RF Engines Limited

White Paper

Pipelined FFT

Reference : W02004-Pipelined FFT White Paper
Revision : 1.1
Date : 04 Oct 2002
Author : IMV

Tel : +44 (0)1983 550330
Fax : +44 (0)1983 550340
E-mail : info@rfel.com
Web : www.rfel.com

RF Engines, Innovation Centre, St Cross Business Park, Newport, Isle of Wight, England, PO30 5WB

RF Engines Pipelined Complex FFT Core for Xilinx Virtex E and II FPGA

Introduction

This is an initial specification sheet for a brand new core being offered by RF Engines Limited (RFEL). It is a fast Pipelined FFT that will continuously process data at a sample rate in excess of 100 MSPS, the complex 4096-point version fits in a single 1M gate FPGA. This document provides details of the basic core, optional items available and our roadmap for development of a full range of multi radix architecture, high specification FFT cores.

These cores are intended for use in applications where processing speed is critical and optimum use of available silicon is required. The cores are fully pipelined for maximum data throughput and complement our range of Pipelined Frequency Transform (PFT) products. The cores are available for licence in net list or bit stream form.

Block Diagram:

The first product in the range is a Radix 2 DIF (Decimate in Frequency) Pipelined FFT. This initial product is optimised for Xilinx Virtex/E and will also work well in VirtexII devices. A new butterfly architecture that takes advantage of the low power dedicated multipliers available in VirtexII will be available shortly. The FFT core is also suitable for implementation in other high performance FPGAs that have sufficient internal RAM to support the designs, such as the Altera Stratix announced recently. The diagram below shows the basic architecture employed in the core.

It is based on successive n stages, where 2^n is the size of FFT. Each stage has switched delay elements and butterflies. The switches and delays re-order the data for processing at the next butterfly. There are n butterflies which implement the complex arithmetic, each performs a 2-point DFT and complex phase rotations (twiddles). The input to the first butterfly has a FIFO n/2 buffer stage, which ensures efficient utilisation of the butterfly arithmetic. The output of the final stage is bit reversed complex (I/Q) data.

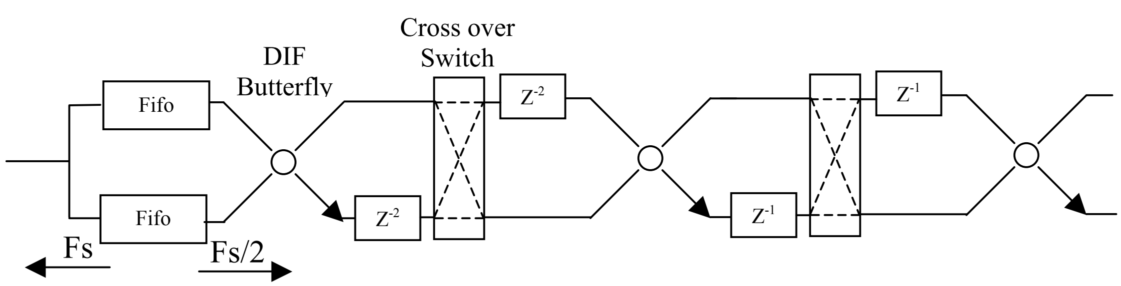

Block diagram of a simple 8 point pipelined DIF FFT

Overall Configuration

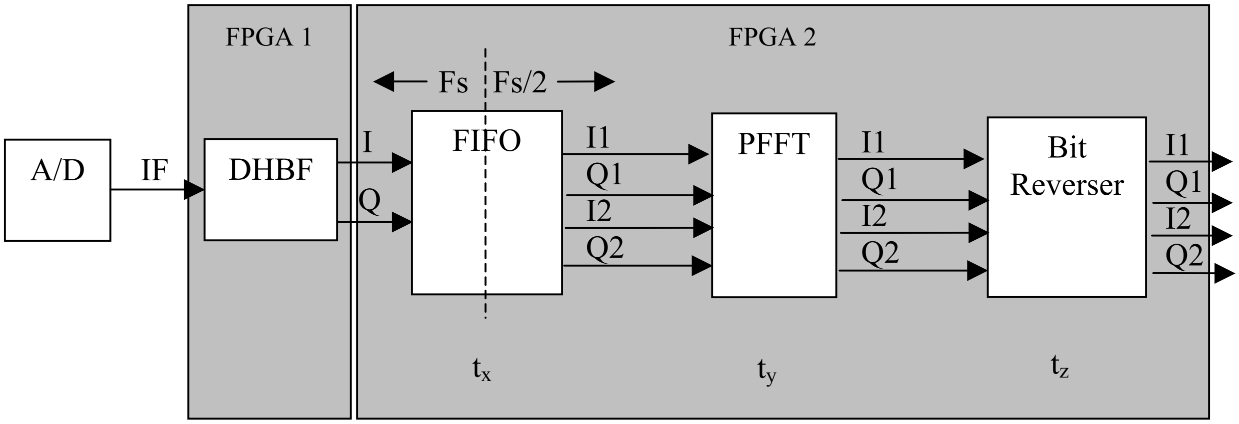

Typical FFT configuration

The Pipelined FFT design has been implemented and tested in the RFEL development system hardware, which has 4 x Virtex1000E-ehq240 FPGAs connected in a pipeline, with the first device fed by an 8 bit A/D with max Fs of 256MHz.

The tested design is a 12 stage (4096 point) Radix 2 DIF complex (I and Q input) FFT. The FFT has an 8 bit I input and 8 bit Q input, increasing by 1 bit per stage until 16 bits and a twiddle resolution of 16 bits.

The FFT core is fed by a distributed half band filter (DHBF) which converts a real IF into a complex baseband signal at up to 128 MSamples / second (MSPS), complex. The DHBF is implemented in the first FPGA of our development system and utilised less than 10% of the Virtex 1000E with no block RAM usage. This design of DHBF could be included in the same FPGA as the FFT design. RFEL can also supply the DHBF as a Licensable IP core.

The FFT core has normally ordered inputs and bit reverse ordered outputs. An optional bit reverser is implemented in the same device to provide a normally ordered output.

Silicon Size and Speed

This whole 4096-point FFT design fits into a single Xilinx Virtex1000E-ehq240, using 8478 slices (68% of the logic resource), 90 out of the 96 block RAM available (these figures include the bit reverser). The bit reverser uses 32 block RAMs and a small amount of logic resource, so if this function were not required the total block Ram used would be 58 out of 96, thus leaving space for other processing.

The FFT core can sustain a constant pipelined data rate of 128 MSPS, complex, which is the maximum the RFEL development system can support, but timing reports from the place and route tools suggests operation up to 160 MSPS (RFEL are carrying out further architecture optimisations which we are confident will increase this even further, to around 200 MSPS).

Processing Delay

One of the most important parameters of the Pipelined FFT is the time it takes to process a block of data or the processing delay. First of all we must describe the meaning of processing delay. In the case given below it is the time between the last sample entering the FIFO and the last sample leaving the bit reverser.

The timing diagram below gives an example of a simple 64 point Pipelined FFT.

In the FIFO buffer, the 64 point block of data at Fs is split into two blocks and stored in the two FIFO's. The reading of data block 'a', is held off until data block 'b' is read, but the data is read at Fs/2. Successive stage's delay elements are half the previous stage until the last stage, which has a delay of one sample. It can be seen that the processing delay takes 1.5 blocks of data in time for the processed data to reach the outputs of the final butterfly.

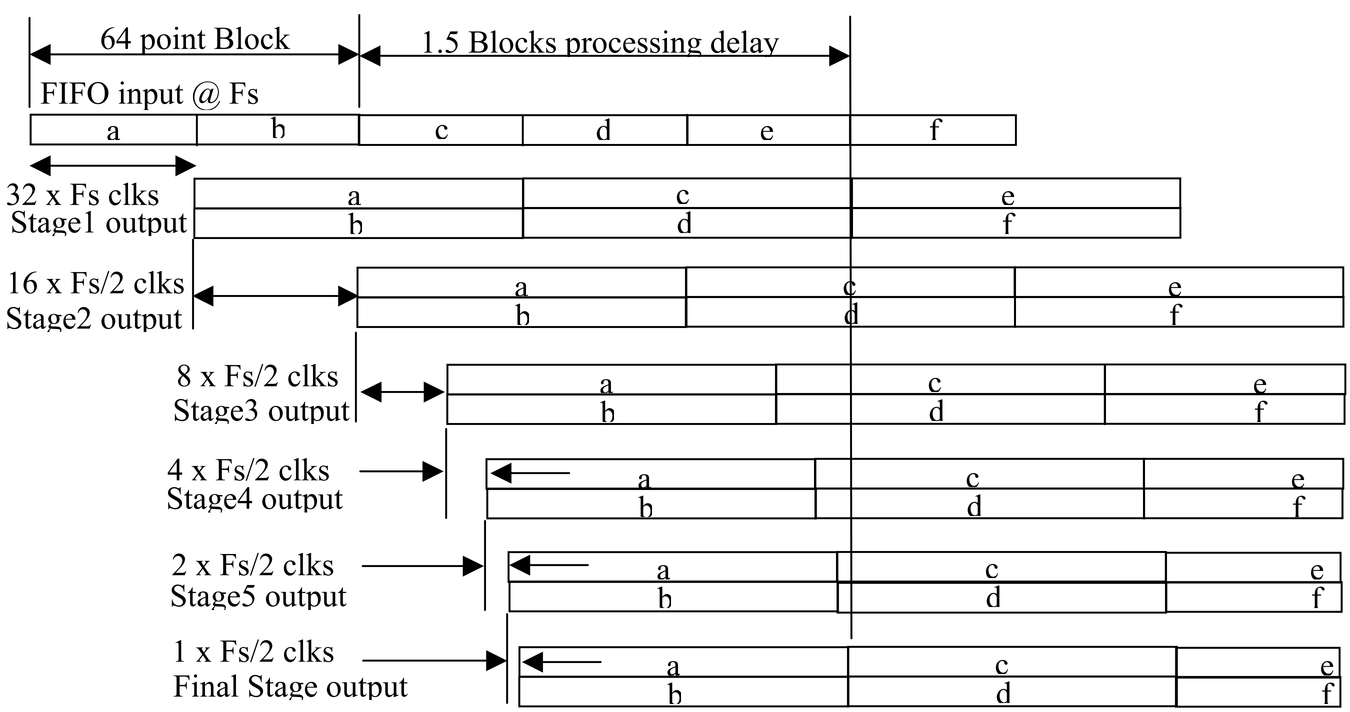

The processing delay (last sample in to last sample out), can be calculated by :

$$1.5 \times n \text{ (clocks at Fs)}$$
Where n = FFT length

As well as the fundamental processing delay for a Radix 2 FFT, additional pipeline delays due to the complex arithmetic are applicable.

The pipeline delays of the cascaded FFT stages (t_y), is 8 clocks at Fs/2 where

$$t_y = \text{stages} \times 8 \times 2 \text{ (clocks at Fs)}$$

The pipeline delay of the Switch (t_w),

$$t_w = 2 \text{ Fs clocks} \times \text{stages}$$

The pipeline delay of the FIFO stage (t_x),

$$t_x = 1 \text{ Fs clock}$$

The delay for the optional bit reverser is also two part. The time take to write a frame of data into memory and then read out in normalised order is:

n Fs clocks

Additionally the pipeline delay of the bit reverser is applicable.

$$t_z = 2 \text{ Fs clocks}$$

So the total delay in Fs clocks is:

$$t_w + t_x + 1.5n + t_y + n + t_z$$

So for example, a 10 stage (1024 point) FFT with FIFO and bit reverser and Fs of 160 MSPS.

Delay (excluding bit reverser) = 20 + 1 + (1.5x1024) + (16x10)
= 1717 clocks at Fs
= 10.73uS

Delay (including bit reverser) = 20 +1 + (1.5x1024) + (16x10) + 1024 + 2
= 2743 clocks at Fs
= 17.14uS

The FFT core is a continuous pipeline so once the initial delay of the pipeline filling up is past (described above), valid data is continuously available at the output. The concept of data blocks is only used in the description above to help describe the initial FFT pipeline delay at start up.

Core Delivery

The Core is delivered as an edif or ngo netlist with .ucf user constraints file. Alternatively, it can be supplied as a bit stream if the FFT target device is fully defined and only holds the FFT core.

VHDL test benches are provided along with a VHDL simulation model.

Supporting documentation including data sheets and user guides are included.

Design support services for implementing the core into your design as well as debugging support can also be provided.

As the FFT core is an optimised pipelined architecture, unlike resource sharing single butterfly architectures, its physical size is dependant on the size of FFT to be performed. Therefore the delivered core is of a set size as it is programmed with that specific n point FFT configuration. It may be possible to provide a core for say a 4096 point transform that can also select smaller transforms (e.g. 512, 256,128 etc). This provides better flexibility of operation but would not be an optimum use of silicon.

What is available now?

The 4K core as described above is available for immediate licensing. Other variants of the core in terms of FFT length, input data width, bit growth per stage, twiddle width and blockRAM / distributed memory split can be supplied under contract. RFEL will optimise the core based on exact customer requirements.

The bit width (word length) of the data or twiddles is not fixed and depending on the system requirements can be tailored throughout the stages of the FFT, to maintain the required dynamic range. Bit width is a factor that affects silicon area, so should be given appropriate consideration.

Accuracy and precision are system design considerations that affect data bit width and twiddle width. RFEL can offer system engineering advice to aid the selection of the optimal core configuration.

Weighted overlapped and add (WOLA) techniques such as the Polyphase-DFT are now available. Please refer to the datasheets for these products

Multi-radix architectures will also be available shortly. These designs will greatly increase the input data sample rates that are currently available for the radix-2 cores.

What does it cost?

The core is supplied under an Application Licence with pricing based on an up front payment and royalty per use. Pricing level is calculated on application and projected royalty revenue. A single price for unlimited use may also be considered.

What should you do now?

Please send a brief email to sales@rfel.com outlining your basic requirements. Alternatively, if you wish to be kept updated of our progress and alerted to new cores as they become available, please request to be put on our email update list via the address above.