). Table 2 shows the results for different specifications and transform-lengths (N). The loop coefficient word-length for Type B design (loopb) is always kept one bit more than the loop coefficient of the Type ABC design (loopa) to maintain the same area consumption when a general-purpose multiplier is used. Pk is the coefficient word-length for the Pk multiplier. The frequency bins k=0 and k=4 are neglected during the calculation of the percentage reduction in noise because these frequency bins only have noise contributions from the Pk multiplier which is outside the recursive kernel. The gain for each frequency bin is different. The maximum and minimum gains that occurred for a set are also shown in Table 2. For particular frequency bins, there is 93 % decrease in the noise. The average gain is a reasonable measure of the overall noise performance enhancement. It is observed that, there is up to 67 % decrease in the MSE levels among the given sets.

Word-length specifications	N	Max	Min	Avg
1. wl=18; loopa=17; loopb=18; Pk=16	8	74%	41%	60%
2. wl=20; loopa=18; loopb=19; Pk=16	8	71%	28%	35%
3. wl=20; loopa=18; loopb=19; Pk=16	16	88%	59%	67%
4. wl=22; loopa=21; loopb=22; Pk=18	32	93%	17%	67%

TABLE 2 The percentage decrease in the MSE noise with reconfigurable Type ABC design as opposed to Type B design.

Figure 5 shows the MSE noise levels for the design given in the 1st set of Table 2. As seen from the figure, although the loop

coefficient precision for $k=3$ and $k=5$ are same for Type B and Type ABC designs, the difference of the magnitude of the intermediate signals leads to an enhancement in the noise performance of Type B kernel outputs in the reconfigurable design.




Figure 5 Comparison of MSE noise figures for $N=8$ and word-length of 18 bit. The coefficient word-lengths in Type B-only and reconfigurable kernel are 18 bit and 17 bit respectively. An average decrease of 67% is achieved.

4. CONCLUSION

A reconfigurable recursive DCT processor, that can compute three kernels (Type A, B and C) has been designed and optimized for Xilinx Virtex FPGA series. It occupies only one more half-slice than Type B kernel does for a single bit, due to high utilization of the slice resources. The loop multiplier coefficient requires one bit less for the same precision. For the same half-slice count, the MSE level is reduced by 67 % on

average of all frequency bins in the proposed design when compared to the Type B kernel. The optimized structure includes design ideas that can be applicable to VLSI implementation too. Future work will focus on the pipelined implementation of this reconfigurable structure to increase the operating clock frequency.

5. REFERENCES

- [1] Demirsoy S., et al., "Novel recursive DCT implementations: A comparative study", *IEEE Int. Conf. on Intelligent Data Acquisition and Advanced Computing Systems (IDAACS'2001)*, pp.120-123, Ukraine, July 2001
- [2] Beck, R. "An Investigation of Finite-Precision Digital Resonators", PhD Thesis, University of Westminster, June 2002
- [3] Turner R. H., T. Courtney and R. Woods, "Implementation of fixed DSP functions using the reduced coefficient multiplier", *IEEE Proc. of ICASSP'2001*, vol. 2, pp. 881-884, May 2001, USA
- [4] http://www.xilinx.com/xlnx/xil_prodcats/landingpage.jsp?title=Virtex_Series
- [5] ISO/IEC, "Information technology-Coding of audio-visual object: Visual ISO/IEC 14496-2 Final Proposed Draft", 14496-2, July 1999
- [6] Wang J.L. et al, "Implementation of the DCT and its inverse by recursive structures", *IEEE Workshop on Signal Processing Systems*, pp 120-130, Oct 1999
- [7] Kozick R.J., and M.F. Aburdene, "Methods for designing efficient parallel -recursive filter for computing discrete transforms", *Telecommunication Systems*, vol. 13, no.1, 2000, pp.69-80
- [8] Demirsoy S. S., A. Dempster and I. Kale, "Design Guidelines for Reconfigurable Multiplier Blocks", submitted to *IEEE ISCAS 2003*