

Four-Point FFT Processor

An Asynchronous/Synchronous Comparison

Radhika Balasubramanian
Anil Kumar Reddy Palaparathi
Arun Tejusve Raghunath Rajan
Trent Rolf

Department of Electrical and Computer Engineering, University of Utah

radhika.balasubramanian@gmail.com anilreddy.syntel@gmail.com
tejus2005@gmail.com trent@rolfed.com

Abstract

The Fast Fourier Transform (FFT) is an efficient algorithm to compute the Discrete Fourier Transform (DFT) and its inverse. It has several applications in the field of signal processing. The traditional butterfly FFT design requires needless computations and data storage which lead to unnecessary power consumption. The architecture used in this project simplifies and pipelines these computations, reducing power and enhancing speed. We implement both an asynchronous and a synchronous version of the 4-point FFT processor as a comparison of speed, power, and area. From this side-by-side comparison we decide which is a more efficient architecture for this application.

1. Introduction

The Fast Fourier Transform is derived from the Discrete Fourier Transform, which in its simplest mathematical form is defined as:

$$X_k = \sum_{n=0}^{N-1} x_n e^{-\frac{2\pi i}{N}kn} \quad k = 0, \dots, N-1$$

A straightforward hardware implementation of the DFT would result in at least N multiplications per output sample, storage of N coefficients, and inefficient data shifting and storage. The FFT was discovered to be a much more computationally friendly algorithm which recursively breaks down larger DFT's into smaller pieces that are easier to compute.[1] The particular model (loosely based on the Steven-Suter algorithm[2]) used in this project further enhances the efficiency of the FFT by pipelining computations, which inherently shares computing resources and reduces overall power consumption.

2. The FFT-4 Model

The block diagram of a 4-point FFT is shown (Figure 1). There are four inputs ($x[0]$ - $x[3]$) and four outputs ($X[0]$ - $X[3]$), each having a real and imaginary part.

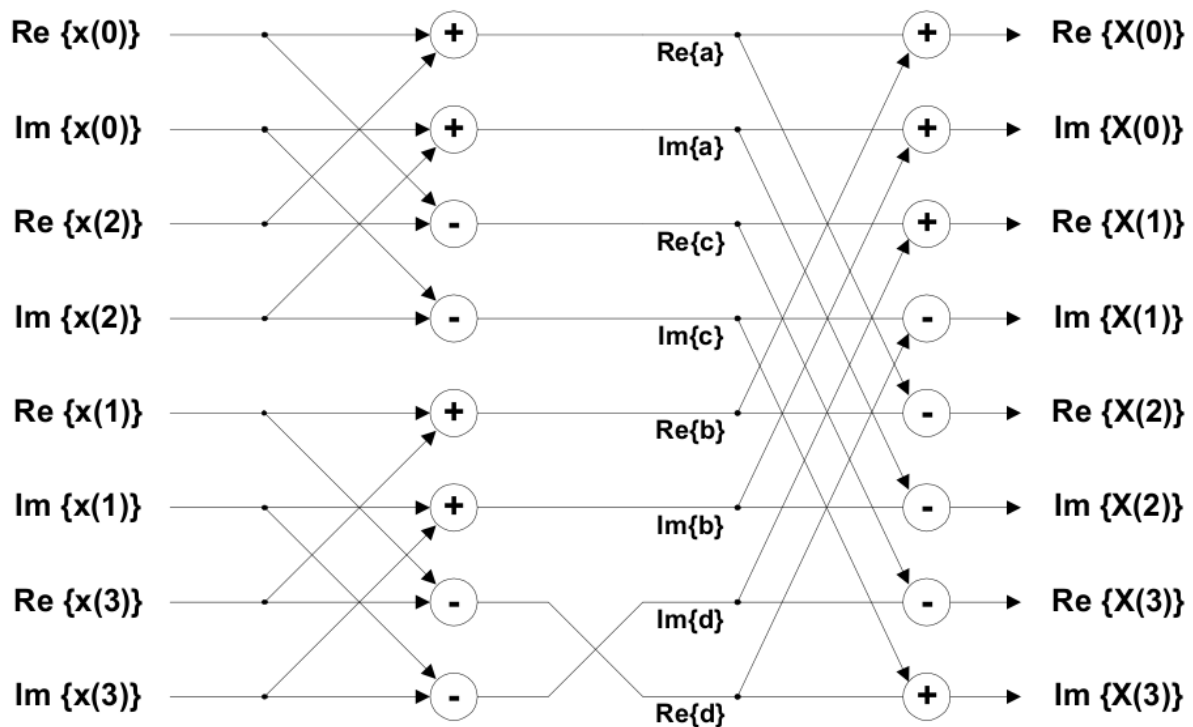


Figure 1: FFT-4 Data flow diagram. [3]

2.1. Flow of Data

Expansion of the matrix of 4-point FFT gives the following equations.

$$X(0) = x(0) + x(1) + x(2) + x(3)$$

$$X(1) = x(0) - jx(1) - x(2) + jx(3)$$

$$X(2) = x(0) - x(1) + x(2) - x(3)$$

$$X(3) = x(0) + jx(1) - x(2) - jx(3)$$

This can be separated into two stages of addition/subtraction.

In the first stage:

$$a = x(0) + x(2)$$

$$b = x(1) + x(3)$$

$$c = x(0) - x(2)$$

$$d = x(1) - x(3),$$

In the second stage:

$$\text{Re}\{X(0)\} = \text{Re}\{a\} + \text{Re}\{b\}$$

$$\text{Im}\{X(0)\} = \text{Im}\{a\} + \text{Im}\{b\}$$

$$\text{Re}\{X(1)\} = \text{Re}\{c\} + \text{Im}\{d\}$$

$$\text{Im}\{X(1)\} = \text{Im}\{c\} - \text{Re}\{d\}$$

$$\text{Re}\{X(2)\} = \text{Re}\{a\} - \text{Re}\{b\}$$

$$\text{Im}\{X(2)\} = \text{Im}\{a\} - \text{Im}\{b\}$$

$$\text{Re}\{X(3)\} = \text{Re}\{c\} - \text{Im}\{d\}$$

$$\text{Im}\{X(3)\} = \text{Im}\{c\} + \text{Re}\{d\}$$

This breakdown of the FFT-4 computation is used directly in our implementation in behavioral Verilog HDL.

2.2. Advantages of FFT-4

The simplicity of 4-point FFT lies in the fact that it requires only 16 individual add or subtract operations and no complex multiplications. By only using hardware resources for addition and subtraction, we have more freedom to implement power-efficient arithmetic that will share computing resources among pipeline stages.

2.3. A "Golden Model" Software Implementation

The golden model of the 4-point FFT has been implemented in C++ language. The golden model generates four random integer inputs (separate real and imaginary values) and computes the FFT using the arithmetic presented above.

The golden model is implemented to mimic the hardware in such a way that it waits for data dependencies to resolve before it continues. As a result, 10 clock cycles pass before valid data appear at the final output (corresponding to the first input sample). The golden model results are then used to compare against the synthesized circuit.

3. Synchronous Implementation

With the software model of the design complete, we moved on to the hardware implementation of the circuit. Our strategy was to first implement the design as a clocked circuit, then remove the clock and replace it with a handshaking protocol to create an asynchronous design.

3.1. Synchronous Architecture

At the heart of the design is the mathematical model of the FFT-4 with its efficient add/subtract-only structure as described above. Equally important, however, is the careful consideration required to convert the serial stream of input samples into segments of four parallel samples that the model depends on. Figure 2 shows a diagram of the top-level design.

Synchronous FFT-4 Architecture

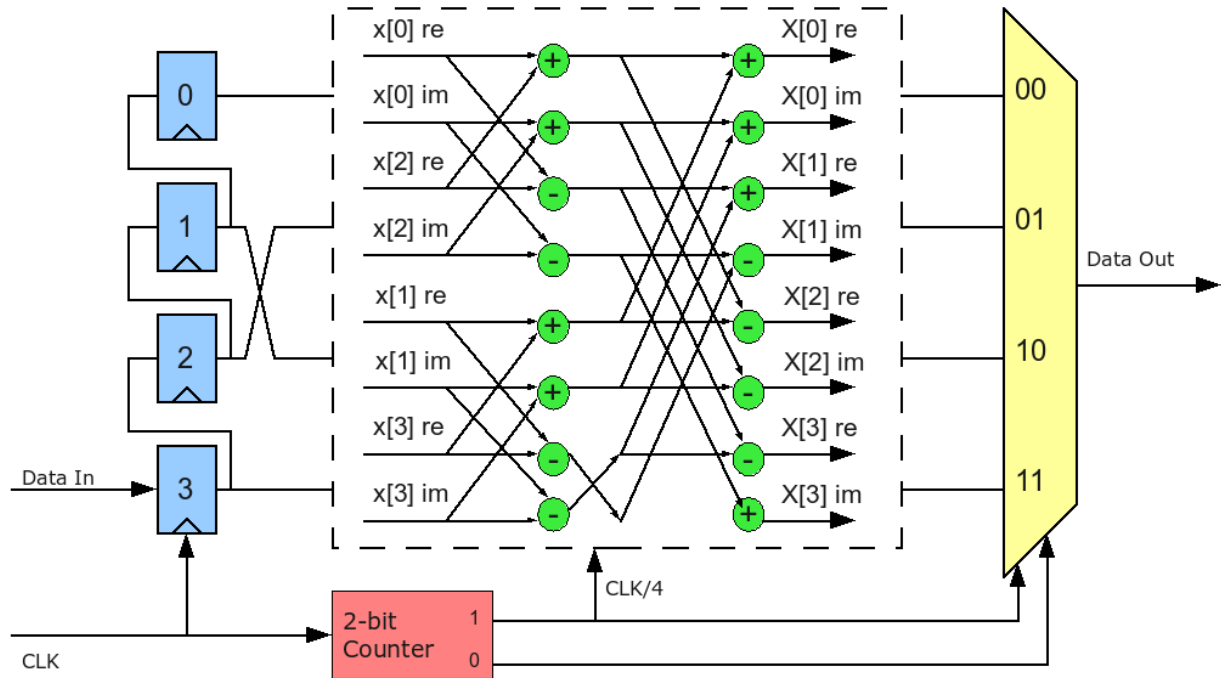


Figure 2: *Top-level synchronous design.*

The data enters the module one sample at a time (consisting of a real and imaginary part) and is latched into the first of four shift registers. Each value is a total of 32 bits wide (16 real, 16 imaginary). After three more clock cycles the input registers 0-3 contain the first four samples of data and the first actual FFT-4 computation can begin. A 2-bit counter acts as a clock divider providing a clock at 1/4 of the original frequency. The 1/4 frequency clock feeds the FFT-4 module.

Inside the FFT-4 module, the data bus expands to 20 bits (from 16) during the arithmetic stages to avoid computational overflow. Besides the adders, there are also buffer registers that exist to allow the synthesizer to re-time the circuit. The bus is truncated back to 16-bits at the final FFT-4 module stage.

At the top-level output stage, the parallelization operation is reversed using a multiplexer that selects the correct FFT result based on the 2-bit counter output. The data stream is serial once again with the correct Fourier transform output appearing at a fixed latency from the corresponding input.

3.2. Synchronous Simulation Results

ModelSim simulation tools were used to simulate the design for a comparison against the golden model. The synthesized Verilog design generated the correct outputs as compared to the golden model, however the latency between input and corresponding output did not match the model exactly. The latency for the golden model was 10 samples while the hardware design had a latency of 22 samples. This discrepancy is probably due to Design Compiler adding re-timing registers and other data buffers.

Here are a few waveforms that demonstrate the correct operation of the FFT-4 design:

Figure 3 shows the simulated input samples, with one new value appearing on each clock cycle.

Notice the top waveform is the real part, and the next waveform is the imaginary part.

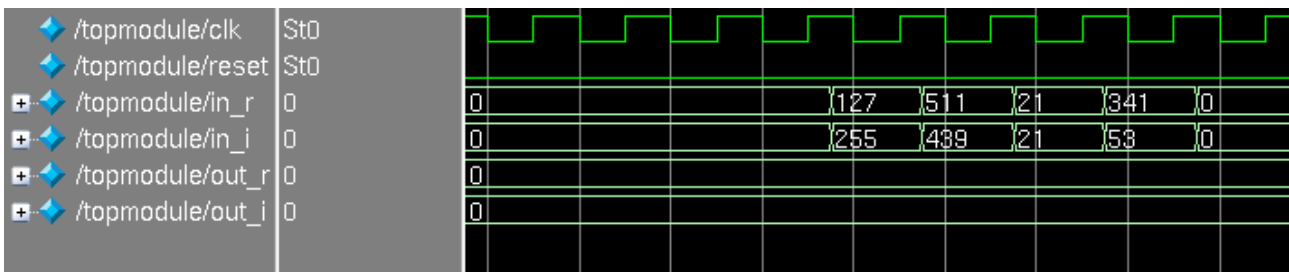


Figure 3: *Simulated input samples.*

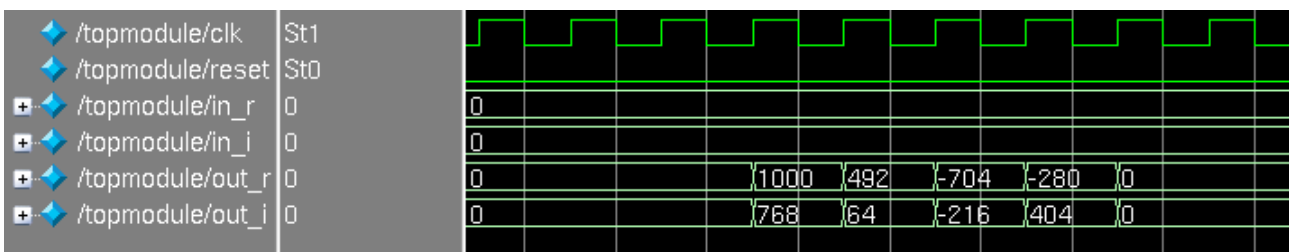


Figure 4: *Resulting output samples (arriving 22 clock cycles later).*

Figure 4 shows the resulting output samples which arrive at the output port 22 clock cycles after the corresponding input. By comparison, here is same data from the golden model:

```
input vector = [127+255i   511+439i   21+21i   341+53i]
output vector = [1000+768i  492+64i   -704-216i  -280+404i]
```

4. Asynchronous Implementation

The structure of the asynchronous implementation is very much similar to the synchronous implementation. The only difference is that there is no clock in the asynchronous version. The data transfer is instead based on a basic handshake protocol between different modules of the design. The asynchronous design consumes less power because the computations and data transfers occur only when request signal is asserted unlike in clocked version where the computations and data transfer occur at every rising edge of the clock. Also, a large design can be divided into simpler and smaller modules which can be connected easily because they use the same handshake protocol. Asynchronous design also allows for multiple clock domains where different modules can run at different frequencies and help to eliminate clock skew.

4.1. Asynchronous Architecture

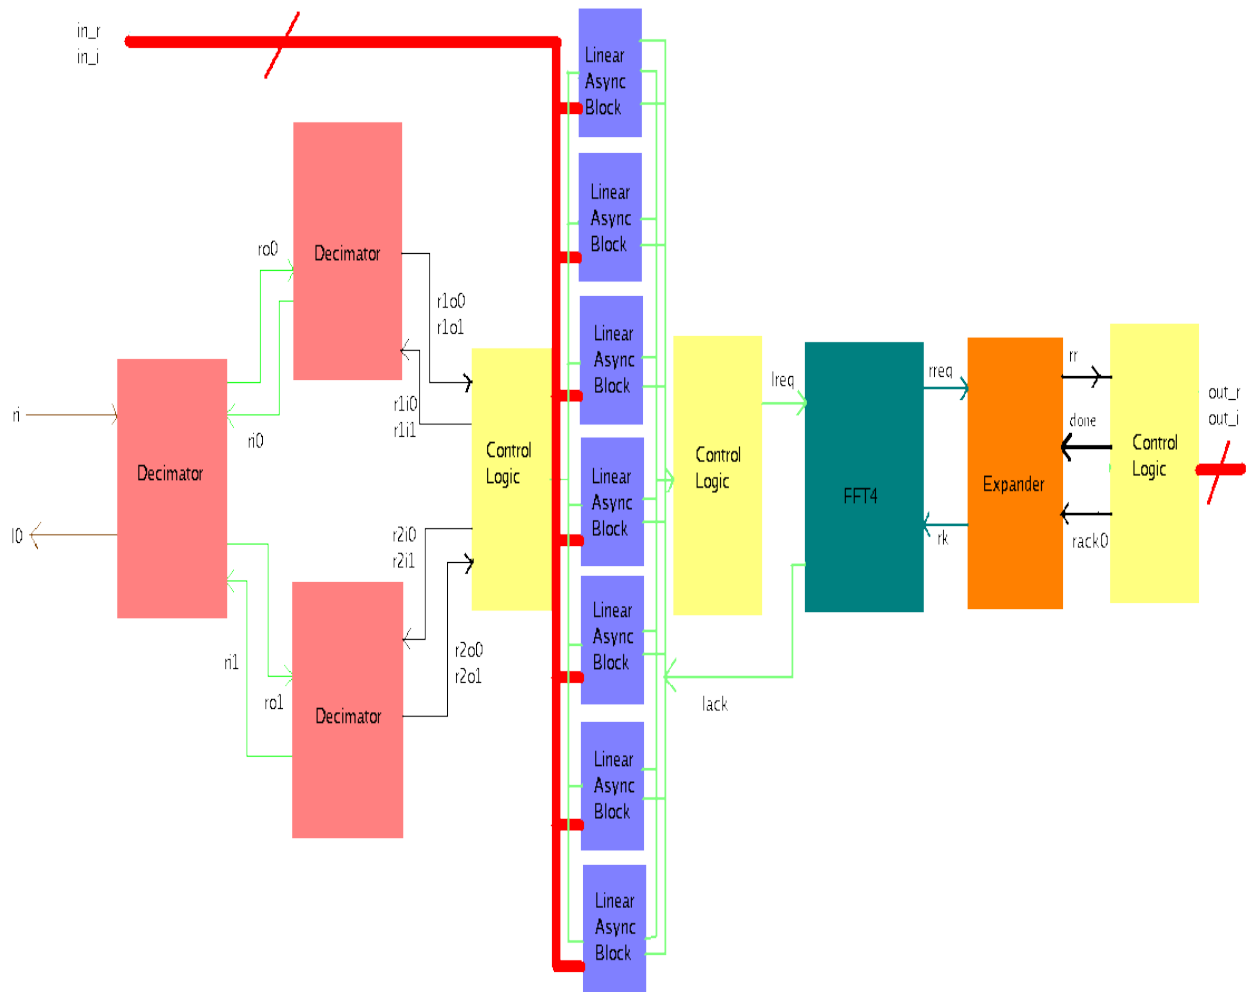


Figure 5: *Top-level asynchronous design.*

The top-level block diagram of the asynchronous version of the FFT-4 design is shown in the figure above. The design mainly consists of:

1. Decimator
2. Linear Asynchronous block
3. Logic Control
4. FFT-4 module, and
5. Expander.

The input data to the FFT-4 is a constant serial stream. The computation of the Fast Fourier Transform should not begin until receiving four input samples. Hence, the FFT-4 module should operate at a frequency four times slower than the input data frequency. The decimator halves the frequency of operation, so for our implementation two stages of the decimator are required. The logical implementation of decimator is as given below.

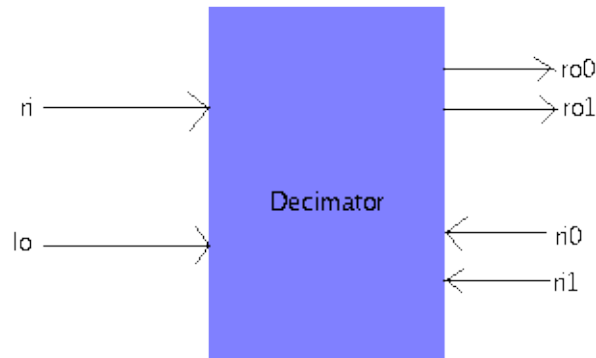


Figure 6: *Decimator.*

1	li+		ro0+
2	ri0+		ro0- lo+
3	li- ri0-		lo-
4	li+		ro1+
5	ri1+		ro1- lo+
0	li- ri1-		lo-

Each decimator (figure below) generates two request signals when li(lreq) goes high and asserts lo(lack) when the data on the next stage is latched (when the ri(rack) signal goes high). The first decimator sends the request signals (ro0,ro1) to the next stage which in turn controls the linear asynchronous blocks. The asynchronous blocks are required to latch the data. Since the data signals are complex valued, eight linear asynchronous blocks are required to latch the data. The linear asynchronous block latches the data as soon as it receives the request signal from the decimator blocks and sends back an acknowledgment to the decimator. The rreq signal from all the linear asynchronous blocks are ANDed together to generate a request signal to the FFT-4 module.

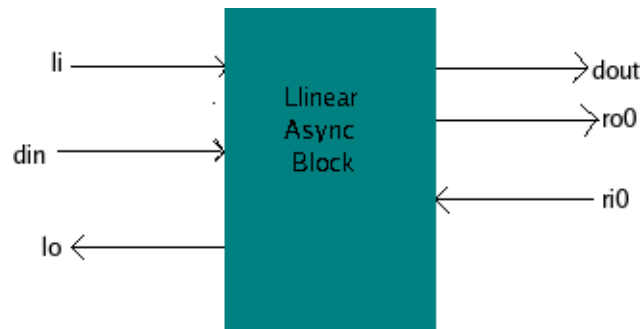


Figure 7: *Linear asynchronous module.*

The FFT-4 module (figure below) computes the Fourier transform of the input data samples. Since the request signals from all the linear asynchronous blocks are ANDed together, the request signal to the FFT-4 block is generated only after all four data signals are attained. The 4-point FFT computation involves two stages of additions and subtractions. The intermediate and final values are latched using the linear asynchronous blocks. The FFT4 module sends a request signal to the expander block which then provides the next set of input samples.

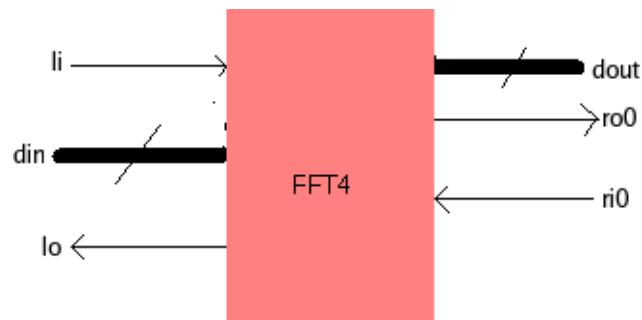


Figure 8: *FFT-4 module.*

The expander (figure below) functionality is similar to that of decimator in the reverse order. The expander sends the data out of the FFT-4 at the same rate as the input data. Although the input and output data run at the same frequency, the computation of the FFT is performed at a rate four times slower than the input signal reducing the overall power cost.

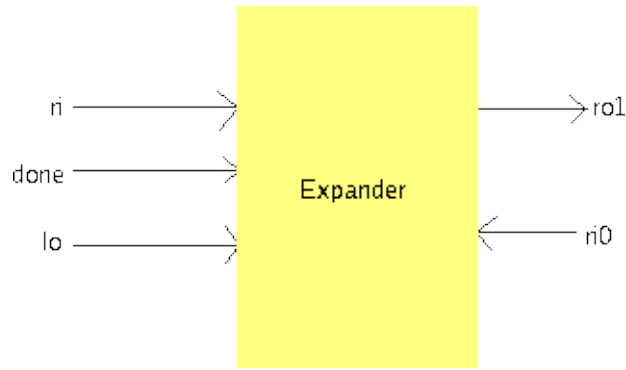


Figure 9: *Expander*.

4.2. Asynchronous Simulation Results

All our module designs were first defined using a state diagram. Then the "3D" asynchronous logic tool was used to generate the logic functions for each module. These generated logic functions were mapped to the UofUDigitalv1_1 library cells and were implemented in Verilog. Each module was tested and verified in the *ModelSim* simulation environment.

Here is an example of a top-level simulation:

◆ /topmodule/req0	St0												
◆ /topmodule/la	St0												
◆ /topmodule/rr	St1												
◆ /topmodule/rack0	St0												
▣ /topmodule/out_r	10	0							540		492		
▣ /topmodule/out_i	10	0							768		524		
▣ /topmodule/in_r	7	0	1	127	5	51	2	21	7	341	7	1	7
▣ /topmodule/in_i	7	0	1	255	5	439	2	21	7	53	7	1	7
◆ /topmodule/reset	St0												

Figure 10: *Input data $xin = [127+255i \ 51+439i \ 21+21i \ 341+53i]$*

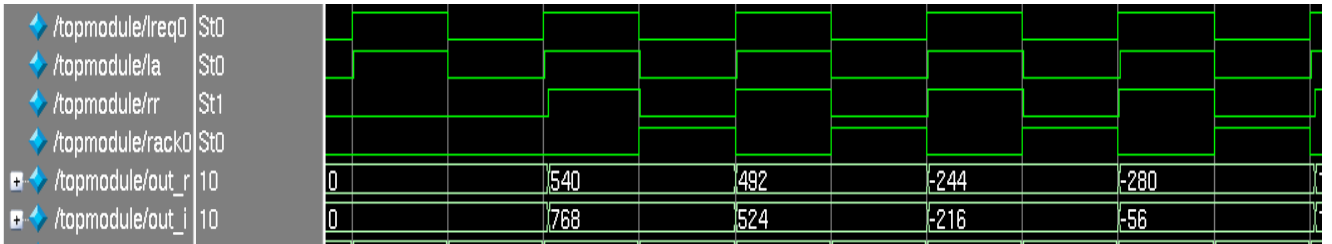


Figure 11: Output data $xout = [540+768i \ 492+524i \ -244-216i \ -280-56i]$

Here are some results from the individual module simulations:

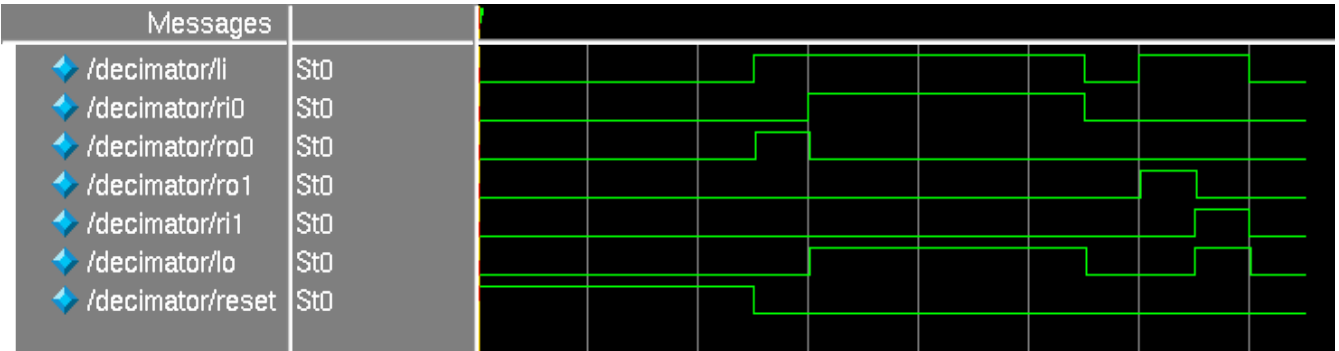


Figure 12: Decimator behavior.

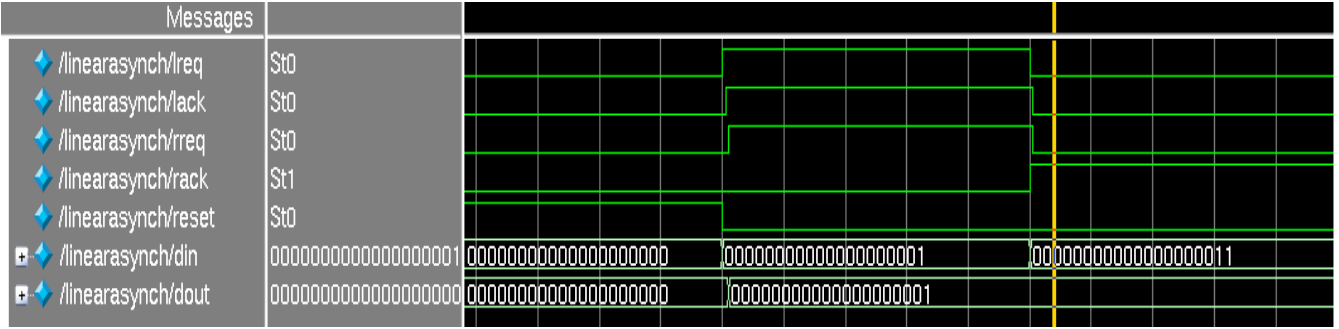


Figure 13: Linear Asynchronous Block behavior.

5. Conclusions

The comparison of the 4-point FFT processor in asynchronous and synchronous implementations has helped us understand the major differences between asynchronous and synchronous design, as well as the benefits of each. The difficulty of synthesizing the custom controllers in the asynchronous design and multiple clock frequencies in the synchronous design proved to be a barrier to our comparisons. However, through our understanding gained in this project we can say that the asynchronous implementation is better in terms of power consumption and overall efficiency. Computations are only performed when necessary, and issues arising from multiple clock domains are eliminated.

6. Acknowledgments

We would like to thank Dr. Ken Stevens for his help revising and fine-tuning our design, as well as sharing his expertise in asynchronous circuit design. We would also like to thank David James Barnhart for sharing his thesis research for our benefit.

7. References

- [1] Welch, T., C. Wright, M. Morrow. *Real-time Digital Signal Processing*. Taylor & Francis Publishers. 2006.
- [2] Hunt, B.W., K.S. Stevens, B.W. Hunter, D.S. Gelosh. "A Single Chip Low Power Asynchronous Implementation of an FFT Algorithm for Space Applications." Air Force Institute of Technology.
- [3] Barnhart, David James. "An Improved Asynchronous Implementation of a Fast Fourier Transform Architecture for Space Applications." 2001.