



CASPER Workshop

Tutorial 3: Wideband Spectrometer

Dev. By : Jason Manley and Danny Price

Doc. By : Irappa M. Halagali `gmrt/ncra/tifr`

Expected completion time: 2hrs

Contents :

1. The Hardware and software required for this tutorial.
2. Introduction
3. Spectrometer basics
4. Setup

PART ONE

5. Simulink design overview
6. Detailed blockumentation

PART TWO

7. Hardware configuration
8. The python script "File3d"
9. iPython walkthrough
10. The python script "File3d" notes
11. Conclusion

1. The Hardware and software required for this tutorial.

1. PC : Dell Intel(R) Core(TM) i3 CPU 530 @ 2.93GHz width 64 bit & 4GB RAM
 2. OS : Linux 2.6.35-30-generic #54-Ubuntu 10.10 SMP x86_64 GNU/Linux
 3. Matlab : 2008a
 4. Xilinx : version 11.5
 5. Casper : gits_100511
 6. corr packages: corr-0.6.5
 7. Python : version 2.6
 8. minicom : version 2.4 (compiled on Jun 3 2010)
 9. Wireshark : Running on Linux 2.6.35-30-generic, with libpcap version 1.1.1, GnuTLS 2.8.6, Gcrypt 1.4.5.
 10. TCPdump : tcpdump-4.1.1
 11. GULP : version 2.0, January 2004
 12. iperf : version 2.0.4 (7 Apr 2008) pthreads
-

13. jperf : version 2.0.2

14. ROACH unit with iADC card. : version 1.0 Rev 3 2009, uboot : [uboot-2010-07-15-r3231-dram](#) , Linux Kernel Image : [ulmage-jiffy-20091110](#) and iADC : BEE2 DUAL 1 GHz ADC BOARD version 1.1

15. Signal generator to feed clock of 800MHz , 0dbm to ROACH unit through iADC's `clk_i` input.

2 Introduction

A spectrometer is something that takes a signal in the *time* domain and converts it to the *frequency* domain. In digital systems, this is generally achieved by utilising the FFT (Fast Fourier Transform) algorithm. However, with a little bit more effort, the signal to noise performance can be increased greatly by using a Polyphase Filter Bank (PFB) based approach.

When designing a spectrometer for astronomical applications, it's important to consider the science case behind it. For example, pulsar timing searches will need a spectrometer which can dump spectra on short timescales, so the rate of change of the spectra can be observed. In contrast, a deep field HI survey will accumulate multiple spectra to increase the signal to noise ratio. It's also important to note that "bigger isn't always better"; the higher your spectral and time resolution are, the more data your computer (and scientist on the other end) will have to deal with. For now, let's skip the science case and familiarize ourselves with an example spectrometer.

In this tutorial, we will build a 400MHz bandwidth, 2048 channel, PFB based spectrometer on the ROACH. You'll need to have done Tutorial 1, 2 and the iADC tutorial¹. You should also have installed python, iPython, corr, aipy, numpy and pylab.

As far as hardware goes, you'll need:

- a ROACH board;
- an iADC, which should be connected to `ZDOK0` on the ROACH; and
- a clock source, such as a signal generator, which should be connected to `clk_i` on the iADC.

You'll need to be familiar with the basic concepts of sampling, have a solid understanding of what a Fourier transform is, and have a vague idea of what a FFT is. A good reference is Smith's free online DSP guide (at <http://www.dspguide.com/>); in particular, have a read of chapters 3, 8 and 12.

3 Spectrometer Basics

When designing a spectrometer there are a few main parameters of note:

- *Bandwidth*: The width of your frequency spectrum, in Hz. This depends on the sampling rate; for complex sampled data this is equivalent to:

$$BW = \text{sampling rate} = \frac{1}{\text{sampling period}}$$

In contrast, for real or Nyquist sampled data the rate is half this:

¹ While the iADC tutorial is for an iBOB, you can follow the iADC tutorial on the ROACH by simply changing the XSG core config from 'iBOB' to 'ROACH'. <http://casper.berkeley.edu/wiki/Tutorials#iBOB>

$$BW = \frac{\text{sampling rate}}{2} = \frac{1}{2 \times \text{sampling period}},$$

as two samples are required to reconstruct a given waveform².

- *Frequency resolution:* The frequency resolution of a spectrometer, Δf , is given by

$$\Delta f = \frac{BW}{\text{no. channels}},$$

and is the width of each frequency bin. Correspondingly, Δf is a measure of how precise you can measure a frequency.

- *Time resolution:* Time resolution is simply the spectral dump rate of your instrument. We generally accumulate multiple spectra to average out noise; the more accumulations we do, the lower the time resolution. For looking at short timescale events, such as pulsar bursts, higher time resolution is necessary; conversely, if we want to look at a weak HI signal, a long accumulation time is required, so time resolution is less important.

4 Setup

The lab at the workshop is pre-configured with the CASPER libraries, Matlab and Xilinx tools. Please refer the file "[LOCATIONSandFILES.pdf](#)" in the [home area](#) or [LOCATIONSandFILES slides displayed](#), for the locations/directories and files information required in the tutorial. Note : The Date and Time portion of the BOF file name will be different! It depends upon when (Date & Time) you compile your model file !

Note : All the following cable connections and entries in the /etc/* files of the workshop PCs are already done.

1. Connect the Serial port cable between the ROACH board's P2 connector and serial port of the PC (on which minicom program exists).
2. Connect the Ethernet cable to J25 port of the ROACH board from the PCs eth1 port. /etc/ethers file should have mac address and corresponding ip address. In the /etc/network/interfaces file , eth1 should be configured. And in the file /etc/hosts , ip address and corresponding roach board(host) name entry to be done.
3. Feed the clock of 800MHz , 0 dbm to the clk_i input of the iADC card from the Signal generator
4. Either copy the mdl file "[File3a](#)" from the the area "[Location3](#)" or follow the steps given below to create the mdl file similar to the file [Location3/File3a](#) after creating your own directory at "[Location23](#)" to save and compile your model file or the bof file "[File3b](#)" is kept in the area "[Location2](#)" to directly program into the FPGA and look at the results.
5. Start the matlab :
\$ cd [Location1](#)
[Location1](#)\$./File0 &

² For a introduction to sampling, read <http://www.dspguide.com/ch3.htm>

PART ONE

Simulink / CASPER Toolflow

5 Simulink Design Overview

If you're reading this, then you've already managed to find all the tutorial files. Jason has gone to great effort to create an easy to follow simulink model that compiles and works. By now, I presume you can open the model file and have a vague idea of what's happening.

The best way to understand fully is to follow the arrows, go through what each block is doing and make sure you know why each step is done. To help you through, there's some "blockumentation" in the appendix, which should (hopefully) answer all questions you may have. A brief rundown before you get down and dirty:

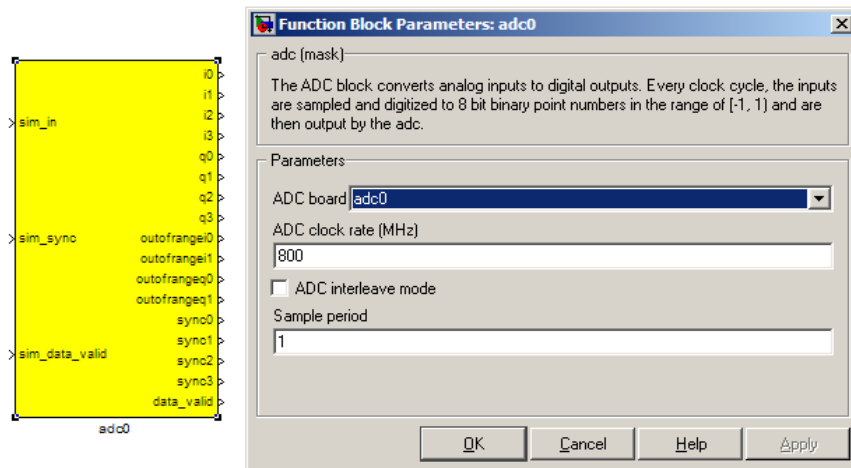
- The all important **Xilinx token** is placed to allow System Generator to be called to compile the design.
- In the **MSSGE block**, the hardware type is set to 'ROACH:sx95t' and clock rate is specified as 200MHz.
- The input signal is digitised by the **ADC**, resulting in four parallel time samples of 8 bits each clock cycle. The ADC runs at 800MHz, which gives a 400MHz nyquist sampled spectrum. The output range is a signed number in the range -1 to +1 (ie 7 bits after the decimal point). This is expressed as `fix_8_7`.
- The four parallel time samples pass through the **pfb_fir_real** and **fft_wideband_real** blocks, which together constitute a polyphase filter bank. We've selected $2^{12}=4096$ points, so we'll have a $2^{11}=2048$ channel filter bank.
- You may notice Xilinx delay blocks dotted all over the design. It's common practice to add these into the design as it makes it easier to fit the design into the logic of the FPGA. It consumes more resources, but eases signal timing-induced placement restrictions.
- The real and imaginary (sine and cosine value) components of the FFT are plugged into **power** blocks, to convert from complex values to real power values by squaring. They are also scaled by a gain factor before being quantised...
- These power values are then requantized by the **quant0** block, from 36.34 bits to 6.5 bits, in preparation for accumulation. This is done to limit bit growth.
- The requantized signals then enter the vector accumulators, **vacc0** and **vacc1**, which are **simple_bram_vacc** 32 bit vector accumulators. Accumulation length is controlled by the **acc_cntrl** block.
- The accumulated signal is then fed into software registers, **odd** and **even**. Without further ado, open up the model file and start clicking on things, referring the blockumentation as you go.

6 DETAILED BLOCKUMENTATION

ADC

<http://casper.berkeley.edu/wiki/Adc>

The first step to creating a frequency spectrum is to digitize the signal. This is done with an ADC – an Analogue to



Digital Converter. In Simulink, the ADC daughter board is represented by a yellow block. Work through the “iADC tutorial” if you're not familiar with the iADC card.

The ADC block converts analog inputs to digital outputs. Every clock cycle, the inputs are sampled and digitized to 8 bit binary point numbers in the range of -1 to 1 and are then output by the ADC. This is achieved through the use of two's-compliment representation with the binary point placed after the seven least significant bits. This means we can represent numbers from -128/128 through to 127/128 including the number 0. Simulink represents such numbers with a `fix_8_7` moniker.

ADCs often internally bias themselves to halfway between 0 and -1 (ie the center of the range of representable values). This means that you'd typically see the output of an ADC toggling between zero and -1 when there's no input. It also means that unless otherwise calibrated, an ADC will have a negative DC offset.

The ADC has to be clocked to four times that of the FPGA clock. In this design the ADC is clocked to 800MHz, so the ROACH will be clocked to 200MHz³. This gives us a bandwidth of 400MHz, as Nyquist sampling requires two samples (or more) each second.

INPUTS

<code>sim_in</code>	Input for simulated data. It's useful to connect up a simulink source, such as “band-limited white noise” or a sine wave.
<code>sim_sync</code>	Simulated sync pulse input. In this design, we've connected up a constant with value '1'.
<code>sim_data_valid</code>	Can be set to either 0 (not valid) or 1 (valid).

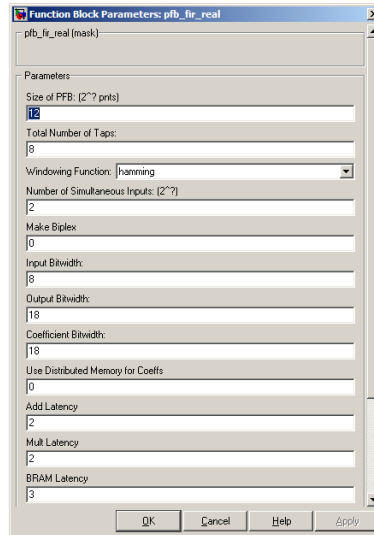

OUTPUTS

The ADC outputs two main signals: *i* and *q*, which correspond to the coaxial inputs of the ADC board. In this tutorial, we'll only be using input *i*. As the ADC runs at 4x the FPGA rate, there are four parallel time sampled outputs: *i0*, *i1*, *i2* and *i3*. As mentioned before, these outputs are 8.7 bit.

³ In the XSG core config block, we have selected `adc0_clk` as our clock, which is provided by the ADC daughter card plugged into connector ZDOK0.

PFB_FIR_REAL

http://casper.berkeley.edu/wiki/Pfb_fir_real



There are two main blocks required for a polyphase filter bank. The first is the **pfb_fir_real** block, which divides the signal into parallel 'taps' then applies finite impulse response filters (FIR). The output of this block is still a time-domain signal. When combined with the **FFT_wideband_real** block, this constitutes a polyphase filterbank.

INPUTS/OUTPUTS

Port	Data Type	Description
<i>sync</i>	<i>bool</i>	A sync pulse should be connected here (see iADC tutorial).
<i>pol1_in1</i>	<i>inherited</i>	As the ADC has four parallel time sampled outputs: <i>i0</i> , <i>i1</i> , <i>i2</i> and <i>i3</i> , we need four parallel inputs for this PFB implementation.
<i>pol1_in2</i>		
<i>pol1_in3</i>		
<i>pol1_in4</i>		

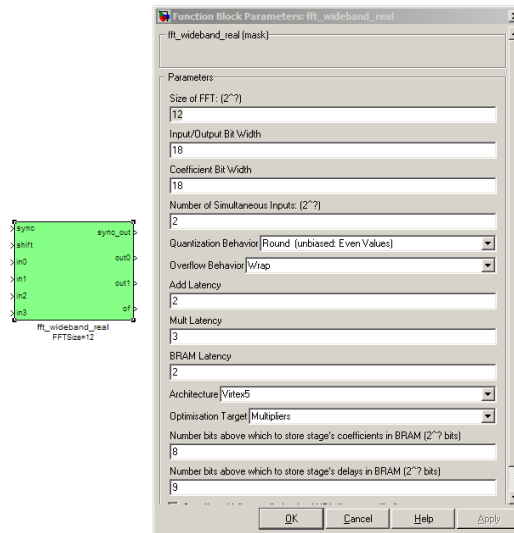
PARAMETERS

Size of PFB		How many points the FFT will have. The number of frequency channels will be half this. We've selected $2^{12}=4096$ points, so we'll have a $2^{11}=2048$ channel filter bank.
Number of taps		The number of taps in the PFB FIR filter. Each tap uses 2 real multiplier cores and requires buffering the real and imaginary streams for $2^{PFBSize}$ samples. Generally, more taps means less inter-channel spectral leakage, but more logic is used. There are diminishing returns after about 8 taps or so.
Windowing function		Which windowing function to use (this allows trading passband ripple for steepness of rolloff, etc). Hamming is the default and best for most purposes.
Number of Simultaneous (2^n) Inputs	of	The number of parallel time samples which are presented to the FFT core each clock. The number of output ports are set to this same value. We have four inputs from the ADC, so set this to 4.
Make biphase		0 (not making it biphase) is default. Double up the inputs to match with a biphase FFT.
Input bitwidth.		The number of bits in each real and imaginary sample input to the PFB. The ADC outputs 8.7 bit data, so the input bitwidth should be set to 8 in our design.

Output bitwidth	The number of bits in each real and imaginary sample output from the PFB. This should match the bit width in the FFT that follows. 18 bits is recommended for the ROACH (18x25 multipliers) and iBOB/BEE2 (18x18 multipliers).
Coefficient bitwidth	The number of bits in each coefficient. This is usually chosen to be less than or equal to the input bit width.
Use dist mem for coefficients	Store the FIR coefficients in distributed memory (if = 1). Otherwise, BRAMs are used to hold the coefficients. 0 (not using distributed memory) is default
Latency	There's normally no reason to change this unless you're having troubles fitting the design into the fabric.
Quantization Behavior	Specifies the rounding behavior used at the end of each butterfly computation to return to the number of bits specified above. Rounding is strongly suggested to avoid artifacts.
Bin Width Scaling	PFBs give enhanced control over the width of frequency channels. By adjusting this parameter, you can scale bins to be wider (for values > 1) or narrower (for values < 1).

FFT_WIDEBAND_REAL

http://casper.berkeley.edu/wiki/Fft_wideband_real



The **FFT_wideband_real** block is the most important part of the design to understand. The cool green of the FFT block hides the complex and confusing FFT butterfly bplex algorithms that are under the hood. You do need to have a working knowledge of it though, so I recommend reading Chapter 8 and Chapter 12 of Smith's free online DSP guide (at <http://www.dspguide.com/>).

Parts of the documentation below are taken from the documentation by Aaron Parsons and Andrew Martens on the CASPER wiki (at <http://casper.berkeley.edu/wik/>).

INPUTS/OUTPUTS

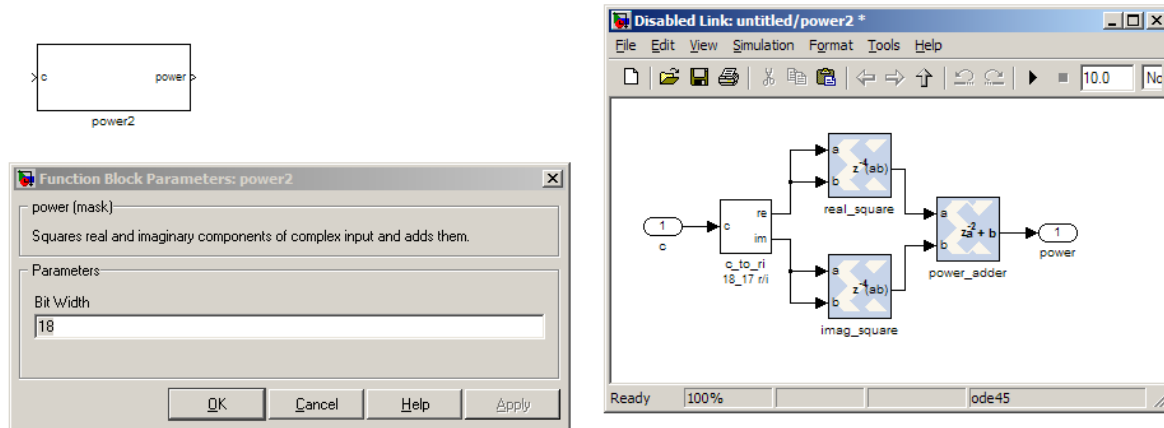
sync	Like many of the blocks, the FFT needs a heartbeat to keep it sync'd
Shift	Sets the shifting schedule through the FFT. Bit 0 specifies the behavior of stage 0, bit 1 of stage 1, and so on. If a stage is set to shift (with bit = 1), then every sample is divided by 2 at the output of that stage. In this design, we've set Shift to $2^{13} - 1 - 1$, which will shift the data by 1 on every stage to prevent overflows.
in0 in1 in2 in3	Four inputs for the parallel data streams coming from the ADC, through the pfb_fir_real filter block, and into here. Just connect them up.
out0 out1	The FFT produces two signals, the real part (out0 , cosine wave values) and the imaginary part (out1 , sine wave values). Following the lines you'll see that these two inputs end up in an "odd" and "even" software register. This is then interleaved in the <i>spectrometer.py</i> script to form a complete spectrum. Data is output in normal frequency order, meaning that channel 0 (corresponding to DC) is output first, followed by channel 1, on up to channel $2^N - 1 - 1$.

PARAMETERS

Size of FFT	How many points the FFT will have. The number of channels will be half this. We've selected $2^{12}=4096$ points, so we'll have a $2^{11}=2048$ channel filter bank. This should match up with the <code>pfb_fir</code> block.
Input/output bitwidth	The number of bits in each real and imaginary sample as they are carried through the FFT. Each FFT stage will round numbers back down to this number of bits after performing a butterfly computation. This has to match what the <code>pfb_fir</code> is throwing out. The default is 18 so this shouldn't need to be changed.
Coefficient bitwidth	The amount of bits for each coefficient. 18 is default.
Number of simultaneous inputs	The number of parallel time samples which are presented to the FFT core each clock. We have $2^2=4$ parallel data streams, so this should be set to 2.
Quantization Behavior	Specifies the rounding behavior used at the end of each butterfly computation to return to the number of bits specified above. Here we'll just use Round (unbiased: Even Values).
Overflow Behavior	Indicates the behavior of the FFT core when the value of a sample exceeds what can be expressed in the specified bit width. Here we're going to use Wrap as Saturate will not make overflow corruption better behaved.
Add Latency	Latency through adders in the FFT. Set this to 2..
Mult Latency	Latency through multipliers in the FFT. Set this to 3.
BRAM Latency	Latency through BRAM in the FFT. Set this to 2.
Convert	Latency through blocks used to reduce bit widths after twiddle and butterfly stages. Set this to 1.
Architecture	Set to Virtex5, the architecture of the FPGA on the ROACH. This changes some of the internal logic to better optimise for the DSP slices.
Use less	This affects if complex multiplications use less multipliers or adders/logic. For the complex multipliers in the FFT, you can use 4 multipliers and 2 adders, or 3 multipliers and a bunch of adders. So you can trade-off DSP slices for logic or vice-versa. Set this to Multipliers.
Number of bits above which to store stage's coeff's in BRAM	Determines the threshold at which the twiddle coefficients in a stage are stored in BRAM. Below this threshold distributed RAM is used. By changing this, you can bias your design to use more BRAM or more logic. We're going to set this to 8.
Number of bits above which to store stage's delay's in BRAM	Determines the threshold at which the twiddle coefficients in a stage are stored in BRAM. Below this threshold distributed RAM is used. Set this to 9.
Multiplier Specification	Determines how multipliers are implemented in the twiddle function at each stage. Using behavioral HDL allows adders following the multiplier to be folded into the DSP48Es in Virtex5 architectures. Other options choose multiplier cores which allows quicker compile time. If selected, you can enter an array of values allowing exact specification of how multipliers are implemented at each stage. Leave this unchecked.
Use DSP48's for adders	The butterfly operation at each stage consists of two adders and two subtracters that can be implemented using DSP48 units instead of logic. Leave this unchecked.

POWER

<http://casper.berkeley.edu/wiki/Power>



The **power** block computes the power of a complex number. The power block typically has a latency of 5 and will compute the power of its input by taking the sum of the squares of its real and imaginary components. The power block is written by Aaron Parsons and online documentation is by Ben Blackman.

In our design, there are two power blocks, which compute the power of the odd and even outputs of the FFT. The output of the block is 36.34 bits; the next stage of the design re-quantizes this down to a lower bitrate.

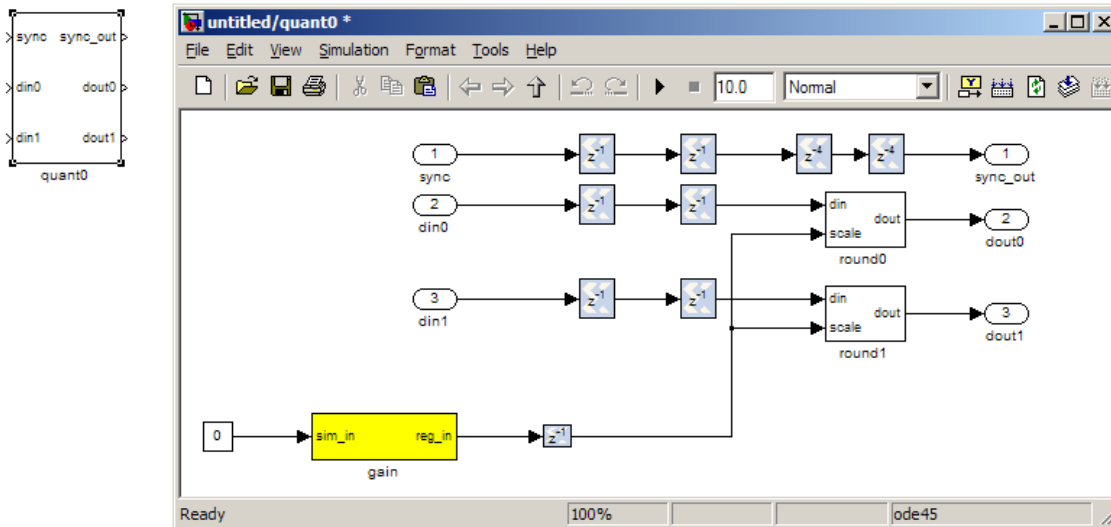
PARAMETERS

Parameter	Variable	Description
Bit Width	BitWidth	The number of bits in its input.

INPUTS/OUTPUTS

Port	Dir	Data Type	Description
c	IN	$2 * \text{BitWidth}$ Fixed point	A complex number whose higher BitWidth bits are its real part and lower BitWidth bits are its imaginary part.
power	OUT	$UFix_{(2 * \text{BitWidth}) - (2 * \text{BitWidth} - 1)}$	The computed power of the input complex number.

QUANT



The **quant0** was written by Jason Manley for this tutorial and is not part of the CASPER blockset. The block re-quantizes from 36.34 bits to 6.5 unsigned bits, in preparation for accumulation by the 32 bit **bram_vacc** block. This block also adds gain control, via a software register. The *spectrometer.py* script sets this gain control. You would not need to re-quantize if you used a larger vacc block, such as the 64bit one, but it's illustrative to see a simple example of re-quantization, so it's in the design anyway.

Note that the **sync_out** port is connected to a block, **acc_ctrl**, which provides accumulation control.

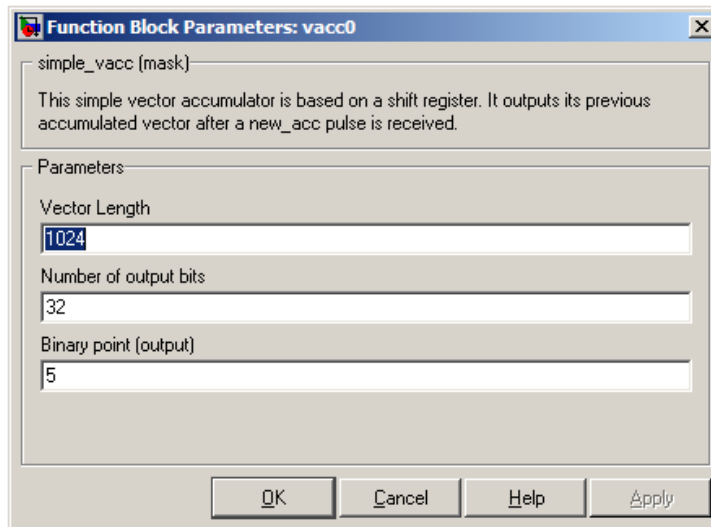

PARAMETERS

None.

INPUTS/OUTPUTS

Port	Description
Sync	Input/output for the sync heartbeat pulse.
din0 din1	Data inputs – odd is connected to <i>din0</i> and even is connected to <i>din1</i> . In our design, data in is 36.34 bits.
dout0 dout1	Data outputs. In this design, the quant0 block requantizes from the 36.34 input to 6.5 bits, so the output on both of these ports is 6.5 unsigned bits.

SIMPLE_BRAM_VACC



The **simple_bram_vacc** block is used in this design for vector accumulation. Vector growth is approximately 28 bits each second, so if you wanted a really long accumulation (say a few hours), you'd have to use a block such as the **qdr_vacc** or **dram_vacc**. As the name suggests, the **simple_bram_vacc** is simpler so it is fine for this demo spectrometer.

The FFT block outputs 1024 cosine values (odd) and 1024 sine values, making 2048 values in total. We have two of these bram vacc's in the design, one for the odd and one for the even frequency channels. The vector length is thus set to 1024 on both.

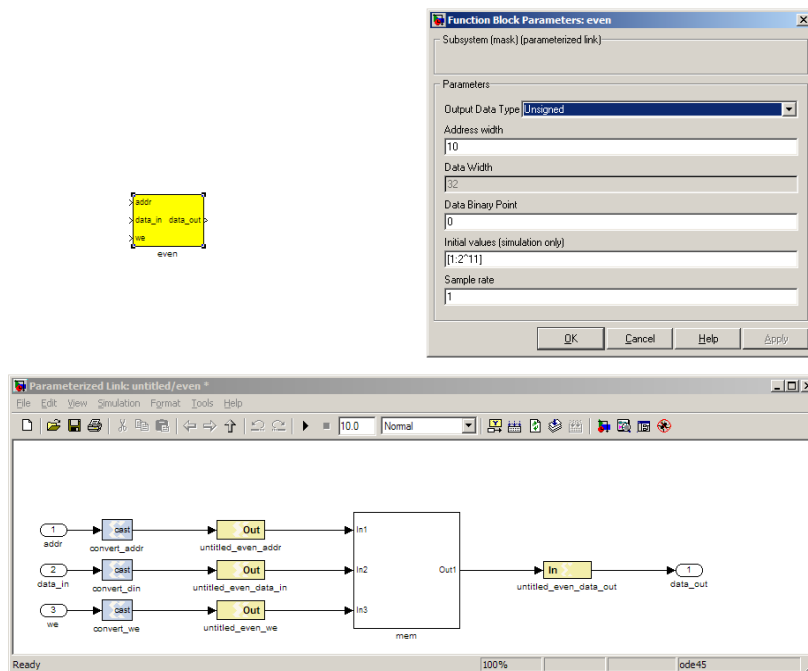
PARAMETERS

Parameter	Description
Vector length	The length of the input/output vector. The FFT block produces two streams of 1024 length (odd and even values), so we set this to 1024.
no. output bits	As there is bit growth due to accumulation, we need to set this higher than the input bits. The input is 6.5 from the quant0 block, we have set this to 32 bits. Note: We could set this to 64 bits and skip the quant block.

INPUTS/OUTPUTS

Port	Description
<i>new_acc</i>	A boolean pulse should be sent to this port to signal a new accumulation. We can't directly use the sync pulse, otherwise this would reset after each spectrum. So, Jason has connected this to acc_cntrl , a block which allows us to set the accumulation period.
<i>din/dout</i>	Data input and output. The output depends on the <i>no. output bits</i> parameter.
<i>Valid</i>	The output of this block will only be valid when it has finished accumulating (signalled by a boolean pulse sent to <i>new_acc</i>). This will output a boolean 1 while the vector is being output, and 0 otherwise.

EVEN AND ODD BRAMS



The final blocks, **odd** and **even** are shared BRAMs, which we will read out the values of using the *spectrometer.py* script.

PARAMETERS

Parameter	Description
<i>Output data type</i>	Unsigned
<i>Address width</i>	$2^{\text{Address width}}$ is the number of 32 bit words of the implemented BRAM. There is no theoretical maximum for the Virtex 5, but there will be significant timing issues at bitwidths of 13. QDR or DRAM can be used for larger address spaces. Set this value to 11 for our design.
<i>Data binary point</i>	The binary point should be set to zero. The data going to the processor will be converted to a value with this binary point and the output data type.
<i>Initial values</i>	This is a test vector for simulation only. We can leave it as is.
<i>Sample rate</i>	Set this to 1.

INPUTS/OUTPUTS

Port	Description
<i>Addr</i>	Address to be written to with the value of <i>data_in</i> , on that clock, if write enable is high.
<i>data_in</i>	The data input
<i>We</i>	Write enable port
<i>data_out</i>	Writing the data to a register. This is simply terminated in the design, as the data has finally reached its final form and destination.

CONTROL REGISTERS

There are a few control registers, led blinkers and snap block dotted around the design too:

1. **cnt_rst**: Counter reset control. Pulse this high to reset all counters back to zero.
2. **acc_len**: Sets the accumulation length. Have a look in *spectrometer.py* for usage.
3. **sync_cnt**: Sync pulse counter. Counts the number of sync pulses issued. Can be used to figure out board uptime and confirm that your design is being clocked correctly.
4. **acc_cnt**: Accumulation counter. Keeps track of how many accumulations have been done.
5. **led0_sync**: In the spirit of Kanye (feat. Dwele), we have some flashing lights. I would argue that ours are *the best flashing lights of all time*, but what do I know⁴? Back on topic: the **led0_sync** light flashes each time a sync pulse is generated. It lets you know your ROACH is alive.
6. **led1_new_acc**: This lights up led1 each time a new accumulation is triggered.
7. **led2_acc_clip**: This lights up led2 whenever clipping is detected.

There are also some *snap* blocks, which capture data from the FPGA fabric and makes it accessible to the Power PC. This tutorial doesn't go into these blocks (in its current revision, at least), but if you have the inclination, have a look at their wiki entry at:

<http://casper.berkeley.edu/wiki/Snap>

In this design, the snap blocks are placed such that they can give useful debugging information. You can probe these through KATCP, as done in tutorial one, if interested.

If you've made it to here, congratulations, go and get yourself a cup of tea and a biscuit, then come back for part two, which explains the second part of the tutorial – actually getting the spectrometer running, and having a look at some spectra.

⁴ Hope you appreciate the joke more than Taylor Swift

PART TWO

Configuration and Control

7 Hardware configuration

The tutorial comes with a pre-compiled bof file “File3b” at “Location4” , which is generated from the tutorial you just went through.

You don't need to telnet in to the ROACH; all communication and configuration will be done by the python control script.

Copy the bof file to be programmed which is compiled by you , in the directory “Location2” after changing the permissions of the file.

eg. for the bof file

File3b in the area Location7/Location10/

```
$ chmod a+x Location7/Location10/File3b
```

```
$ cp Location7/Location10/File3b Location2/
```


The tutorial comes with a python file called “File3d”. To use this, you need to have installed a few python libraries. If you haven't already, go through the instructions on

<http://casper.berkeley.edu/wiki/Corr>

I'd recommend installing **all** the packages, and documenting any trouble you have on the discussion page of the wiki. Also, iPython is used later on in this tutorial.

Next, you need to set up your ROACH. Switch it on, making sure that:

- You have your ADC in *ZDOK0*, which is the one nearest to the power supply.
- You have your clock source connected to *clk_i* on the ADC, which is the second on the right. It should be generating an 800MHz sine wave with 0dBm power.



If set up correctly, it should look like the photo on the right (sans the 10GbE cable).

8 The python script “File3d”

Once you've got that done, it's time to run the script. First, check that you've connected the ADC to *ZDOK0*, and that the clock source is connected to *clk_i* of the ADC.

Now, if you're in linux, browse to where the “File3d” file is in a terminal and at the prompt type

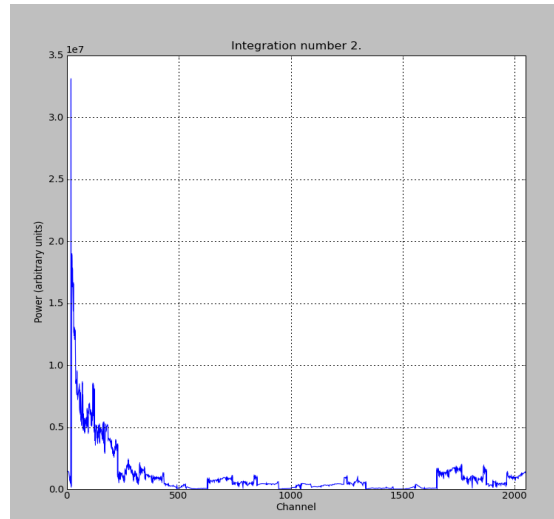
```
./File3d <roach IP or hostname> <boffile name>
```

replacing '<roach IP or hostname>' with the IP address of your ROACH and boffile name with your boffile “File3b”.

eg. \$ Location5/File3d roach030172 File3b

#Enter the corresponding Location/File names and roach name/IP.

You should see a spectrum like this:



In the plot, there should be a fixed DC offset spike; and if you're putting in a tone, you should also see a spike at the correct input frequency. If you'd like to take a closer look, click the icon that is below your plot and third from the right, then select a section you'd like to zoom in to.

Now you've seen the python script running, let's go under the hood and have a look at how the FPGA is programmed and how data is interrogated. To stop the python script running, go back to the terminal and press *ctrl + c* a few times.

9 iPython walkthrough

The *spectrometer.py* script has quite a few lines of code, which you might find daunting at first. Fear not though, it's all pretty easy. To whet your whistle, let's start off by operating the spectrometer through iPython. Open up a terminal and type:

```
ipython --pylab
```

and press enter. You'll be transported into the magical world of iPython, where we can do our scripting line by line, similar to MATLAB. Our first command will be to import the python packages we're going to use:

```
import corr,time,numpy,struct,sys,logging,pylab
```

Next, we set a few variables:

```
katcp_port = 7147
roach='enter IP address or hostname here' eg.roach = 'roach030172'
timeout = 10
```

Which we can then use in *FpgaClient()* such that we can connect to the ROACH and issue commands to the FPGA:

```
fpga = corr.katcp_wrapper.FpgaClient(roach, katcp_port, timeout)
```

We now have an *fpga* object to play around with. To check if you managed to connect to your ROACH, type. We will true if connected :

```
fpga.is_connected()
```

Let's set the bitstream running using the `progdev()` command. We will get 'ok' after programming FPGA :

```
fpga.progdev('File3b')
```

Now we need to configure the accumulation length and gain by writing values to their registers. For two seconds and maximum gain: accumulation length, $2 \cdot (2^{28}) / 2048$, or just under 2 seconds:

```
fpga.write_int('acc_len', 2*(2**28)/2048)
fpga.write_int('gain', 0xffffffff)
```

Finally, we reset the counters:

```
fpga.write_int('cnt_rst', 1)
fpga.write_int('cnt_rst', 0)
```

To read out the integration number, we use `fpga.read_uint()`:

```
fpga.read_uint('acc_cnt')
```

Do this a few times, waiting a few seconds in between. You should be able to see this slowly rising. Now we're ready to plot a spectrum. We want to grab the *even* and *odd* registers of our PFB:

```
a_0=struct.unpack('>1024l', fpga.read('even', 1024*4, 0))
a_1=struct.unpack('>1024l', fpga.read('odd', 1024*4, 0))
```

These need to be interleaved, so we can plot the spectrum. We can use a for loop to do this:

```
interleave_a=[]

for i in range(1024):
    interleave_a.append(a_0[i])
    interleave_a.append(a_1[i])
```

Press Enter key two times to comeout of this for loop in ipython. This gives us a 2048 channel spectrum. Finally, we can plot the spectrum using pyLab:

```
pylab.figure(num=1, figsize=(10,10))
pylab.plot(interleave_a)
pylab.title('Integration number')
pylab.ylabel('Power (arbitrary units)')
pylab.grid()
pylab.xlabel('Channel')
pylab.xlim(0,2048)
pylab.show()
```

Voila! You have successfully controlled the ROACH spectrometer using python, and plotted a spectrum. Bravo! You should now have enough of an idea of what's going on to tackle the python script. Type `exit()` to quit ipython.

10 The python script “File3d” notes

Now you're ready to have a closer look at the “**File3d**” script. Open it with your favorite editor. Again, line by line is the only way to fully understand it, but to give you a head start, here's a few notes:

Connecting to the ROACH

To make a connection to the ROACH, we need to know what port to connect to, and the IP address or hostname of our ROACH. The connection is made on line 98:

```
88. fpga = corr.katcp_wrapper.FpgaClient(...)
```

The *katcp_port* variable is set on line 16, and the *roach* variable is passed to the script at the terminal (remember that you typed *python “File3d” roachname*). We can check if the connection worked by using *fpga.is_connected()*, which returns true or false:

```
101. if fpga.is_connected():
```

The next step is to get the right bitstream programmed onto the FPGA fabric. The *bitstream* is set on line 15:

```
18. bitstream = 'File3b'
```

Then the *progdev* command is issued on line 101:

```
110. fpga.progdev(bitstream)
```

Passing variables to the script

Starting from line 66, you'll see the following code:

```
from optparse import OptionParser

p = OptionParser()
p.set_usage('“File3d” <ROACH_HOSTNAME_or_IP> [options]')
p.set_description(__doc__)

p.add_option('-l', '--acc_len', dest='acc_len',
             type='int', default=2*(2**28)/2048,
             help='Set the number of vectors to accumulate between dumps. default is 2*(2^28)/2048, or just under 2 seconds.')

p.add_option('-g', '--gain', dest='gain',
             type='int', default=0xffffffff,
             help='Set the digital gain (6bit quantisation scalar). Default is 0xffffffff (max), good for wideband noise. Set lower for CW tones.')

p.add_option('-s', '--skip', dest='skip', action='store_true',
             help='Skip reprogramming the FPGA and configuring EQ.')

opts, args = p.parse_args(sys.argv[1:])
```

```
if args==[]:
    print 'Please specify a ROACH board. Run with the -h flag to see
    all options.\nExiting.'
    exit()
else:
    roach = args[0]
```

What this code does is set up some default parameters which we can pass to the script from the command line. If the flags aren't present, it will default to the values set here.

11 Conclusion

If you have followed this tutorial faithfully, you should now know:

- What a spectrometer is and what the important parameters for astronomy are.
- Which CASPER blocks you might want to use to make a spectrometer, and how to connect them up in Simulink.
- How to connect to and control a ROACH spectrometer using python scripting.

In the “Wideband poco” tutorial, you'll be looking at a wideband 'FX' correlator design on the ROACH. You'll see quite a few of the blocks you were introduced to today reused, as the 'F' of a 'FX' correlator is essentially a spectrometer. Alternatively, tutorial “Wideband GPU/Zooming spectrometer” will take you through how to build a “million channel spectrometer”, along with a few simulink tricks.