



# TangerineSDR



---

*TangerineSDR*  
*Data Engine Firmware / Software*  
**Architecture Document and Test Plan**

---

Version Number: 0.1 *PRELIMINARY*

Version Date: October 6, 2022

Document Number: TBD

# VERSION HISTORY

Version Number	Implemented By	Revision Date	Approved By	Approval Date	Description of Change
0.1	T. McDermott	September 19 - October 6, 2022			Original Issue. Much missing content.

<b>1. INTRODUCTION.....</b>	<b>4</b>
<b>1.1. Scope.....</b>	<b>4</b>
<b>1.2. FPGA Overview.....</b>	<b>4</b>
<b>1.3. Phased Development Approach.....</b>	<b>5</b>
1.3.1. Phase 1A.....	5
1.3.2. Phase 1B.....	5
1.3.3. Phase 1C.....	6
<b>1.4. Phase 2.....</b>	<b>6</b>
<b>2. DMA OF DATA I/O .....</b>	<b>6</b>
<b>2.1. Phased DMA implementation .....</b>	<b>8</b>
<b>2.2. Longer term DMA implementation .....</b>	<b>8</b>

## **1. Introduction**

### **1.1. Scope**

The Data Engine (DE) Version 1 provides the connection of TangerineSDR modules (Receiver, Clock, Magnetometer, VLF receiver) with an FPGA. The network interface connects the DE to the Local Host (LH) using 1 Gigabit Ethernet. The control interfaces to the hardware peripheral modules are via I2C and SPI. The clock module also uses USB connections but these are not provided by the DE. There are other module-specific data interfaces to the various receiver modules. The DE Version 1 hardware contains a MAX10 FPGA to control the system, and to communicate with the Local Host computer via 1 GbE.

The MAX10 Development Kit contains a subset of the interfaces provided by the DE. An adaptor board is used to connect the DevKit to the Receiver, Clock, and other peripherals. The DevKit will be used for initial prototype development until the actual DE is available and tested.

This document outlines the architecture of the FPGA firmware and software to accomplish the TangerineSDR initial objectives. It also discusses some alternative approaches for controlling the DMA engines.

### **1.2. FPGA Overview**

The FPGA firmware is created using Altera/Intel Quartus II version 20.1. written in Verilog, and synthesized to a binary loadable image for the FPGA. There are three main sections of FPGA firmware:

1. The digital signal processing, buffer management and signal I/O.
  - a. This code will be hand written in Verilog.
2. The NIOS II processor, peripherals, DDR3RAM, QSPI flash memory, internal RAM, DMA controllers, timers, and the Ethernet interface.
  - a. This code is designed using the Platform Designer (Qsys) graphical GUI tool. It is then synthesized to Verilog by the Quartus Qsys tool.
  - b. About 50% of the FPGA capacity is consumed by the processor and peripherals synthesized by the Qsys tools.
3. C-code running on the NIOS processor. This software initializes the system, controls the peripherals, and processes the Local-Host-to-Data-Engine protocol packets.
  - a. This code is written using the Quartus Eclipse-based design environment which provides the proper software libraries and linkage.

- b. The code relies on the Altera-supplied RTOS (uCOSII) and TCP/IP stack software (Nichestack). Nichestack is deprecated in later versions of the Quartus toolset.
- c. The NIOS executable code may be bundled with the Verilog from 1) and 2) into a single unified binary for download to the FPGA. The Quartus toolset also allows separating the software binary from 3) into a separate download should only the software portion need to be updated.

The Verilog generated for 1) and 2) above needs to co-exist in the Verilog module namespace. Traditionally this would be done by creating a top\_level Verilog module that instantiates both sections (NIOS and Tangerine). Currently however Quartus requires the Qsys Verilog to be defined as the top module. To solve this problem (perhaps temporarily perhaps not) the DSP and other Verilog from section 1) is included via a file include directive inside the Qsys generated top\_level so as to appear as a top level module along with the Qsys module. The Qsys top level module is called m10\_rgmii.v while the module from 1) is called Tangerine.V. In this way, all the code in the Tangerine module is fully isolated from the Qsys modules and appears as a separate top\_level module. If Qsys were to over-write the synthesized top\_level module, only a single `include directive needs to be retyped into m10\_rgmii.v

### **1.3. Phased Development Approach**

The initial FPGA firmware and software will be targeted at debugging the receiver module using the MAX10 DevKit.

#### **1.3.1. Phase 1A**

Phase 1A will test the I2C and SPI registers on the receiver. This will demonstrate proper connection of the MAX10 development kit, adaptor board, and receiver module as well as proper communication with the FPGA using gigabit Ethernet, processing the LH-DE protocol discovery of the DevBoard, and processing of the Module Read (MR) and Module Write (MW) commands. The test should be able to:

- Acquire a DHCP address (static IP not yet supported, MAC address hard-coded).
- Respond to OpenHPSDR discovery broadcast.
- Turn each of the 6 relays on and off and turn the two LEDs on and off.
- Read the Ident PROM unique serial number.
- Read and write the ADC via the SPI interface.

#### **1.3.2. Phase 1B**

Phase 1B will be to pack roughly 1500 bytes of data from a single receiver channel into a sequence of Ethernet frames on the data socket. These frames will consist of raw

ADC samples. Initially the ADC can be programmed to send various test patterns on the receiver DDR interface. This test will verify that the DDR interface is operating properly and the patterns can be properly received. Once that is achieved, the actual receiver data will be sent from one channel via the DDR interface. This will consist of a sequence of data frames that are contiguous in time for 16384 samples, then idle. These will be sent to gnuradio via a UDP socket block. The delimiting 0<sup>th</sup> sample will have a special mark in the header that allows a custom gnuradio OOT module to properly group this sequence into a single gnuradio vector. The sequence will be repeated at a slow rate of a few vectors per second. This is essentially a wideband spectrum from the receiver. A gnuradio flowgraph will be constructed to analyze these vectors to look for the presence of sampling and other spurs amid the base noise sequence.

### **1.3.3. Phase 1C**

Phase 1C will use a NCO / Sin-Cos generator and a sequence of CIC and FIR filters to source a single downconverted RF channel to the Ethernet stream. This stream will be sent to gnuradio using a standard UDP socket where the samples can be analyzed. Gnuradio flowgraphs will be constructed to analyze the Noise Figure of the receiver at a couple of discrete RF channel frequencies. Additionally (depending on test equipment availability) the 3<sup>rd</sup> order dynamic range of the receiver will be estimated.

It is anticipated that the tests through Phase 1 will be able to use processor-directed DMA engines and FIFOs to pass a limited amount of data to the attached gnuradio platform, which is anticipated to be a Linux based host, likely a general purpose desktop computer.

### **1.4. Phase 2**

Continued development during Phase 2 will include additional LH-DE protocol processing, multiple channels, clock module testing, etc. The needed data throughput is much higher. The details are TBD at this point.

## **2. DMA of Data I/O**

One key architectural issue is that the dual ADC converter receiver module can produce samples at a very high rate of speed. That sample rate needs to be reduced before the frames can be sent over the 1 GbE UDP socket. The concern is that the NIOS processor may not be fast enough to handle the rate of packets that could be emitted.

Some back-of-the-envelope numbers help illustrate the issue. The two receivers are synchronously clocked at 122.88 MHz and produce 14-bit receiver samples plus over-range. For simplicity this will be transferred as a 16-bit word for each channel per clock

cycle. The FPGA hardware clocks samples into the FPGA using a 16-bit wide differential DDR (Double Data Rate interface) operating at a 245.76 MHz clock rate. The total data sample rate is:

$$122.88 \text{ Ms/s} * 4 \text{ bytes} * 8 \text{ bits/byte} = 3.932 \text{ 16 Gb/sec.}$$
$$1 / 122.88 \text{ Ms/s} = 8.14 \text{ nanoseconds per 4 bytes.}$$

The 1 GbE interface cannot stream at 4 Gb/s obviously, so the data streams need to be decimated.

Each packet needs to have a UDP socket header added (presumably from a table constructed in RAM by the NIOS processor).

The protocol spec lists 1024 bytes of data per frame as desirable. For Ethernet, a 1500-byte frame is largest standard-sized frame allowed. Jumbo frames can allow up to about 9k bytes for frame. At a speed of 1 Gb/s, the frame time is:

$$1024 \text{ bytes} + 22 \text{ bytes overhead} = 8.514 \text{ microseconds per Ethernet frame}$$
$$1500 \text{ bytes} + 22 \text{ bytes overhead} = 12.176 \text{ microseconds}$$
$$9000 \text{ bytes} + 22 \text{ bytes overhead} = 72.176 \text{ microseconds.}$$

The NIOS processor may not be able to service an interrupt and handshake the frame each 8, or 12, or 72 microseconds. It needs to find the packet in memory, append the appropriate UDP header, then link the start address and length of the packet into the DMA controller descriptor table. After transmission (interrupt) it needs to de-link and recycle the sent buffer. With jumbo frames this problem becomes easier.

The NIOS processor has a maximum asynchronous clock speed of 160 MHz, but the actual processor core needs to run on the system clock (so that it can synchronize memory and peripheral reads and writes).

The processor is estimated to be able to process an assembly instruction roughly once each 40 nsec. period (25 MHz.) thus providing roughly:

- 212 instructions for 1024 byte data frames.
- 304 instruction cycles for normal frames.
- 1800 instruction cycles for Jumbo frames.

The actual interrupt handler must be significantly shorter in order to allow time for non-interrupt processing to occur.

## **2.1. Phased DMA implementation**

The first prototype of the implementation will try processor-based interrupt handling of the DMA controllers to test how practical that approach is, and to quickly get to a minimal level of receiver performance in order to characterize the receiver.

The receiver characterization needs two different types of data streams:

1. Downconverted, decimated receiver data. This is used to calculate the Noise Figure (NF) of the receiver. Different frequency bands will be hand-selected to provide NF estimates across the 100 KHz to 60 MHz range of the receiver. It also provides a way to characterize the dynamic range of the receiver.
2. Wideband interrupted data. This consists of a run of 16384 contiguous time samples with a long dead time in between each run. This allows measuring the presence of spurious signals across the entire receiver frequency range.

## **2.2. Longer term DMA implementation**

If shorter packets are needed (for example 1024 bytes) then the NIOS performance required increases. After Phase 1 it can be decided how the final DMA architecture should be done.

A phase 2 architectural approach might then be to have the NIOS processor setup all the receiver conditions (number of channels to be downconverted, frequencies, decimation rate, number of wideband channels and how they are time-gapped) and to set up buffer descriptors for all the above. Then Scatter-Gather DMA controllers would be responsible for transferring the data blocks to the Ethernet Verilog module. The Verilog would need to handle packing buffers and queuing them to the SGDMA engines, and de-queuing and recycling spent buffers.

The NIOS processor would still talk directly to the Ethernet frames via a socket interface, but it would only process command, control, and provisioning packets, not actual data packets.