



Implementing FFT Accelerators with SDSoC™ Using Open-Source Software and C-Callable IP

Tools:	2018.2 Vivado & SDSoC
Version:	v2018.2-1
Date:	29 June 2018

© 2018 Avnet. All rights reserved. All trademarks and registered trademarks are the property of their respective owners. All specifications are subject to change without notice.

NOTICE OF DISCLAIMER: Avnet is providing this design, code, or information "as is." By providing the design, code, or information as one possible implementation of this feature, application, or standard, Avnet makes no representation that this implementation is free from any claims of infringement. You are responsible for obtaining any rights you may require for your implementation. Avnet expressly disclaims any warranty whatsoever with respect to the adequacy of the implementation, including but not limited to any warranties or representations that this implementation is free from claims of infringement and any implied warranties of merchantability or fitness for a particular purpose.

Introduction

The Fast Fourier Transform (FFT) is one of the fundamental building blocks of Digital Signal Processing (DSP) and Signal Analysis. Due to its frequent use, many device manufacturers offer code libraries and intellectual property (IP) optimized for their architecture in order to achieve the highest possible performance. Xilinx, for example, offers a fully customizable FFT IP core that is optimized for their FPGA, Zynq® SoC, and Zynq MPSoC programmable logic devices. When it comes to general purpose computers there are a few open source code libraries. One such open source library is the Fastest Fourier Transform in the West (FFTW) library which can be obtained from www.fftw.org.

This paper looks at the performance of two different single-precision floating-point FFT implementations using a Xilinx Zynq UltraScale+™ MPSoC device and the SDSoC Development Environment. The first implementation uses version 3.3.7 of the FFTW library compiled for the ARM® Cortex®-A53 processor within a Xilinx ZU3EG device. The second implementation is an FPGA accelerator using the Xilinx LogiCore™ IP FFT version 9.0 (XFFT) running in the programmable logic of the ZU3EG device.

Environment

The Avnet® UltraZed-EG™ Starter Kit was chosen as the test platform for evaluating FFT performance. The UltraZed features the Xilinx Zynq UltraScale+ MPSoC ZU3EG device. The ZU3EG device features a Processing System (PS) containing quad-core ARM Cortex-A53 application processors (APU), dual-core ARM Cortex-R5 real-time processors (RPU), and an ARM Mali™-400 MP2 GPU. Linux was chosen as the host operating system which provides symmetric multi-processing support. For an apples-to-apples comparison, this paper only looks at running the FFTW library on a single APU core running at 1.1GHz.

The ZU3EG device also features FPGA programmable logic (PL) with 360 dedicated DSP slices to support heavy computational workloads. The XFFT running in PL operates at a 300 MHz clock rate using the pipelined-streaming I/O architecture. During execution of the XFFT, the PS offloads processing to the PL and then waits for the results to be computed. This processing model is characteristic of a heterogeneous compute system with a host processor and co-processor accelerator. The Xilinx Zynq UltraScale+ MPSoC architecture combines the heterogeneous system into a single chip solution.

The SDSoC platform for the UltraZed-EG Starter Kit can be downloaded from the zedboard.org website using the link <http://www.zedboard.org/content/ultrazed-3eg-starter-kit-sdsoc-platform-sdsoc-20172>. The platform is built using the 2017.2 Xilinx tool suite, but can be used with the 2018.2 release without issue. For more details on the UltraZed-EG Starter Kit visit <http://www.zedboard.org/product/ultrazed-eg-starter-kit>.

Test Setup

Xilinx offers a feature-rich software defined development environment for implementing complex algorithms in embedded processors and FPGA fabric. One such tool is the SDSoC Development Environment. SDSoC stands for Software Defined System-on-Chip and allows for easy acceleration of software defined functions using FPGA programmable logic. SDSoC takes functions defined in C/C++ and moves them to FPGA fabric for acceleration. More details on SDSoC can be found on the Xilinx website at <https://www.xilinx.com/sdsoc>.

For the purpose of this paper, SDSoC is the perfect environment to run FFT implementations in both software and in FPGA fabric. Thus, a software based test bench was created as the main application calling both the FFTW and XFFT libraries. FFT execution time performance was measured using a built-in high-resolution timer operating at the 1.1GHz APU clock. Figure 1 depicts a block diagram of the SDSoC based test bench.

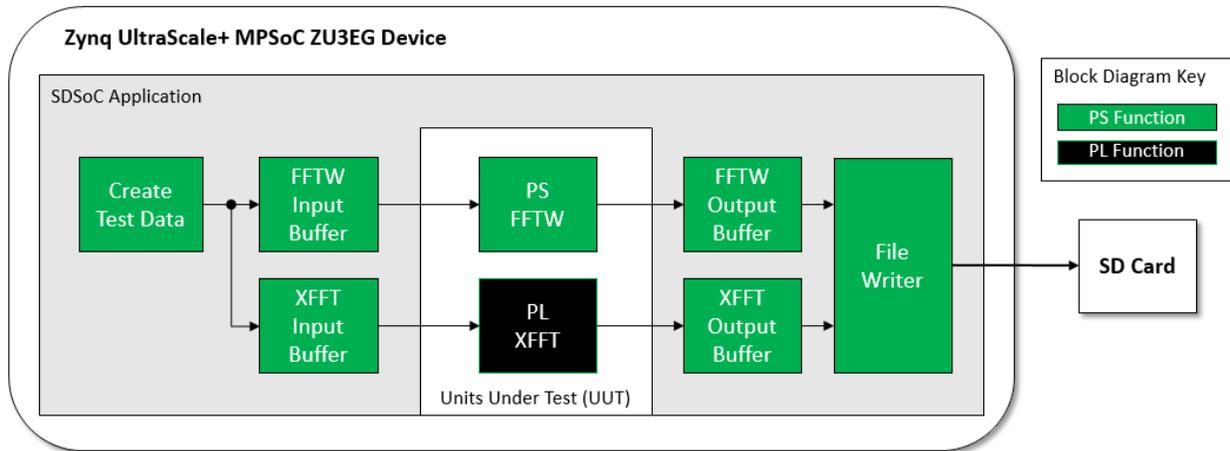


Figure 1 – SDSoC Test bench

Two methods of measuring execution time performance were used for XFFT analysis. The first method measures the average time of a blocking XFFT function call. The second method measures the average time of a non-blocking XFFT function call. Both methods perform 1000 function calls and compute the average execution time per call. With non-blocking function calls the SDSoC application continues execution after calling the PL accelerator without waiting for the accelerator to finish processing. A blocking function call waits for the PL accelerator to complete processing before proceeding. Performing non-blocking function calls helps overcome the overhead of transferring data between the PS and PL because the next set of data can be loaded while the current data is being processed. Figure 2 demonstrates the processing flow for the two measurement methods.

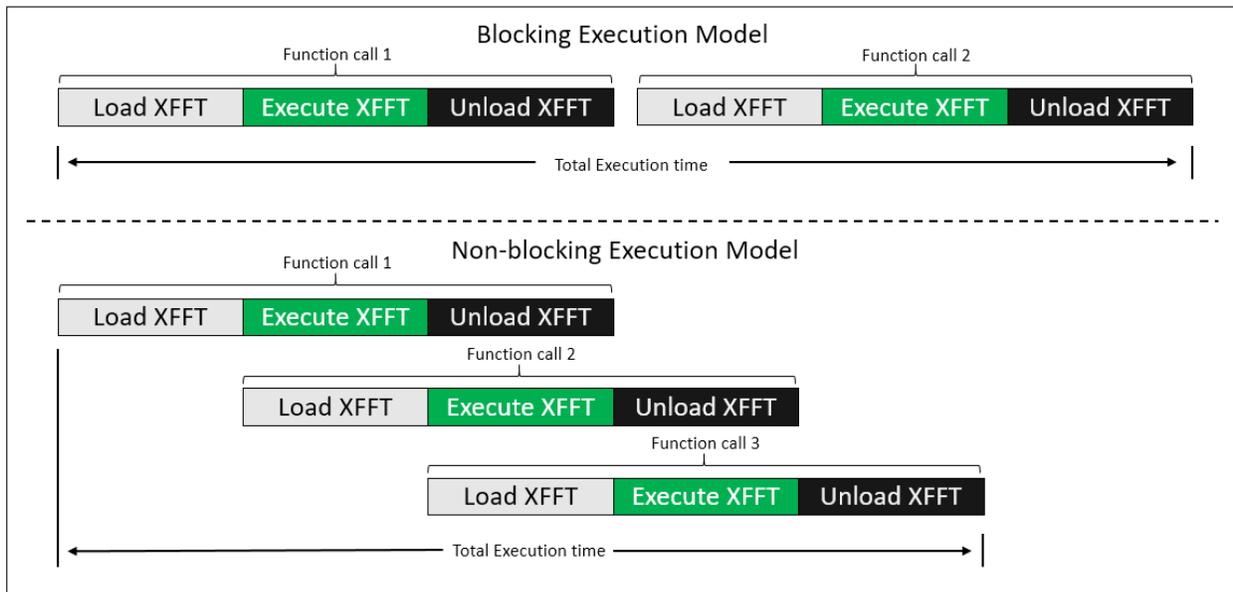


Figure 2 – PL Accelerator Execution Model

Results

Table 1 below gives the average round-trip processing times per function call for the FFTW and XFFT implementations for power-of-two FFT sizes ranging from 8 to 16384. Round-trip processing time includes any overhead associated with data movement between PS and PL.

FFT execution performance using FFTW and XFFT (FFTW in ARM Cortex-A53 @ 1.1GHz; XFFT in Xilinx PL @ 300MHz)			
FFT Size	FFTW Execution Time (us)	Blocking XFFT Execution Time (us)	Non-blocking XFFT Execution Time (us)
8	0.2	1.8	9.6
16	0.3	1.9	9.6
32	0.8	2.4	9.7
64	1.4	3.0	9.7
128	2.9	3.9	9.9
256	8.0	5.9	10.0
512	20.5	9.1	10.5
1024	56.9	16.0	13.0
2048	153.3	29.7	14.2
4096	399.1	57.0	25.0
8192	834.7	111.8	44.5
16384	2140.5	221.2	76.8

Table 1 – FFT Execution Time Performance

As seen in Table 1 above and Figure 3 below, the XFFT implementation has a longer execution time than the FFTW library for FFT sizes smaller than 256. For FFT sizes equal to or larger than 256 the XFFT implemented in the PL outperforms the FFTW library running on the ARM processor. Figure 3 also shows XFFT sizes 8 through 32 have approximately the same processing time. This indicates that the overhead of moving the data from the PS to the PL is overshadowing the FFT computation time. Thus, for small FFT sizes it does not make sense to offload a single FFT operation to a PL accelerator.

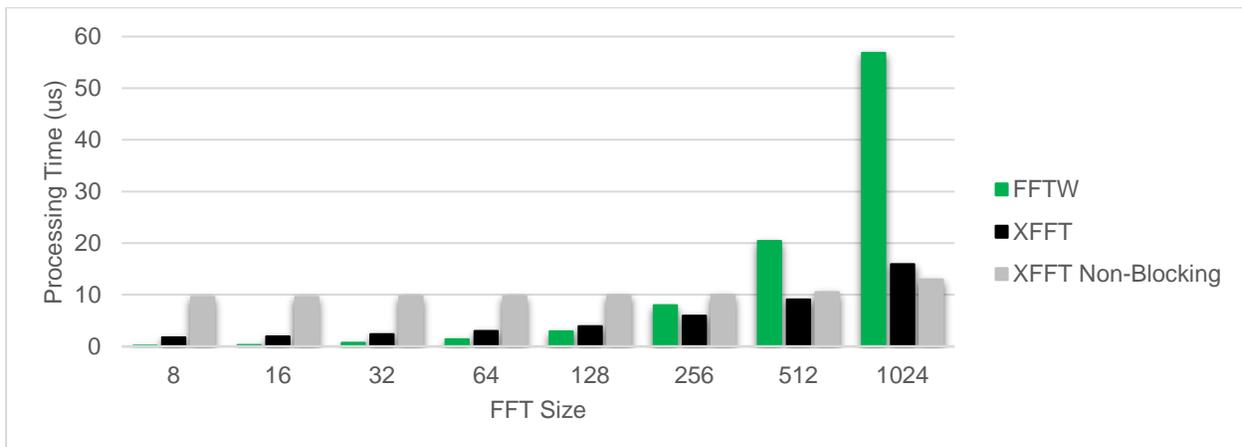


Figure 3 – FFT Processing Times (8 – 1024 point)

Figure 4 shows the processing times for FFT sizes larger than 1024 points. For the 16384 point FFT the processing time for FFTW is 2141 microseconds, the blocking XFFT is 221 microseconds, and the non-blocking XFFT implementation comes in at 77 microseconds. That corresponds to a speed up between 10X and 28X for executing FFT operations in the PL over the ARM processor.

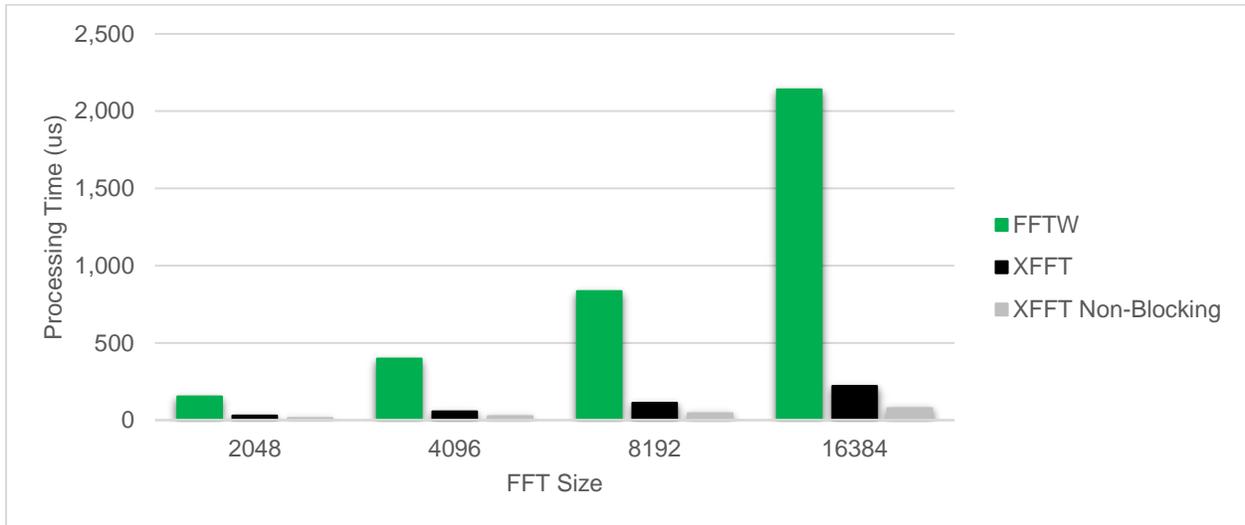


Figure 4 – FFT Processing Times (2048 – 16384 point)

As mentioned previously, smaller FFT sizes have a difficult time overcoming the data movement overhead. One way to get around this is to group multiple small transfers into one large transfer, i.e. group 100 16-point FFT data sets into one 1,600 word data transfer. Figure 5 below shows the results of doing this with 1,024 16-point FFTs. As shown in Figure 5, the average execution time for a 16-point FFT running in programmable logic is drastically reduced and now out performs the FFTW implementation.



Figure 5 – 16-point FFT Processing Times with Large Data Transfer
 (Note: HLS FFT was hand-written in C and compiled into programmable-logic using SDSoC)

Table 2 below summarizes the implementation accuracy using the mean-square-error (MSE) when compared to a double-precision MATLAB® model. The MSE statistic provides the average of the square difference between the single-precision floating-point FFT implementation and the double-precision floating-point MATLAB model. A small MSE indicates a good overall match between the model and the implementation. As indicated in Table 2, the FFTW and XFFT implementations closely match the MATLAB model.

FFT accuracy compared to double-precision MATLAB model		
FFT Size	FFTW Mean Square Error	XFFT Mean Square Error
8	1.6e-12	1.6e-12
16	3.4e-12	3.5e-12
32	5.7e-12	6.5e-12
64	1.3e-11	1.2e-11
128	2.4e-11	2.2e-11
256	4.6e-11	4.6e-11
512	8.9e-11	8.9e-11
1024	1.8e-10	1.7e-10
2048	3.5e-10	3.4e-10
4096	7.4e-10	6.9e-10
8192	1.5e-09	1.4e-09
16384	2.9e-09	2.8e-09

Table 2 – FFT Implementation Accuracy

Conclusion

This paper presented two options for implementing the Fast Fourier Transform in the Zynq UltraScale+ MPSoC device family from Xilinx. The first option was a software only implementation using the FFTW open-source code library which runs on the ARM Cortex-A53 processor. The second option used the Xilinx LogiCore IP to implement the FFT in programmable logic as an accelerator function. Significant gains were achieved with the accelerator which boosted execution time performance by a factor of 28 for the 16384-point FFT size. Smaller FFT sizes have difficulty overcoming the data movement overhead, but optimization can be achieved through strategic planning and control of data movement between the processing system and programmable logic. For the 16-point FFT size, execution speedup of 2.6X is possible with an accelerator when compared to the FFTW library running on the A53 processor.

Appendix A: Compiling FFTW for the ARM Cortex-A53 processor

This section describes the steps necessary to compile the FFTW library for use with the Zynq UltraScale+ MPSoC application processor. The following steps assume that appropriate Xilinx tools have been installed and the OS being targeted is Linux. The procedure below can be performed on a Windows machine, but it is recommended that a Linux machine is used (Note: the steps outlined below have not been tested in the Windows environment). Anywhere the <> brackets are used requires a replacement with a path for your specific environment – the text has also been changed to orange for your convenience.

1. Download the FFTW 3.3.7 source code from www.fftw.org/download.html and extract to your desired location, i.e. ~/fft_lib/fftw/.

Note: The rest of these instructions assume that ~/fft_lib/fftw is the location you extracted the source files to. If you used a different location then adjust accordingly.

2. Create a build directory in ~/fft_lib/fftw, i.e. ~/fft_lib/fftw/build
3. Change directory to ~/fft_lib/fftw/build
4. Create an install directory in ~/fft_lib/fftw/build i.e. ~/fft_lib/fftw/build/install
5. Setup your environment by executing the following on the Linux command line (assumes bash shell is being used)
 - a. export CROSS_COMPILE=aarch64-linux-gnu-
 - b. source <Xilinx-install-path>/SDK/2018.2/settings64.sh
6. Create a toolchain.make file with the following contents (this is for the A53 architecture, the A9 architecture uses a different toolchain):

```
set( CMAKE_SYSTEM_NAME Linux )
set( CMAKE_SYSTEM_PROCESSOR arm )
set( CMAKE_C_COMPILER aarch64-linux-gnu-gcc )
set( CMAKE_CXX_COMPILER aarch64-linux-gnu-g++ )
set( CMAKE_INSTALL_PREFIX <full-path-to-your-fftw-directory>/build/install )
set( CMAKE_FIND_ROOT_PATH <Xilinx-install-
path>/SDK/2018.2/gnu/aarch64/lin/aarch64-linux/bin )
```

7. At the Linux command prompt execute:
cmake -D CMAKE_TOOLCHAIN_FILE=toolchain.make -D CMAKE_INSTALL_PREFIX=<full-path-to-your-fftw-directory>/build/install <full-path-to-your-fftw-directory>
8. When cmake completes execute:
ccmake .

9. A menu will open that looks like:

```
BUILD_SHARED_LIBS ON
BUILD_TESTS ON
CMAKE_BUILD_TYPE Release
CMAKE_INSTALL_PREFIX /home/test/fft_lib/fftw/build/install
CMAKE_TOOLCHAIN_FILE /home/test/fft_lib/fftw/build/toolchain.make
ENABLE_AVX OFF
ENABLE_AVX2 OFF
ENABLE_FLOAT OFF
ENABLE_LONG_DOUBLE OFF
ENABLE_OPENMP OFF
ENABLE_QUAD_PRECISION OFF
ENABLE_SSE OFF
ENABLE_SSE2 OFF
ENABLE_THREADS OFF
WITH_COMBINED_THREADS OFF
```

Navigation within the cmake menu is performed using the arrow keys on your keyboard. To change/edit an option, navigate to the appropriate line and then press enter on the keyboard.

- a. Change `BUILD_SHARED_LIBS` to `OFF` by pressing enter on your keyboard
- b. Move the cursor to the `ENABLE_FLOAT` line and change it to `ON` by pressing enter on your keyboard
- c. Type `c` to configure
- d. Type `g` after configuration completes to generate the Makefile and exit the cmake menu

10. At the Linux command prompt execute:

```
make
```

11. When make completes execute the following at the Linux command prompt:

```
make install
```

12. The following files will be populated in your `~/fft_lib/fftw/build/install` directory

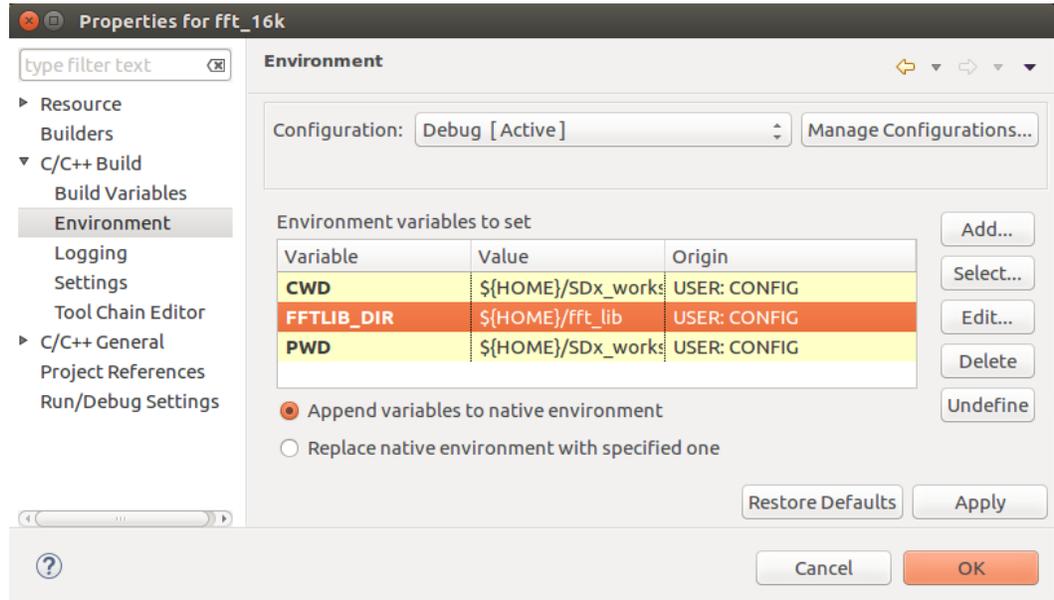
```
lib/libfftw3f.a
include/fftw3.h
```

These two files are needed for your SDSoc environment to correctly include and link the FFTW library

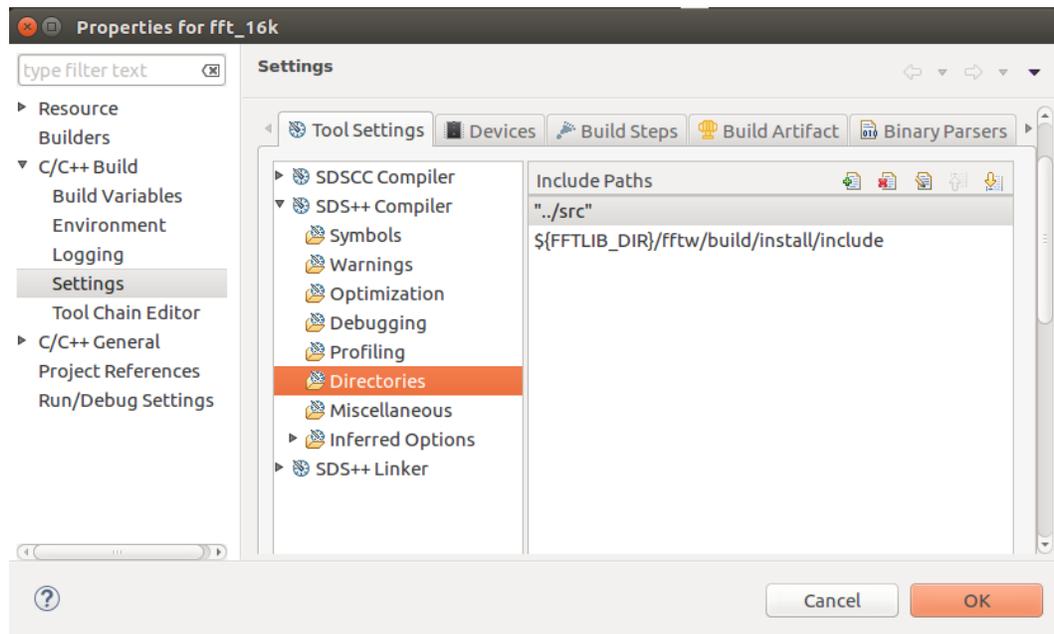
13. In your SDSoc project (assuming you have a project already set up that you want to include FFTW in, if not create one)

- a. Right-click on the project you want to add FFTW to and select `C/C++ Build Settings`

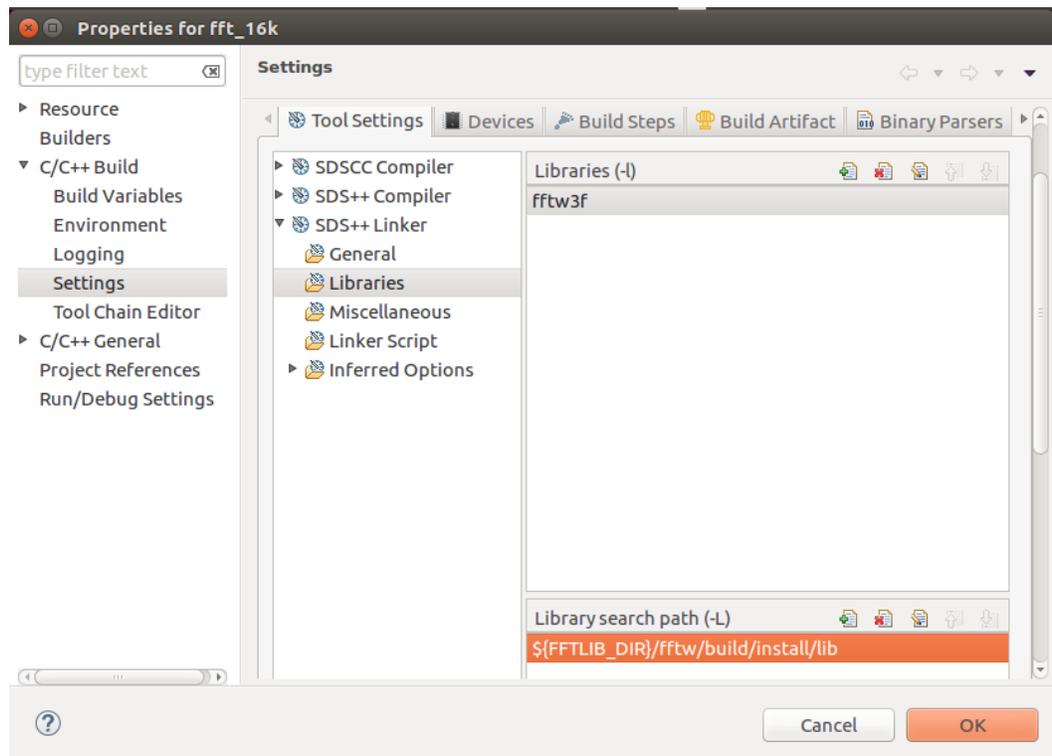
- b. Under the C/C++ Build → Environment menu set a build variable named FFTLIB_DIR that points to your fft_lib directory



- c. Click on Settings under C/C++ Build
- d. Assuming you are using C++, expand the SDS++ Compiler → Tool Settings. If you are using C then expand the SDSCC Compiler → Tool Settings.
- e. Under SDS++ Compiler → Directories add the path location of your include directory, i.e. \${FFTLIB_DIR}/fftw/build/install/include directory



- f. Under SDS++ Linker → Libraries add fftw3f to the Libraries dialog and the path location of the library (i.e. `$(FFTLIB_DIR)/fftw/build/install/lib`) to the Library search path dialog



14. Click OK and exit the project settings menu

15. You have successfully included the FFTW library for your ARM Cortex-A53 processor, see the `fftw3.h` header file for function prototypes and check out <http://www.fftw.org/index.html#documentation> for appropriate usage of the library functions

Appendix B: Creating a C-callable library for the Xilinx LogiCore IP FFT (v9.0)

Two different methods were employed when creating the C-callable XFFT library. This section covers both methods and why they were used. Anywhere the <> brackets are used requires a replacement with a path for your specific environment – the text has also been changed to orange for your convenience.

Method 1: Direct packaging of the XFFT IP

This method packages the XFFT IP directly without opening Vivado® and allows for full customization of the XFFT IP. There are some issues with using the run-time configurable transform length option, which will be explained in more detail for method 2. The following steps outline packaging of a C-callable library using a static FFT size and using the Xilinx LogiCore XFFT (v9.0) IP. These steps follow the command-line method defined in Chapter 4 of Xilinx UG1027 (v2018.2).

1. Create a folder to put your files, i.e. ~/fft_lib/xfft

Create a header file containing constants and function prototypes for the XFFT functions being compiled. For this example there are two functions

```
/* xfft.hpp */

#ifndef XFFT_H_
#define XFFT_H_

#include "stdint.h"

#define N_FFT 16384

typedef uint64_t data_t;
typedef uint16_t config_t;

void xfft( data_t x[N_FFT], data_t y[N_FFT] );
void xfft_config( config_t config[1] );

#endif
```

The `xffft()` function is used to call the accelerator while `xffft_config()` is used to configure the accelerator (i.e. forward versus inverse, etc.). Each argument must map to a physical AXI Interface on the underlying accelerator. For detailed port information for the XFFT IP please see PG109.

Note: The data type for `x` & `y` is given as `uint64_t`. The XFFT core has a 64-bit AXI input interface that expects the imaginary component in the upper 32-bits and the real component in the lower 32-bits. We could create a `struct` or use the `complex<float>` type instead of `uint64_t` and rely on the SDSoc compiler to pack the data into a 64-bit container before sending it to the XFFT accelerator. The important thing is making sure the size of the data being transferred is aligned with the size expected by the accelerator.

2. Create a makefile with function argument mapping and parameter definitions

a. Add the `sdx_pack` command to the makefile

```
sdx_pack -header xfft.hpp -lib libxffft.a \
```

b. Add the function argument map to the makefile

```
-func xfft -map x=S_AXIS_DATA:in -map y=M_AXIS_DATA:out -func-end \  
-func xfft_config -map config=S_AXIS_CONFIG:in -func-end \
```

The purpose of the function argument map is to map software arguments to physical interfaces on the accelerator.

There needs to be a separate function map for each software function, but they can be co-located in the same makefile as shown above. You can see that the `xffft_config()` function parameter `config` maps to the `S_AXI_CONFIG` (configuration) interface of the XFFT IP core. It is possible to get rid of the `xffft_config()` function and add a new argument to `xffft()`. However, it's best to separate configuration from processing so the accelerator does not get configured each time it is called. This in turn reduces the processing overhead. The accelerator can be configured once at the beginning of the application and then called multiple times before reconfiguring.

c. Add the IP repository path to the makefile

```
-ip  
<install_directory>/Vivado/2018.2/data/ip/xilinx/xffft_v9_0/component.xml  
-control none \
```

d. Add parameter customizations to the makefile

```
-param transform_length="16384" \  
-param implementation_options="pipelined_streaming_io" \  
-param data_format="floating_point" \  
-param input_width="32" \  
-param phase_factor_width="24" \  
-param output_ordering="natural_order" \
```

The `-param` flag sets parameters for the IP which would typically correspond to a HDL generic parameter. For example, the XFFT IP has a parameter called `data_format` which is set to `floating_point` above. Chapter 4 of PG109 lists all of the available User Parameters for the XFFT IP as well as acceptable values.

e. Add the target device family and operating system to the makefile

```
-target-family zynqplus -target-cpu cortex-a53 -target-os linux \
```

The final Makefile is shown below

```
all: libxfft.a

libxfft.a: xfft.hpp
    sdx_pack -header xfft.hpp -lib libxfft.a \
    -func xfft -map x=S_AXIS_DATA:in -map y=M_AXIS_DATA:out -func-end \
    -func xfft_config -map config=S_AXIS_CONFIG:in -func-end \
    -ip /opt/Xilinx/Vivado/2018.2/data/ip/xilinx/xfft_v9_0/component.xml -control none \
    -param transform_length="16384" \
    -param implementation_options="pipelined_streaming_io" \
    -param data_format="floating_point" \
    -param input_width="32" \
    -param phase_factor_width="24" \
    -param output_ordering="natural_order" \
    -target-family zynqplus -target-cpu cortex-a53 -target-os linux \
    -verbose

clean:
    rm -rf reports

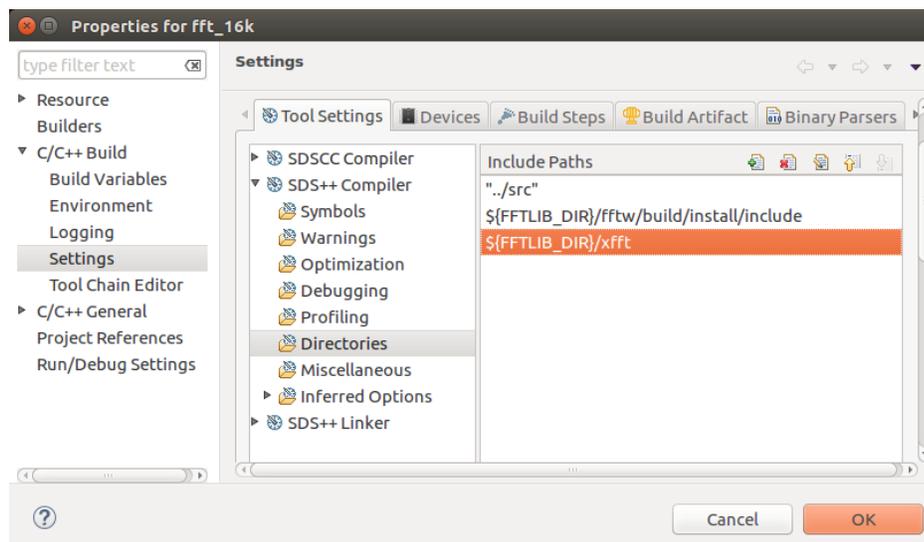
ultraclean:
    rm -rf reports
    rm -rf _sds .Xil
```

3. Packaging the IP

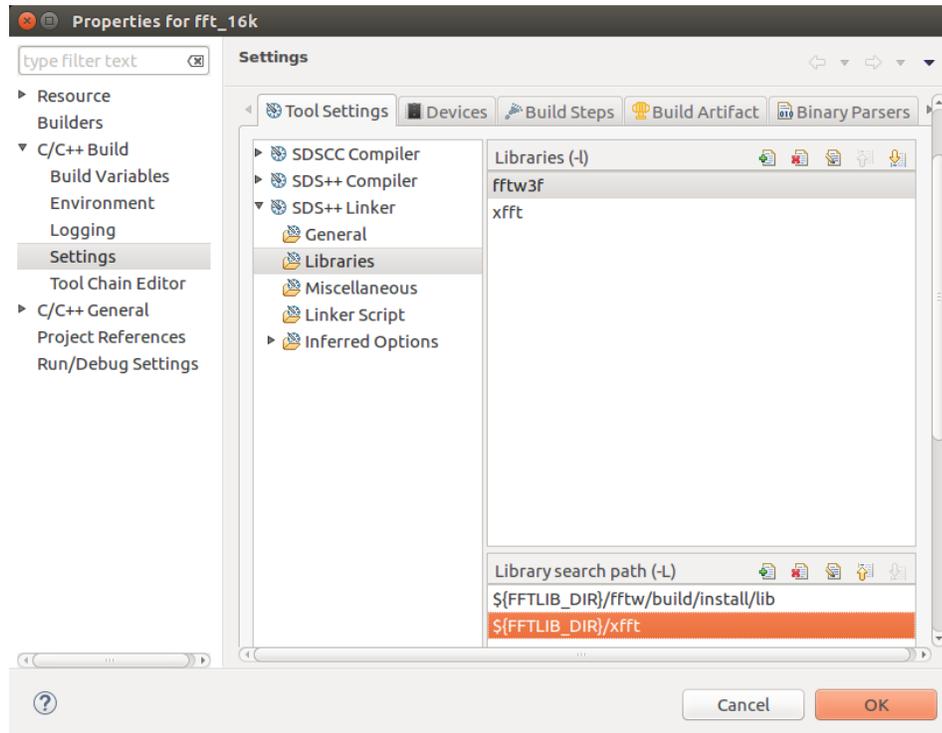
- a. Open the Linux terminal if not already open and execute
`source <install directory>/SDx/2018.2/settings64.sh`
- b. From the terminal navigate to the location where you saved your header file and Makefile – if they are not all in the same directory move them there now (i.e. ~/fft_lib/xfft/)
- c. At the command line type `make` and press `enter`
- d. After the build completes you should see a `libxfft.a` file in the directory. This file is used by the SDS++ linker to include the accelerator in your SDSoC project.

4. To include your C-callable IP in SDSoC

- a. Right-click on the project you want to add XFFT to and select `C/C++ Build Settings`
- b. Assuming you are using C++, expand the `SDS++ Compiler` under the `Tool Settings` tab. If you are using C then expand the `SDSCC Compiler`.
- c. Under `SDS++ Compiler` → `Directories` add the path location of your directory containing the `xfft.hpp` file (i.e. /home/<user>/fft_lib/xfft – Note: figure below uses environment variable `FFTLIB_DIR` defined in Appendix A as /home/<user>/fft_lib)



- d. Under SDS++ Linker → Libraries add `xfft` to the Libraries (-l) dialog and the path location of the `libxfft.a` library (i.e. `$(FFTLIB_DIR)/xfft`) to the Library search path (-L) dialog (figure below)



5. Click OK and exit the project settings menu
6. You have successfully included the XFFT library
 - a. To use XFFT make sure to call `xfft_config()` with the correct configuration parameters before calling `xfft()`
 - b. You do not need to specify the `xfft()` function for hardware acceleration within the SDSoC GUI – C-callable IP are hardware accelerated by definition
 - c. Configuration parameters are dependent on your implementation and will vary based on how you parameterized the XFFT IP in your `params.xml` file – for more details on the XFFT configuration interface see PG109

Method 2: Packaging Using a HDL Wrapper

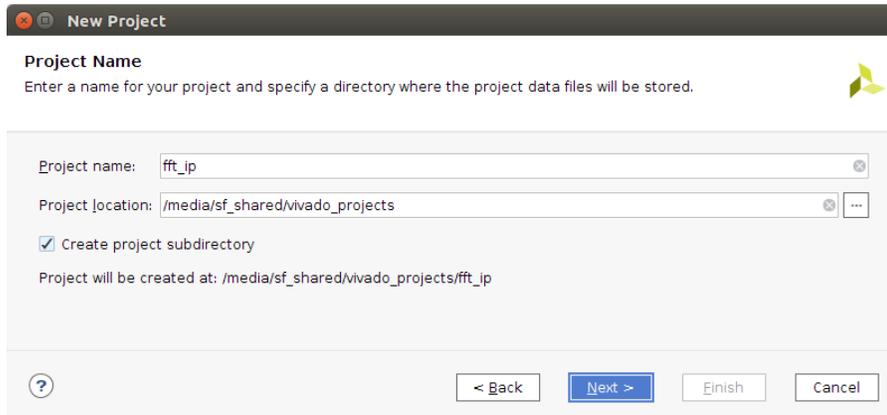
There are some issues with packaging the XFFT IP directly with the run-time configurable transform length configuration option. In order to transfer a variable amount of data using SDSoC from the PS to the PL, the data copy (or zero_copy) pragma must be used to specify the amount of data. The copy pragma can use a run-time determinable parameter to decide how much data to move, but the parameter used to compute this information must be part of the function call. The following example illustrates this point.

Example – using copy pragma to specify variable length transfer between PS and PL:

```
#pragma SDS data access_pattern(x:SEQUENTIAL, y:SEQUENTIAL)
#pragma SDS data copy(size[0:0], x[0:*size], y[0:*size])
#pragma SDS data data_mover(x:AXIDMA_SIMPLE, y:AXIDMA_SIMPLE)
void xfft( int size[], data_t x[], data_t y[] );
```

In the above example the size parameter is needed in the function prototype to determine the amount of data to move from PS to PL. For C-callable IP, there must be a one-to-one mapping between the function arguments and the ports on the IP. Thus, there must be an interface port on the IP that we can map the size argument to. This is where the HDL wrapper comes in. The HDL wrapper provides an AXI-Stream interface for the size argument that is left unconnected internally. That way SDSoC requirements are satisfied and we can use a run-time configurable FFT size. The following steps work through creating and packaging the XFFT IP using Vivado.

1. Create a Vivado project named `fft_ip` and click Next



New Project

Project Name
Enter a name for your project and specify a directory where the project data files will be stored.

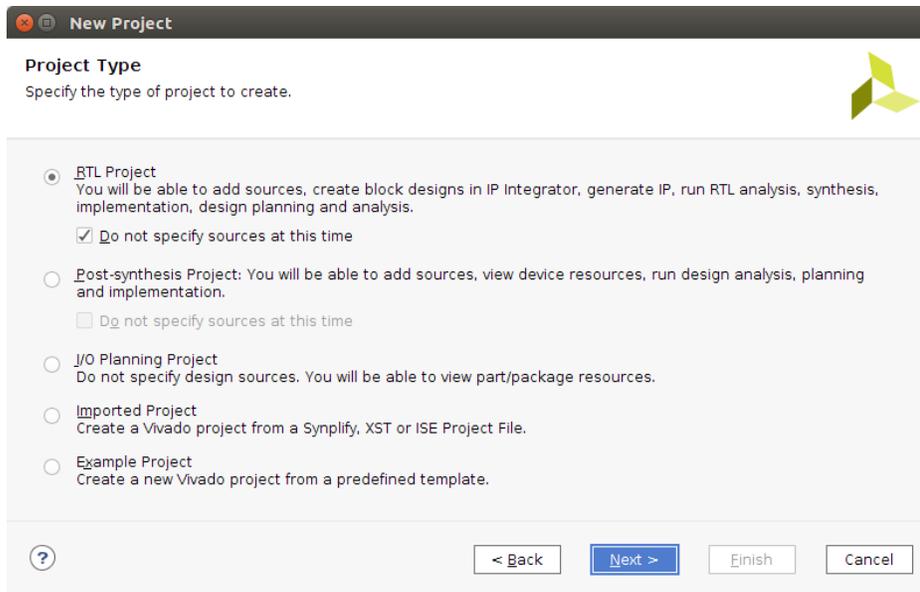
Project name:

Project location:

Create project subdirectory

Project will be created at: /media/sf_shared/vivado_projects/fft_ip

2. For the project type select **RTL Project** and check the “Do not specify sources at this time” box then click Next



New Project

Project Type
Specify the type of project to create.

RTL Project
You will be able to add sources, create block designs in IP Integrator, generate IP, run RTL analysis, synthesis, implementation, design planning and analysis.
 Do not specify sources at this time

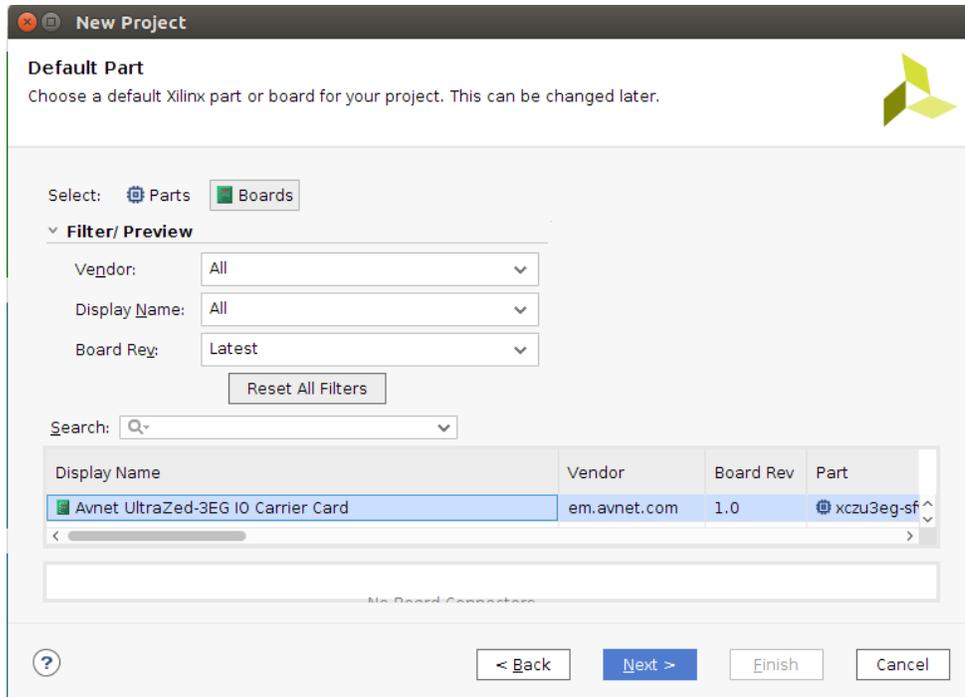
Post-synthesis Project: You will be able to add sources, view device resources, run design analysis, planning and implementation.
 Do not specify sources at this time

I/O Planning Project
Do not specify design sources. You will be able to view part/package resources.

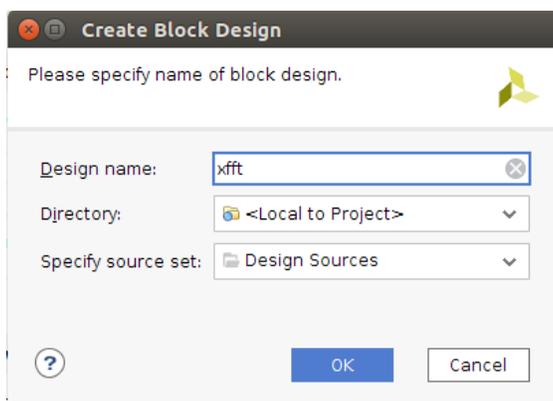
Imported Project
Create a Vivado project from a Synplify, XST or ISE Project File.

Example Project
Create a new Vivado project from a predefined template.

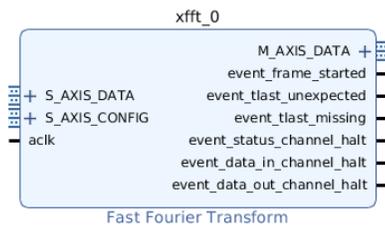
3. Select the Avnet UltraZed-3EG IO Carrier Card board and click Next
Note: this step assumes that you have already installed the UltraZed IO Carrier Card board definition files which can be found at http://zedboard.org/sites/default/files/documentations/UltraZed Board Definition Files v2017_2_Release_All_CC_5_0.zip)



4. Click Finish to create the project
5. Create a new block design by clicking on "Create Block Design" under the Flow Navigator window
 - a. Name the block design xfft and click OK



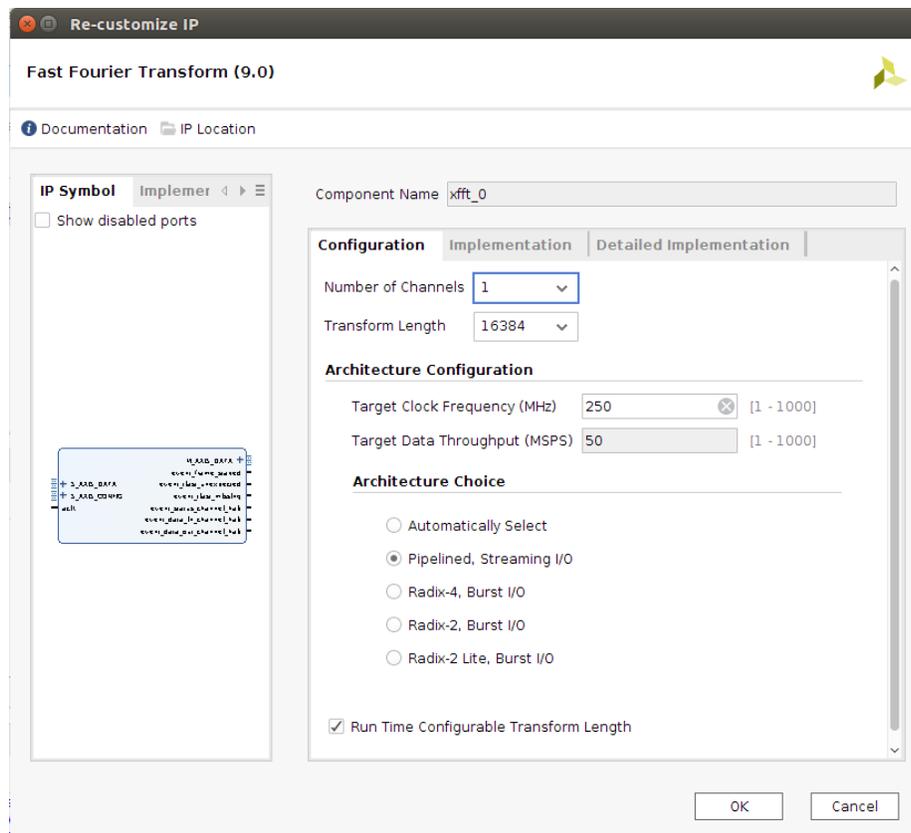
6. Add the Fast Fourier Transform IP to the block design canvas



7. Double-click the IP to configure

a. On the Configuration tab

- i. Set the Transform Length to the desired maximum.
For this study the maximum was set to 16384
- ii. Set the Architecture to "Pipelined, Streaming I/O"
- iii. Check the "Run Time Configurable Transform Length" box



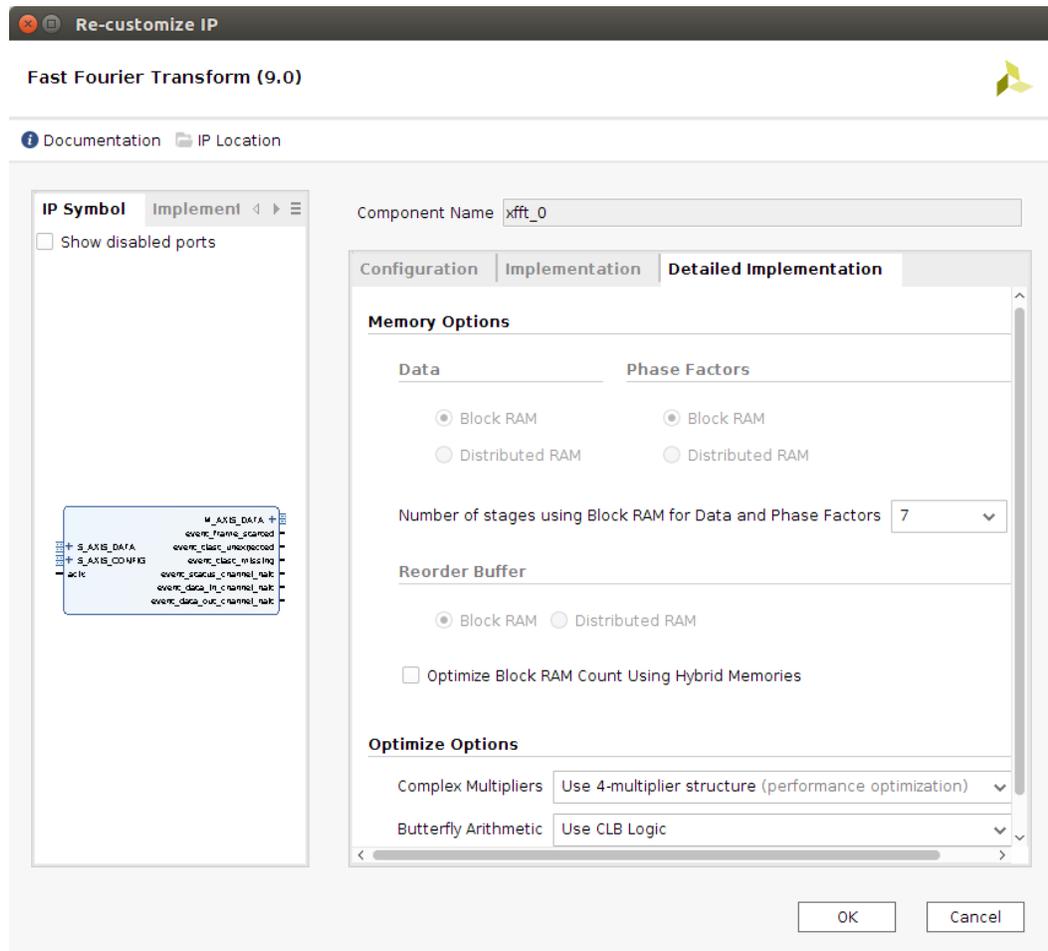
- b. On the Implementation tab
 - i. Set Data Format to “Floating-Point”
 - ii. Set Output Ordering to “Natural Order”
 - iii. Leave other fields with the default values

The screenshot shows the 'Re-customize IP' dialog for the 'Fast Fourier Transform (9.0)' IP block. The 'Implementation' tab is selected, and the component name is 'xfft_0'. The configuration options are as follows:

- Data Format:** Floating Point
- Scaling Options:** Scaled
- Rounding Modes:** Truncation
- Precision Options:**
 - Input Data Width: 32
 - Phase Factor Width: 24
- Control Signals:**
 - ACLKEN
 - ARESETn (active low)
 - ARESETn must be asserted for a minimum of 2 cycles
- Output Ordering Options:**
 - Output Ordering: Natural Order
 - Cyclic Prefix Insertion
- Optional Output Fields:**
 - XK_INDEX
 - OVFLO
- Throttle Scheme:**
 - Non Real Time
 - Real Time

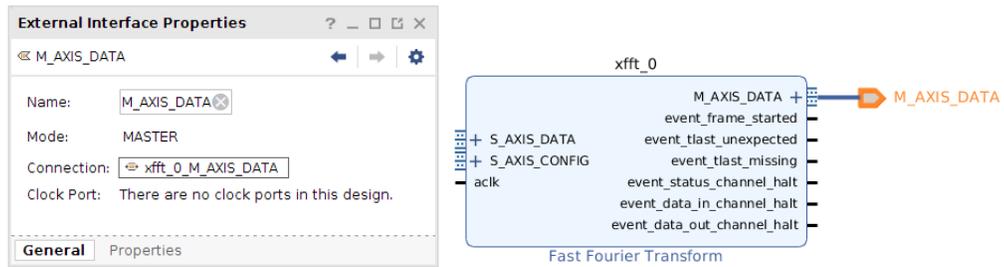
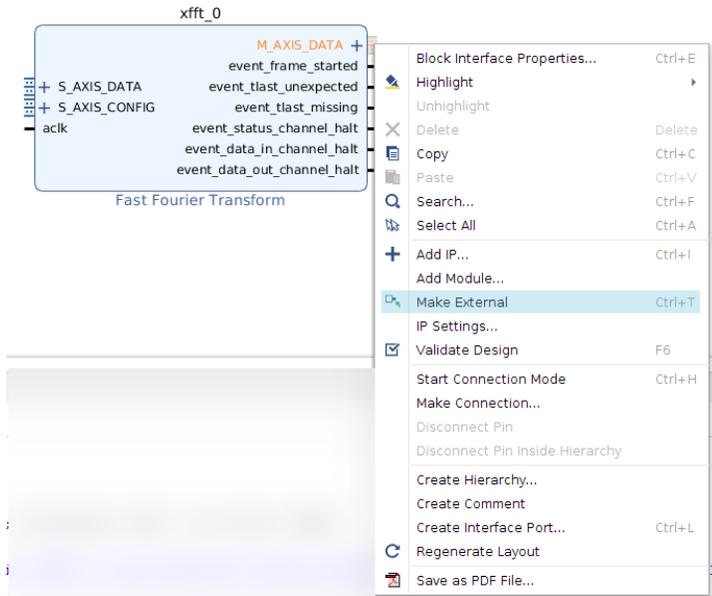
At the bottom of the dialog, there are 'OK' and 'Cancel' buttons.

- c. On the Detailed Implementation tab
 - i. Under Optimize Options set the Complex Multipliers option to “Use 4-multiplier structure”

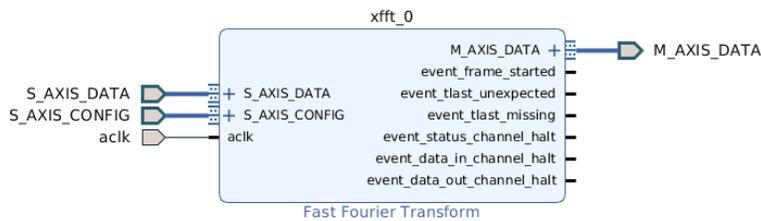


- d. Click OK to finish customizing the IP

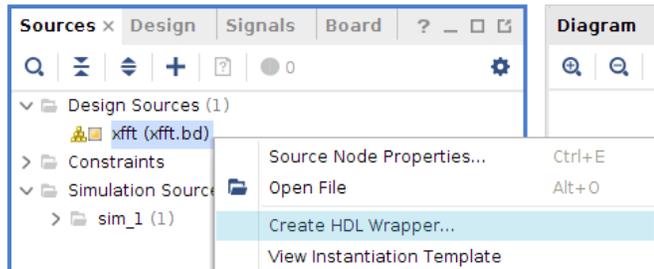
8. Make the XFFT IP ports external
 - a. Right-click on M_AXIS_DATA and select "Make External". Vivado will add "_0" to the end of the port name, so modify the port to remove the "_0"



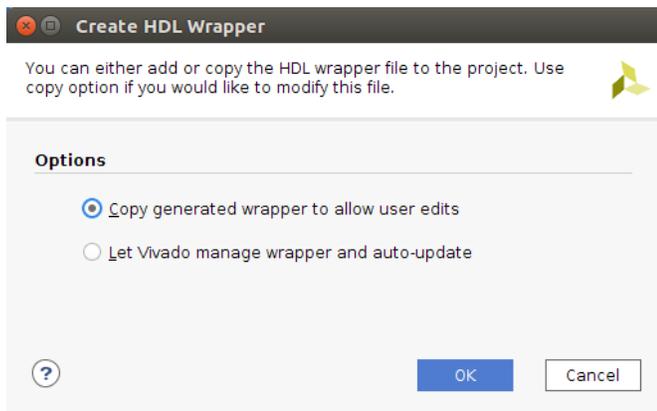
- b. Repeat step (a) for the aclk, S_AXIS_DATA, and S_AXIS_CONFIG ports



9. Create a HDL wrapper file
 - a. In the Sources window right click on xfft and select “Create HDL Wrapper”



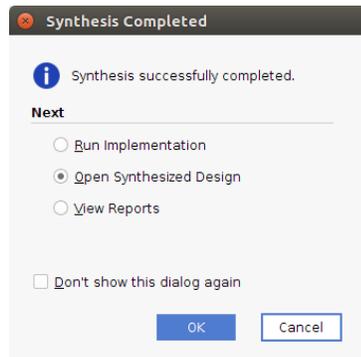
- b. Choose “Copy generated wrapper to allow user edits” option in the pop-up dialog and click OK



10. Add an AXI-Stream interface to the HDL wrapper for the “size” argument to map to

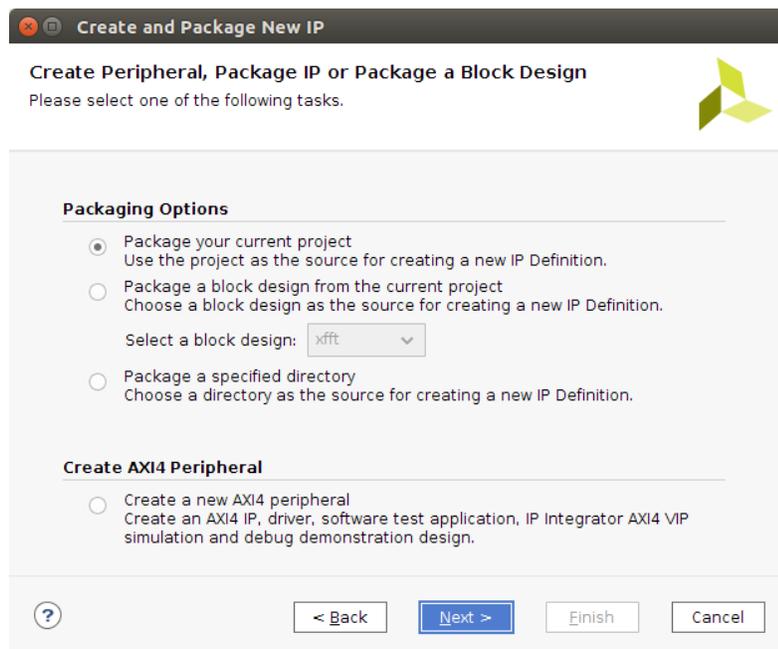
```
entity xfft_wrapper is
port (
  M_AXIS_DATA_tdata : out STD_LOGIC_VECTOR ( 63 downto 0 );
  M_AXIS_DATA_tlast : out STD_LOGIC;
  M_AXIS_DATA_tready : in STD_LOGIC;
  M_AXIS_DATA_tvalid : out STD_LOGIC;
  S_AXIS_CONFIG_tdata : in STD_LOGIC_VECTOR ( 23 downto 0 );
  S_AXIS_CONFIG_tready : out STD_LOGIC;
  S_AXIS_CONFIG_tvalid : in STD_LOGIC;
  S_AXIS_SIZE_tdata : in STD_LOGIC_VECTOR ( 31 downto 0 );
  S_AXIS_SIZE_tvalid : in STD_LOGIC;
  S_AXIS_DATA_tdata : in STD_LOGIC_VECTOR ( 63 downto 0 );
  S_AXIS_DATA_tlast : in STD_LOGIC;
  S_AXIS_DATA_tready : out STD_LOGIC;
  S_AXIS_DATA_tvalid : in STD_LOGIC;
  aclk : in STD_LOGIC
);
end xfft_wrapper;
```

11. Synthesize the design. When synthesis completes click on Cancel to close the pop-up

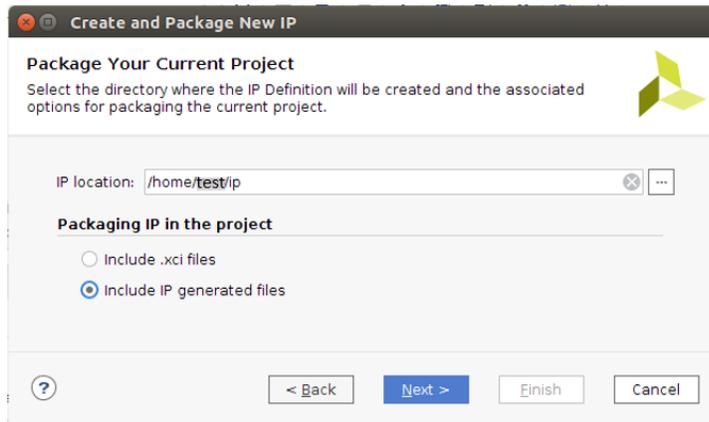


12. Package the project as IP

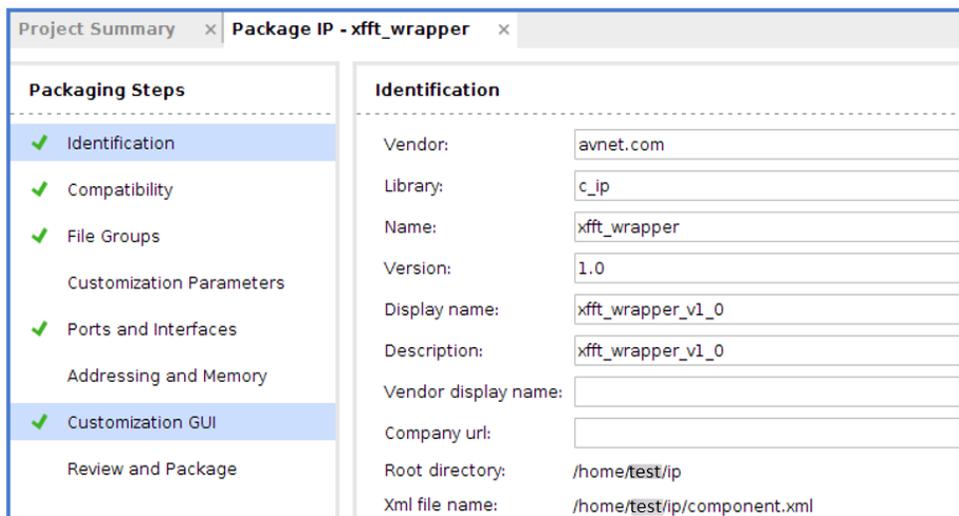
- a. Go to Tools → Create and Package New IP
- b. Click Next
- c. Choose “Package your current project” and click Next



- d. Set the IP Location to your desired output path (i.e. ~/ip) and select “Include IP generated files” then click Next



- e. Click Finish and OK if you get a pop-up message about copying IP to your output location
- f. On the “New IP Creation” window click Finish
- g. A new Vivado project will open
 - i. On the “Package IP” window under “Packaging Steps > Identification” fill out the information for your IP. This will be used later during the C-Callable IP packaging steps in SDSoc



- ii. Under “Packaging Steps > Review and Package” click Package IP
- h. Close the Vivado project when the IP packaging has completed

13. Change directory to the IP output location (i.e. ~/ip)
 - a. Create a new directory using the VLNV format (Vendor, library, name, and version). This information was used in step 11.g.i.

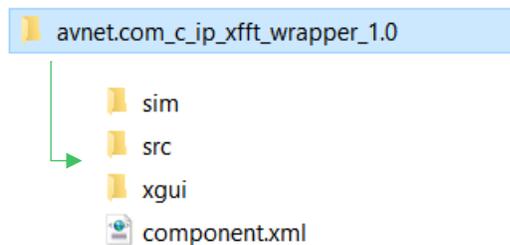
For example, if your VLNV was

- Vendor = avnet.com
- Library = c_ip
- Name = xfft_wrapper
- Version = 1.0

Then you would create a directory named avnet.com_c_ip_xfft_wrapper_1.0.

Copy the contents of the IP output to your newly created directory.

Your directory structure should look like:



14. The remainder of this procedure follows all steps outlined in Method 1 with a few minor differences noted below
 - a. Update xfft.hpp to account for port changes and to add data copy pragmas (updates highlighted in yellow below)

```
/* xfft.hpp */

#ifndef XFFT_H_
#define XFFT_H_

#include "stdint.h"

#define N_FFT 16384

typedef uint64_t data_t;
typedef struct{char data[3];} config_t;

#pragma SDS data access_pattern(x:SEQUENTIAL, y:SEQUENTIAL)
#pragma SDS data copy(size[0:0], x[0:*size], y[0:*size])
#pragma SDS data data_mover(x:AXIDMA_SIMPLE, y:AXIDMA_SIMPLE)
void xfft( int size[], data_t x[], data_t y[] );

void xfft_config( config_t config[1] );

#endif
```

Note: 1 – Argument *config* of function *xfft_config* changed from a 16-bit to 24-bit data type to allow for run-time selection of the FFT size
2 – AXIDMA_SIMPLE data mover should only be used when doing blocking *xfft* function calls

- b. Update the Makefile to incorporate the size argument, update the IP location, and remove parameter settings (since these were done when the IP was packaged in Vivado). An updated Makefile is shown in the figure below – updates are highlighted.

```
all: libxfft.a

libxfft.a: xfft.hpp
    sdx_pack -header xfft.hpp -lib libxfft.a \
    -func xfft -map size=S_AXIS_SIZE:in -map x=S_AXIS_DATA:in -map y=M_AXIS_DATA:out -func-end \
    -func xfft_config -map config=S_AXIS_CONFIG:in -func-end \
    -ip ~/ip/avnet.com_c_ip_xfft_wrapper_1.0/component.xml -control none \
    -target-family zynqplus -target-cpu cortex-a53 -target-os linux \
    -verbose

clean:
    rm -rf reports

ultraclean:
    rm -rf reports
    rm -rf _sds .Xil
```

- c. Complete IP packaging as defined in Method 1 starting with step 3.

Revision History

Date	Version	Revision
6/13/2018	V2017.4-1	Initial release for 2017.4
7/2/2018	V2018.2-1	Updated for SDx 2018.2 C-Callable IP packaging updates