

# A New Class of High Performance FFTs

Dr. J. Greg Nash,  
Centar, jgregnash@centar.net

## Introduction

An FPGA implementation of a block floating point (BFP), streaming, 256-point FFT is described as an example of a new class of parallel FFT architectures that can provide better performance and functionality than commercially available pipelined FFTs. It is based on a row/column factorization of the discrete Fourier transform (DFT) coefficient matrix along with an index remapping that converts the direct transform into structured sets of arithmetically simple 4-point transforms that are computed on a systolic array [1].

## Mathematical Foundation

The index remapping starts with the DFT defined as  $Z(k) = \sum_{n=0}^{N-1} W_N^{nk} X(n)$  where  $X(n)$  are the time domain input values,  $Z(k)$  are the frequency domain outputs and  $W_N = e^{-j(2\pi/N)}$ . In matrix terms an alternative representation is  $Z = CX$  where  $C$  is a coefficient matrix containing elements  $W_N^{nk} = e^{-j(2\pi nk/N)}$ . If  $N$  can be factored as  $N = 4N_1$  ("base-4" form) with  $N_1$  divisible by 4, then using the reindexings  $n = n_1 + N_1 n_2$  and  $k = k_1 + N_1 k_2$  with  $n_1 = 0, 1, \dots, N_1 - 1$ ,  $k_1 = 0, 1, \dots, N_1 - 1$ ,  $n_2 = 0, 1, 2, 3$ ,  $k_2 = 0, 1, 2, 3$ , then  $Z$  can be obtained from the matrix equations

$$\begin{aligned} Y_b &= W_M \bullet C_{M1} X_b \\ Z_b &= C_{M2} Y_b^t \end{aligned} \quad (1)$$

where  $W_M$  is an  $N_1 \times N_1$  matrix with elements  $W_M[k_1, n_1] = W_N^{n_1 k_1}$ ,  $C_{M1}$  is an  $N_1 \times 4$  coefficient matrix with elements  $C_{M1}[k_1, n_2] = W_4^{n_2 k_1}$ ,  $X_b$  is an  $4 \times N_1$  matrix with elements  $X_b[n_2, n_1] = X(n_1 + N_1 n_2)$ ,  $C_{M2}$  is a  $4 \times N_1$  coefficient matrix with elements  $C_{M2}[k_2, n_1] = W_4^{n_1 k_2}$ ,  $Z_b$  is an  $4 \times N_1$  matrix containing the transform outputs  $Z_b[k_2, k_1] = Z(k_1 + k_2 N_1)$ , and " $\bullet$ " means element-by-element multiply.

The computational advantages of the manipulation leading to the matrix algorithm form (1) are evident when compared to the traditional direct form  $Z = CX$ . First the matrix products  $C_{M1}X$  and  $C_{M2}Y_b^t$  involve only exchanges of real and imaginary parts plus additions because the elements of  $C_{M1}$  and  $C_{M2}$  contain only  $\pm 1$  or  $\pm j$ , whereas the product  $CX$  requires complex multiplications. Also, the size of the coefficient matrix

$W_M$  in (1) is  $(N/4) \times (N/4)$  vs. the  $N \times N$  size of  $C$ ; consequently the number of overall direct multiplications is reduced by a factor of x16 compared to the direct form on which past systolic FFT implementations are based. Note that distribution of the elements in  $C_{M1}$  and  $C_{M2}$  does not impose significant bandwidth requirements because full complex numbers are not used. More details are provided in [1].

## Architecture

The base-4 circuit design described here (others are possible) contains two separate  $N_1 \times 4$  systolic adder arrays that compute the matrix products  $C_{M1}X$  and  $C_{M2}Y_b^t$  and a single  $N_1$  element linear array in between the two to do the complex multiplies by  $W_M$ . The single biggest drawback to past use of systolic arrays has been the substantial arithmetic hardware that is required because systolic approaches use a number of complex multipliers equal to the size of the transform. Thus, a 256-point DFT would require 256 complex multipliers, compared to only 4 complex multipliers used in a base-4 implementation.

Compared to traditional pipelined FFTs, the base-4 architecture possesses the following advantages:

- *It can compute the DFT for any N-point sequence divisible by 256.* In contrast traditional FFT circuits require  $N$  to be a power of two, which limits the number of reachable values of  $N$  and their spacing uniformity.
- *It uses fewer clock cycles per transform.* The number of clock cycles per DFT is  $\approx N/2 + 4\sqrt{N} + 40$  compared to  $N$  for a typical pipeline FFT.
- *Higher clock rates are possible* because
  - the circuit consists of a regular array of locally connected small processing elements, each PE containing only a few registers, an adder and memory.
  - the circuit architecture matches that of FPGAs with their embedded linear arrays of hardwired multipliers and memory.
  - there are only three global lines (2 clocks and a global clear)
  - a large number of smaller, faster and more power efficient memories are used
- *It can do both 1-D and 2-D DFTs.*

- Any FFT implementation can do any FFT size. This flexibility also provides a strategy to match circuit architecture to required application throughputs.
- It offers better dynamic range (DR) and signal-to-noise (S/N) due to an enhanced BFP circuitry. This feature is a result of having many different block floating point circuit regions, i.e., each region has its own exponent.
- It provides low latency as well as high throughput. Computational latency in cycles/DFT is only slightly higher than the same throughput measure.
- The design is scalable and reconfigurable. Larger FFTs are obtained by replicating identical blocks of PEs.
- Design, testing and maintainability are simplified because the circuit is based on only two small PE types.

### Streaming 256-Point FFT Design

To demonstrate this base-4 architecture a design that accepts a continuous input stream  $X(n)$ , while generating a continuous output stream  $Z(n)$  at the same rate (“streaming”) was chosen since this mode is common to many signal processing applications. To add application generality a BFP capability was added and a word size of 16-bits was chosen since this a common choice. The design has circuit pins for real and imaginary inputs/outputs,  $Z/X$ , a single global reset, and two clocks. Because the base-4 design computes FFTs using a number of clock cycles that is less than the transform size, a separate higher speed clock is used to read out the data.

An evaluation of the base-4 design is best done by comparison with another start-of-art pipelined FFT design built to the same specifications, using exactly the same underlying hardware. For this purpose a streaming BFP 256-point FFT design from Altera was chosen (FFT v2.2.0), since their pipelined BFP FFT circuits are the fastest of which we are aware. Both designs were targeted to an Altera Stratix II EP2S15F484C3 FPGA. Altera Quartus II tools (v5.1) were used to design and evaluate the two FFT circuits. The base-4 circuit operation was verified by comparing the Quartus simulator result with a Centar bit-accurate simulation model. The Quartus II timing analyzer finds the critical path that determines the maximum clock frequencies.

The base-4 and Altera results are summarized in Table 1. Because the base-4 design provides higher dynamic range for a given bit-length, a 20-bit Altera FFT was used in the comparison. To demonstrate this both the dynamic range and a signal-to-noise ratio of each circuit are included in the table. Each entry in the table is the mean value obtained from 62 different single frequency, full range, input data sets (random frequency and phase). The results for the Altera circuit were obtained from a bit-accurate Matlab model that’s created by the Altera FFT generator.

There was no noise added to the single frequency inputs, so the “noise” represents only internally generated round-off noise. Dynamic range was determined from the ratio of the maximum output coefficient and maximum noise output value.

For the Altera design the clock speed listed in Table 1 is also the complex sample rate. For the base-4 design the complex sample rate is larger than the clock speed in the table by a ratio of 256/240 or 389 MHz. However, the timing simulations show that the maximum read rate could be as high as 429 MHz.

Category	Altera (20 bit)	Base-4 (16-bit)
Throughput (cycles/DFT)	256	240
Clock speed (MHz)	302	365
Throughput ( $\mu$ sec)	0.85	0.66
Dynamic Range (db)	98.2	99.7
Signal/Noise (db)	86.9	90.2
Total ALUTs	7555	7731
ALMs	4192	4286
Memory Bits	48640	78708
18-bit multipliers	12	16

Table 1. Comparison of Altera and base-4 FFT circuits.

In the table the number of “adaptive logic modules” (ALMs) is included because this value is roughly comparable to a Xilinx “slice”. The base-4 memory size is significantly higher, but almost half of the difference is due to the fact that the control store consists of completely unrolled code. The control code is repetitive, so added control structures here could improve this value.

This comparison shows that for a 256-point transform the base-4 architecture can provide a better throughput with comparable FPGA resource usage than traditional pipelined FFTs for applications that require a high dynamic range. Similar comparisons can be expected for other FFT sizes. For example a base-4 1024-point circuit uses only ~680 clock cycles per DFT vs. 1024 for pipelined FFTs. Because the base-4 structure is local and regular, similar clock speeds as shown in Table 1 should be expected, yielding a complex sample rate of ~542 MHz. Alternatively, timing analyses of even a smaller 16-bit Altera 1024-pt FFT indicate the complex sample rate would only be ~325MHz.

Finally, this base-4 circuit design can also do 2-D DFTs and with additional memory it could do larger FFTs, including DFTs where  $N$  is not a power of 2.

### References

[1] J. G. Nash, “Computationally efficient systolic architecture for computing the discrete Fourier transform,” IEEE Transactions on Signal Processing, Volume 53, Issue 12, Dec. 2005, pp. 4640 – 4651.