

WHITE PAPER

Software Defined Radio in a multi-purpose COTS platform that can reduce the time to market your 5G products.



BOB MURO
Application Specialist
Mercury Systems



COTS Software Defined Radio for 5G Development

Commercial-Off-The-Shelf (COTS) Software-Defined-Radio (SDR) products have been traditionally used for Military RADAR and communication applications because of their high performance, and design flexibility. The latest SDR products offer SoC solutions with integrated I/O, ARM processors, and large FPGAs that include intellectual property (IP) for accessing, routing, and processing digital data. Combining these attributes with superior signal integrity, phase-coherent sampling, and multi-channel transceivers a COTS SDR system is an ideal choice for your 5G development platform. This article is designed to familiarize a new user about Software Defined Radio in a multi-purpose COTS platform that can reduce the time to market their 5G products.

COTS BACKGROUND

During the early 90’s congressional spending cuts meant the US military had to reduce R&D spending but had to retain global technological leadership for America’s security. In 1994, then Secretary of Defense, William Perry wrote a memo that requested the military increase the purchase of commercial products and the use of commercial practices & specifications. The memo was officially enacted into law as part of the Federal Acquisition Streamlining Act (FASA) of 1994 and was a departure from the more stringent military specification requirements of the past that prohibited the purchase of more general purpose commercial products. This created a new industry based upon commercially available high performance products that could be re-purposed for “dual-use”. As a result, an abundance of COTS products are available today with the latest technology available for both military & commercial purposes.

COTS SDR DEFINED

Each section of this article is divided into hardware, firmware, and software sub-sections for additional clarity. The hardware section is the SDR PCB including components, firmware is internal FPGA code to create a logic design and implement digital signal processing (DSP) functions, and software is C code to control an FPGA with firmware to perform any additional DSP functions.

HARDWARE

SDR replaces legacy analog systems that consisted of an RF Filter, analog down-converter (Lo + Mixer), band-pass filters, and a demodulator (see figure 1a). These fixed analog systems were limited to a specific platform like AM, or FM radio and needed to be replaced if another platform was required.

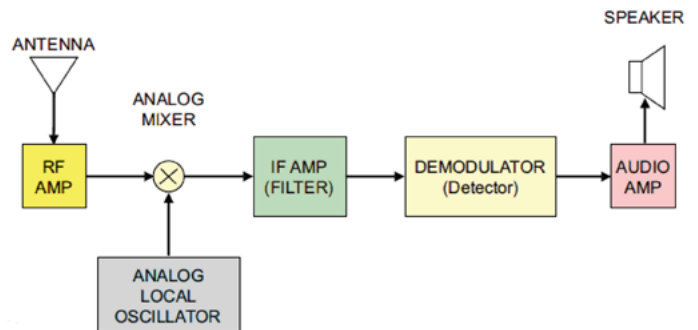


Figure 1a

The primary function of SDR is to exploit digital signal processing techniques to support the ever-increasing complexity, precision, and bandwidth of today’s radio traffic. Suitable data conversion is required between the antenna signals and the DSP operations for both receive and transmit functions [1].

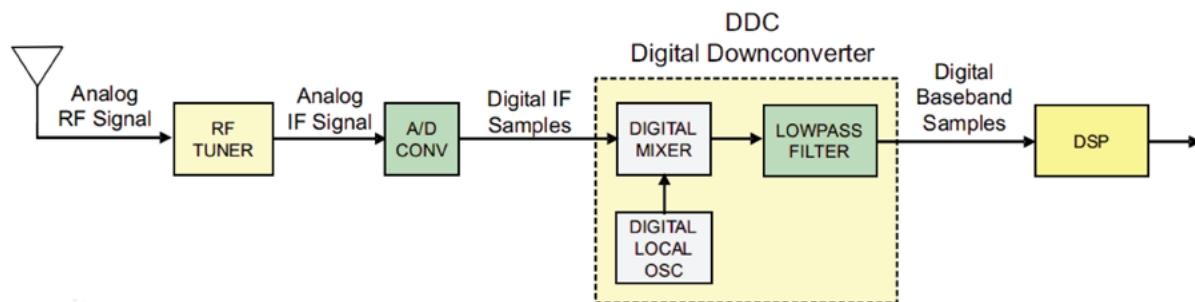


Figure 1b

An SDR receiver converts an RF signal from an antenna into digital samples with an A/D converter and uses subsequent DSP operations to extract the required information from the signal (see figure 1b).

An SDR transmitter accepts digital information to be transmitted and then performs the necessary DSP operations to produce digital samples for a D/A converter, whose analog output signal drives a PA for delivery to the antenna (see figure 1c).

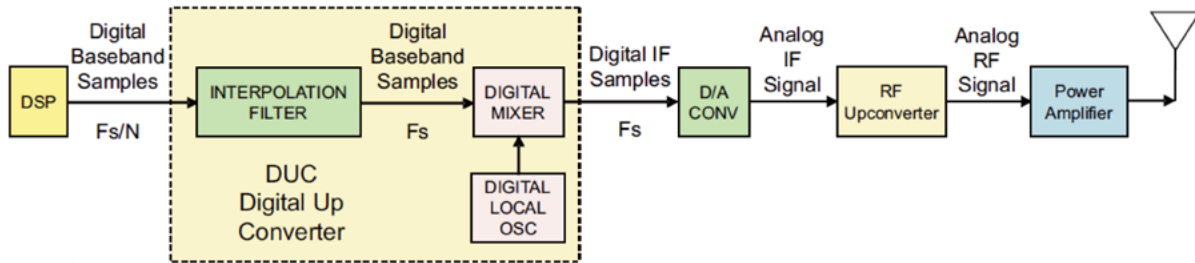


Figure 1c

Because these radios are software defined, they can be programmed on-the-fly with new parameters in microseconds, or re-configured for many different purposes by simply loading a new firmware image from internal or external memory.

An SDR is often implemented on a specialized PCB board called a “mezzanine card”. The current generation of these cards is a switched-fabric mezzanine card, an XMC or an FPGA mezzanine card, an FMC. Figure 2a contains an image of an XMC and FMC mezzanine card with the corresponding functional block diagram to the right. Image A is an XMC card

with four 200 MHz A/D channels, and image B is an FMC card with two 3.0 GHz A/D channels, and two 2.8 GHz D/A channels. Either of these mezzanine cards can be combined with a different form-factor carrier for installation in a different chassis, or on a PC motherboard. The primary difference is the FMC mezzanine card on Figure 2b does not include an FPGA and requires a more complex carrier card. This concept allows the same mezzanine card and carrier combination to be used on different platforms in multiple systems.

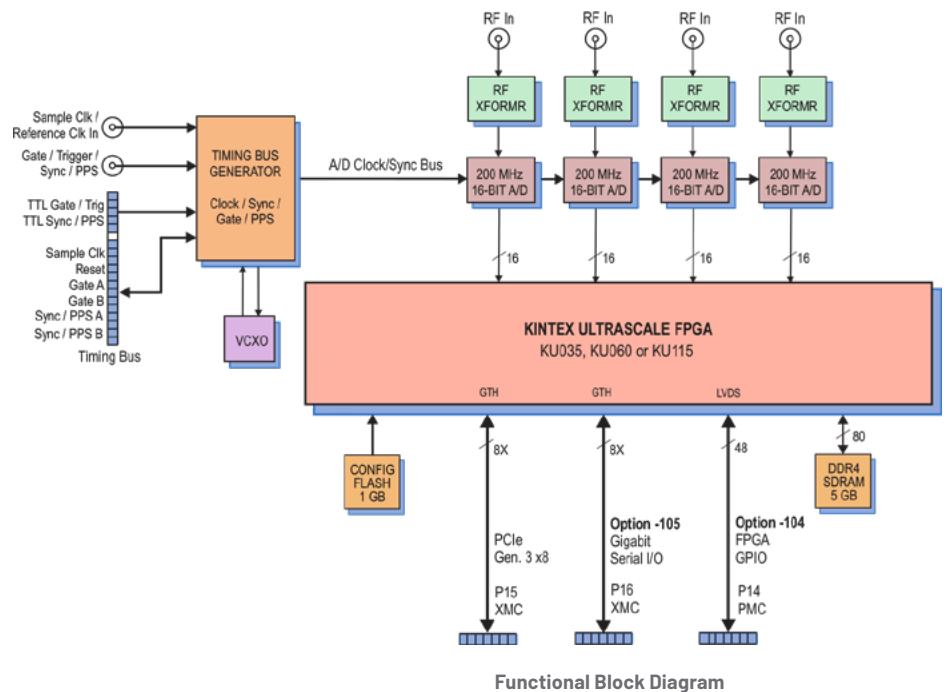



Figure 2a

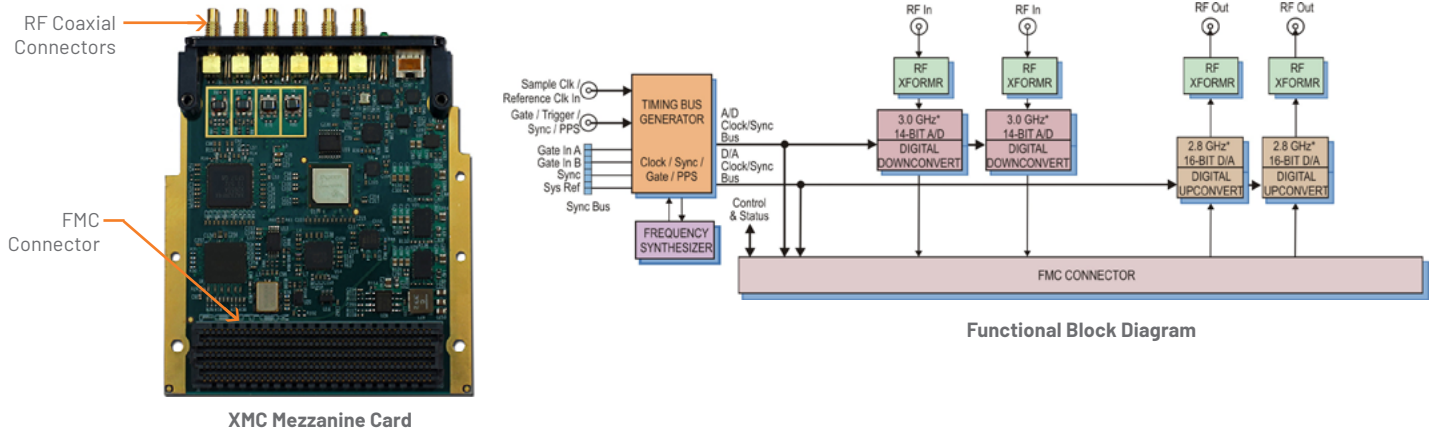


Figure 2b

Each SDR board includes a precision timing system with a multi-bit, fractional synthesizer for variable sampling rates locked to an on-board OCXO, or a reference input signal (see Figure 2b). These timing systems usually accept external synchronization signals from a Network Time Protocol (NTP) server or GPS receiver for the highly precise timing requirements of a radar or cellular system. The precision is also required for phase coherent sampling of the A/Ds, FPGA DSP data synchronization, and D/A signal transmission.


The XMC example A/D has a 200 MSPS maximum sampling rate that can capture a 100 MHz Nyquist BW excluding filtering. A common technique with digital radio is to acquire channel information, or intermediate frequency (IF) BW by under-sampling the signal. Figure 3 and text below explain the “the Fan-fold” concept using multiple Nyquist zones.

THEORY SIDE-BAR: NYQUIST ZONE SAMPLING


- 1) "Bandwidth" is highlighted in the Nyquist theorem below to distinguish it from frequency when explaining the concept of under-sampling.

NYQUIST'S THEOREM
 "Any signal can be represented by discrete samples if the sampling frequency is at least twice the **bandwidth** of the signal."


- 2) Traditional Fan-fold printer paper illustrates the location of "Nyquist Zones", which are defined as multiples of half the sampling frequency, F_s . In our XMC example $F_s = 200$ MHz, and $f_s/2 = 100$ MHz, so successive Nyquist zones occur every 100 MHz.




- 3) All signal energy must fall within one Nyquist zone to satisfy both the bandwidth and frequency requirements of the Nyquist Theorem. This example of a wideband signal shown in red crosses multiple zones and violates the "single zone" rule.



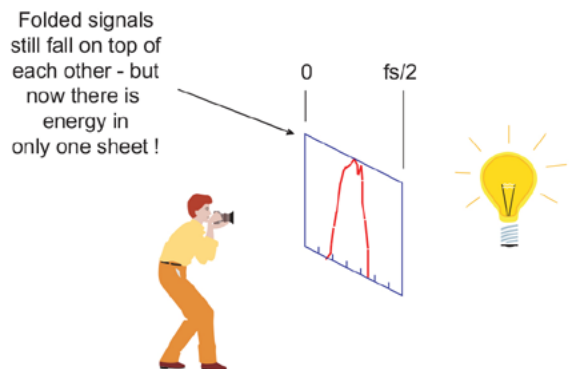
- 4) To illustrate the result of sampling this signal, collapse the fan-fold paper and back light it. All the signal energy above $f_s/2$ is aliased into the first zone. This can be corrected by using a LPF to remove all signal energy above $f_s/2$.



- 5) Another example is a narrow-band signal that falls entirely within Nyquist zone 4 (between 300 MHz to 400 MHz in our case). The signal can be properly sampled using a suitable band-pass filter that eliminates signal energy from all other zones.



- 6) Although the signal frequency is $> F_s/2$, all the energy is contained within one zone, satisfying the Nyquist Theorem. Sampling above Zone 1 is called "under-sampling" [1].



Under-sampling allows an A/D with a lower sample rate and higher dynamic range to capture a narrow bandwidth signal centered at a higher frequency without loss of information. In order for this to work correctly, the RF input path and the A/D must accommodate these higher frequency signals. Our previous 200 MHz A/D example requires an A/C transformer with > 400 MHz passband and adequate Band-pass-filtering (BPF) to reduce noise and additional harmonics from all Nyquist zones excluding the fourth.

After meeting the Nyquist criterion for A/D sampling, the next stage is typically the digital-downconverter (DDC). The DDC is often implemented as IP firmware within the FPGA. It performs frequency translation and bandwidth reduction as described in detail in the next section.

FIRMWARE

An FPGA consists of unconnected logical, arithmetic, and signal processing building blocks that must be configured with firmware (IP) for operation. This is ideal for extreme programming flexibility, but complex because it requires development of the firmware. Some COTS SDR manufacturers provide FPGA IP for basic operation of their board to simplify the development process. This usually includes analog and digital I/O functions for acquiring and transmitting data, along with DSP IP for specific radio functions like DDCs, filters, channelizers and engines to transfer data to the system.

The DDC function requires three IP building blocks: The NCO local oscillator, a complex mixer, and digital filters to replace those functions of the legacy analog radio system (see fig 1a, 1b). Figure 1b is a functional block diagram of the SDR with the DDC. The tuning stage of this DDC uses a complex digital mixer to translate the frequency of interest down to baseband. A pair of multipliers driven by a direct digital synthesizer (DDS) numerically controlled oscillator (NCO) allows the user to "tune" the receiver to the desired frequency. The samples are then passed through a LP FIR filter to decimate the signal for a finite (channel) BW.

Two key benefits of the DDC are higher SNR as a result of decimation, and the ability to tune to the narrow-band signal center frequency. Decimating the signal effectively lowers the sample rate and reduces un-correlated, white noise, referred to as process gain, and the NCO digitally tunes to a specific carrier frequency within a single Nyquist zone.

THEORY SIDE-BAR: SNR AND DIGITAL PROCESS GAIN

The calculation for theoretical SNR of an ADC; $SNR = 6.02 * N(\text{bits}) + 1.76\text{dB}$, or for our 200 MHz 16-bit A/D, $SNR = 6.02 * 16 + 1.76\text{dB}$, or 98dB [2]. Due to A/D non-linearity, additional noise from the LNA input, and sample clock jitter the measured SNR for a good quality A/D is limited to $\approx 75\text{dB}$. This calculation also assumes the input signal is full-scale, which often varies because of the reduced gain, or back-off setting of the front-end LNA to compensate for high crest-factor communication signals. The dynamic range can be improved by reducing the full Nyquist BW down to the channel bandwidth of interest using a DDC. Our 200MHz A/D, with $F_s/2 = 100\text{ MHz}$, and an example 5 MHz LTE channel has the following calculated process gain, or reduction in RMS noise. $SNR_{pg} = SNR_{nyquist} + 10 * \log_{10}(F_s/2 / F_{filter}) = 75\text{dB} + 10 * \log_{10}(100\text{ MHz}/5\text{ MHz}) \approx 88\text{dB}$, or a 13dB improvement.

SOFTWARE

Depending upon the application, the vendor-provided FPGA IP might meet the application specifications, but requires controlling software to operate the radio. The FPGA IP needs operational parameters sent across the system interface from a software program. This is the function of a board support package (BSP) normally written as "C" callable routines for a Windows, or Linux OS environment. The BSP contains library functions and pre-compiled example code that can be executed to test board functionality.

An example case for SDR, is commanding the A/D to capture and transfer data to the FPGA for

further processing in the DDC. This processed data can be stored to memory or transferred to the D/A section for conversion back to an analog signal and output for transmission. This is an example of a software program developed using the BSP software library functions and drivers. If any new FPGA IP is created by the user than additional control software must be written for inclusion in the BSP package.

The Latest COTS SDR Technology

HARDWARE

Over the past decade, FPGA manufacturers like Xilinx have lowered the Size, Weight, and Power (SWaP) values of their chips by reducing the silicon fabrication structure size. In late 2008 the Xilinx Virtex-6 family was constructed using a 40nm process, and had an average 2000 DSP slices per FPGA. By 2017 the Ultrascale family used a 20nm process and the FPGA DSP slices had increased to approximately 5.5K. The latest System-on-Chip (SoC) device from Xilinx, the RFSoc, consists of FPGA fabric with ARM processors, A/Ds, and D/As, all on the same chip. The 16nm technology has over 4.2K DSP slices, four 1.5 GHz A53 ARM processors, two 600 MHz R5 ARM processors, eight 4 GHz, 12-bit A/D's, and eight 6.4 GHz, 14-bit D/As per device. This game-changing technology can be

used by COTS manufacturers to provide multi-channel, SDR transceivers for engineers developing 5G radio products.

Figure 4 is a functional block diagram of one COTS implementation of the Xilinx RFSoc and is the central component of the 5950 3U VPX board from Mercury. The "gray" area is a fully connectorized RFSoc or System-on-Module (SOM) that plugs into the 3U VPX carrier. This device can be controlled via the Gigabit ENET port similar to the previous generation FPGA, but the on-board, ARM processors allow autonomous operation and the ability to communicate with, or control devices locally, or on an external network [3].

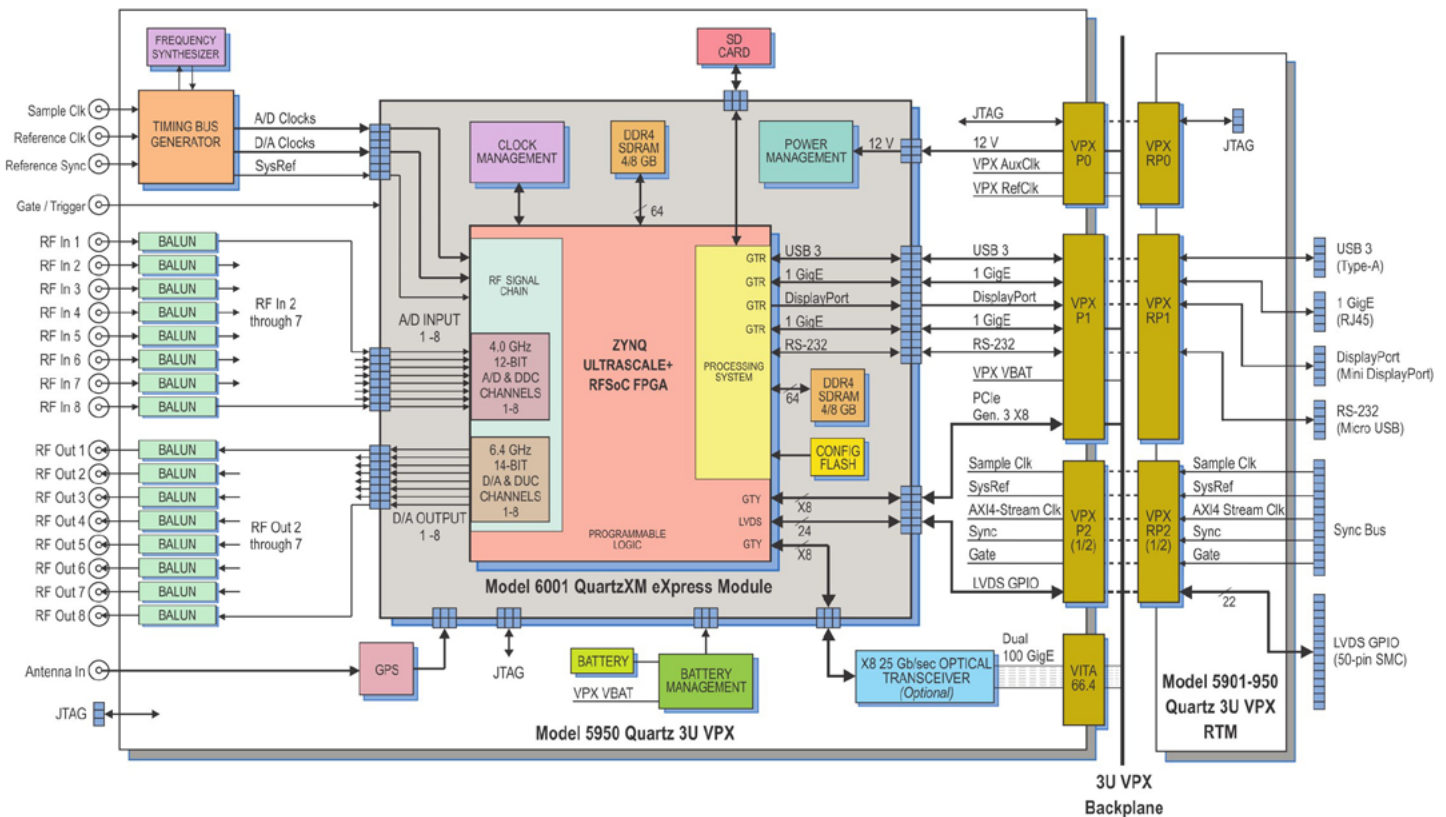


Figure 4

FIRMWARE

Previous generation FPGAs were programmed using a textual hardware description language, or HDL like VeriLog, or VHDL. The latest AXI4 compliant IP blocks are included in Vivado from Xilinx. The IP Integrator tool from Xilinx has virtual graphical blocks that represent HDL code, which can be connected to one another via drag-and-drop wiring. Figure 5 shows VHDL code on the left in contrast to the graphical

block representation on the right. This more intuitive way to program allows someone new to FPGAs to wire together logical blocks that represent hardware like FIR filters, and DDCs to create an SDR. This programming method supports fast integration of vendor supplied, hardware-specific IP blocks with Xilinx IP blocks to create a working SDR system. Both IP block types can be combined to create a common library.

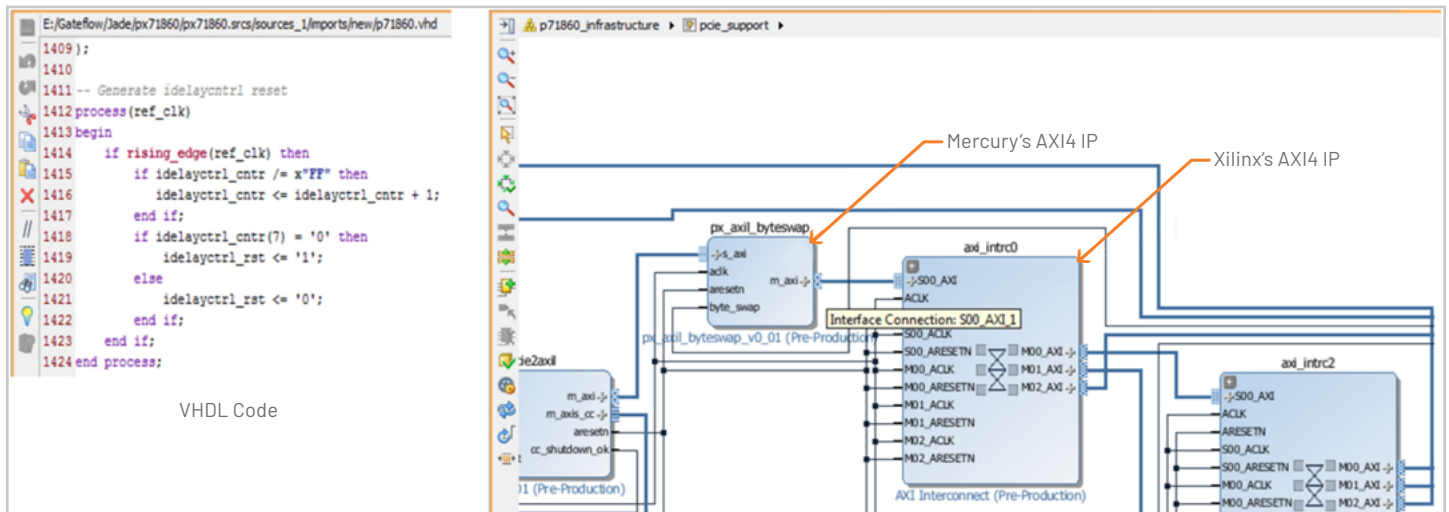


Figure 5
Xilinx Vivado IP Integrator

SOFTWARE

These IP programming advances provided an opportunity for COTS vendors to create a single BSP module that corresponds to one IP module with all the necessary FPGA program parameters in one location. An example would be a "Clock Control BSP Module" that corresponds directly to a "Clock Control IP Module".

5G APPLICATION SPECIFIC EXAMPLE

Figure 6 below illustrates the difference between a Distributed, and Centralized-Radio-Area-Network. The traditional D-RAN “cell sites” were initially being replaced by newer C-RAN’s to improve data transfer efficiency and reduced radio costs, but the latest 5G millimeter wave massive MIMO applications require the separation to move the RRH closer to the end user because of increased RF path loss.




Figure 6
5G Application Specific Example

Figure 7 is a functional block diagram of a Centralized-radio-area-network, or C-RAN consisting of a base-band-unit (BBU), remote-radio-head (RRH), GPS time & frequency reference, and an interconnect module. The selections in “blue” are possible COTS SDR locations. The BBU is located at a central office, or is a virtual network (“the cloud”) with access to multiple optical data lines for back-haul, and the RRH is in an external location closer to the end user. The BBU and the RRH in this front-haul connection example can use a common-public-radio-interface, or CPRI, Open-base-station-architecture-initiative, or OBSAI, or standard Ethernet depending upon system requirements. New front-haul concepts like Extensible-radio-access-networks, or xRAN, and Open-radio-access-networks, or oRAN will be replacing these legacy interfaces in the future.

These various transfer mode options combined with legacy cellular, 5G TF (Verizon specification), or the 3GPP 5G NR (New Radio) can be configured to form a complex heterogeneous network that will require a flexible development platform [4] [6].

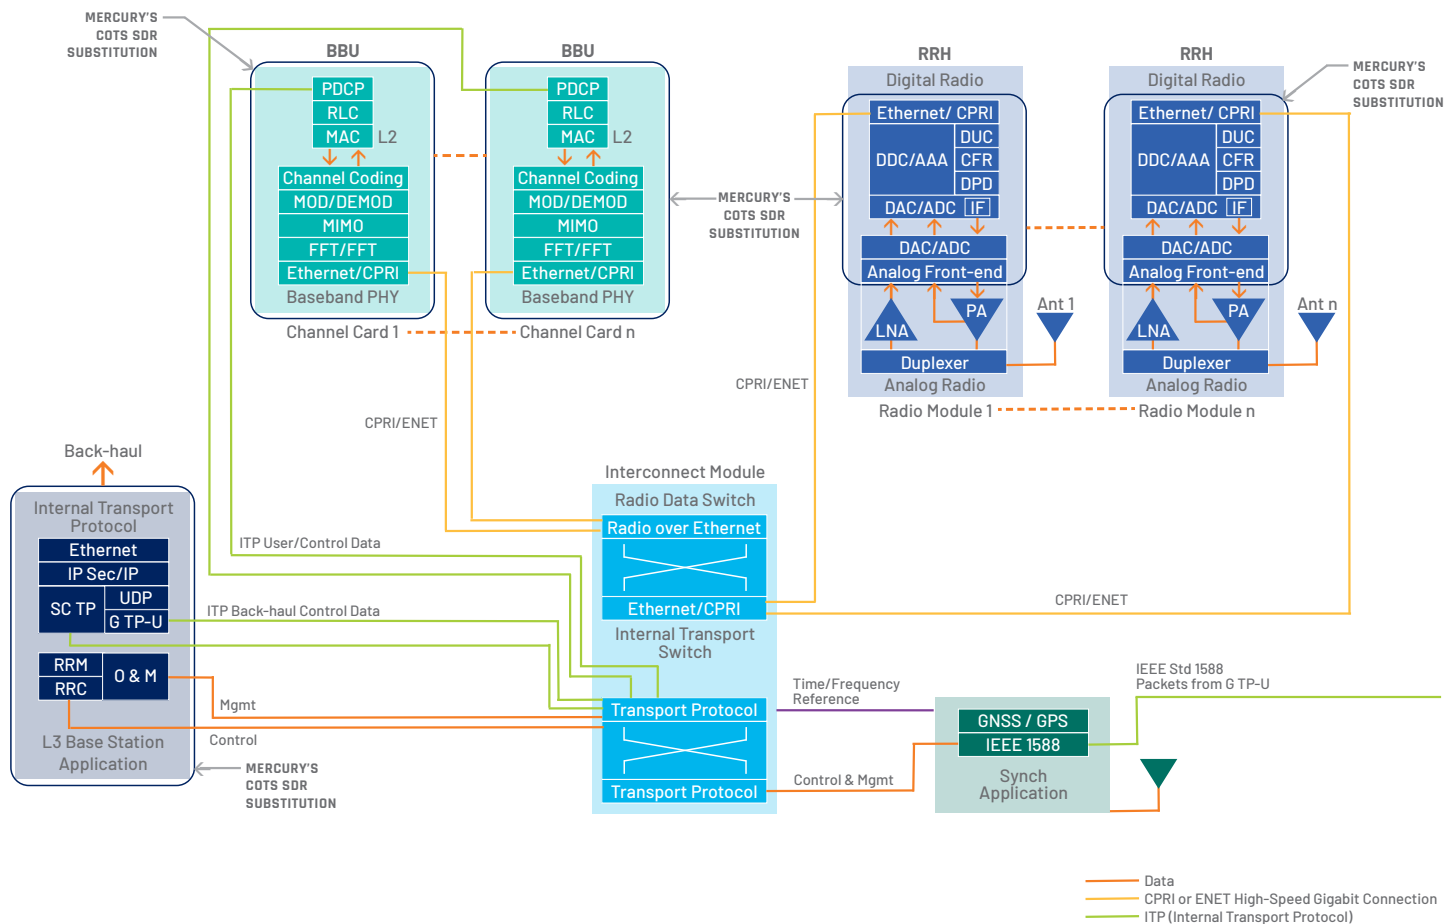


Figure 7
Example C-RAN with COTS SDR Solution

THEORY SIDE-BAR: CELLULAR PHONE CALL

A User Equipment (UE) signal is received by the LNA via the antenna in the RRH. This RF signal is then filtered and adjusted for gain before input to the A/D. The digitized I/Q sample data from the A/D is packetized in the digital radio for front-haul transport to the BBU via a radio data switch. The packetized data is converted into a bit-stream for FFT, MIMO algorithm, demodulation, and channel coding. This data is then managed by the internal transport switch, and re-packetized for back-haul transport to the main cellular network for identification, and further processing. If a phone call is in progress, the user data will be sent out to another RRH using a CPRI, or Ethernet protocol over fiber for OTA transmission to the other party by the reverse process. The previous explanation is a “very” simplified version of this process.

HARDWARE

Figure 8 is an example COTS SDR board used to emulate an RRH in a C-RAN system. A sub-section of the original C-RAN with the RRH is pictured on the left side of the figure and the COTS SDR RRH to the right. The “blue” encircled areas are equivalent. The custom modular carrier card (light green area) contains Rx, and Tx amplifiers, a GPS receiver, and an O/E transceiver module. The SOM (gray area) contains the

RFSoc, and all the connections for power management, data storage, and analog & digital I/O. The incoming RF signal from the antenna is connected to the Rx LNA via a duplexer, isolating it from high power amplifier (PA) transmit levels, and connecting it to one A/D channel. This SOM & custom carrier combination can emulate the original RRH provided it has the necessary IP described in the next section.

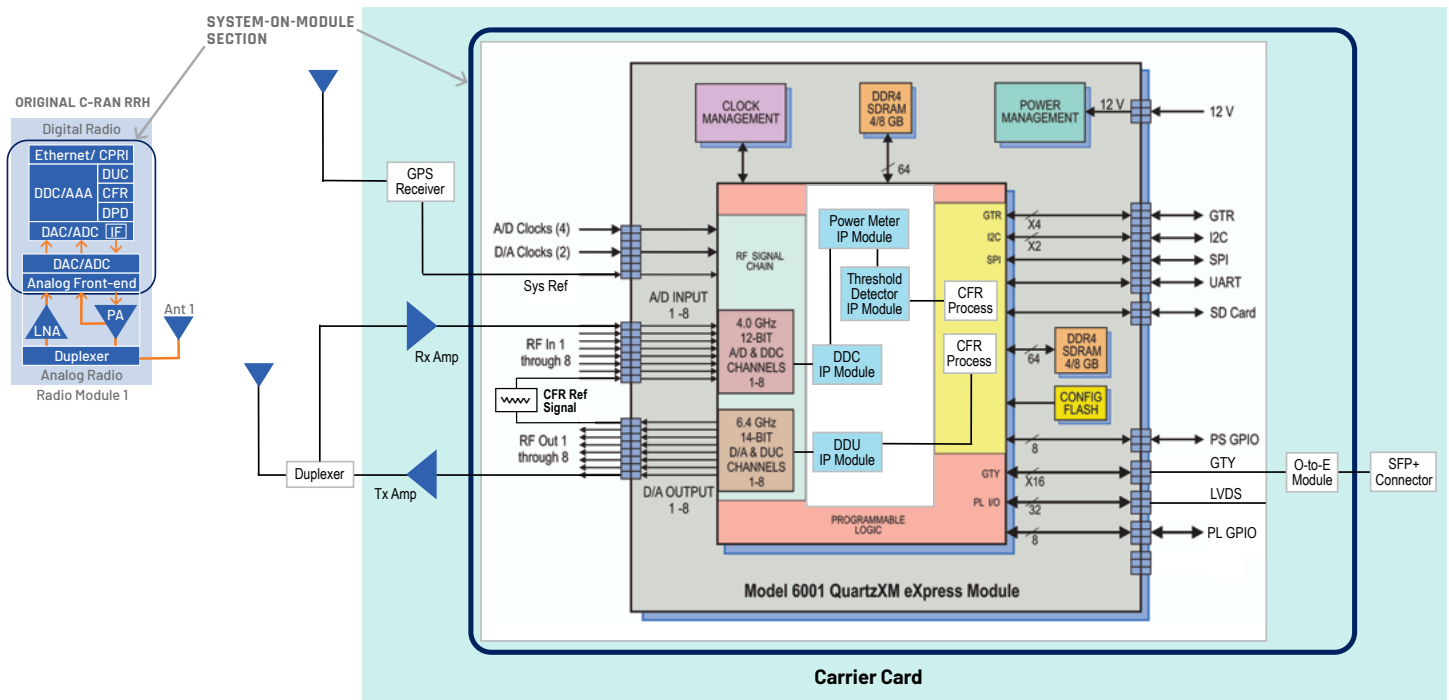


Figure 8
System-on-Module RRH Substitution Example

FIRMWARE

Once inside the FPGA fabric the digital samples are decimated, frequency selected, or tuned, and filtered in the DDC. The DDC output samples can be streamed to the power meter module for measurement, and sorted in the threshold detector IP module. These processed samples can be streamed to the ARM processors for crest factor reduction and digital pre-distortion routines before being up-converted in the DUC

for re-transmission. The DUC is the reverse of the DDC using frequency translation, and interpolation instead of decimation. The digitized I/Q sample data is packetized in the digital radio for transport to the BBU via a radio data switch as in the previous description of a cellular phone call. Because of the variety of channels and data transfer protocols it is necessary to calculate the maximum data through-put of your signal.

Theory Side-Bar: Maximum Data Transport Requirements

Example 1

A remote RRH with two antennas and a 5MHz LTE channel BW will have the following data transfer requirements:

- a. The 5MHz channel requires at least 10MHz sampling, or 10 MSPS to capture the information. There are two bytes per 16-bit sample, and two samples for I&Q.
- b. $SR_{max} = 10 \text{ MS/s} \times 2 \text{ bytes/S} \times 2 = 40 \text{ Mbytes/s}$
 $\times 8 \text{ bits/byte} = 320 \text{ Mb/s}$ per antenna
- c. Two antenna inputs = $320 \text{ Mb/s} \times 2 = 640 \text{ Mb/s}$ of data throughput and no issue for a CPRI port with a 10 to 25 Gb/s capacity

Example 2

A new 5G link with a 100MHz channel and 8 antenna inputs increases the data transfer requirement to $\approx 52 \text{ Gb/s}$ requiring multiple CPRI ports.

(Calculations ignore encoding variations).

SOFTWARE

Depending upon the desired level of control, either BSP routines need to be created for the new IP and ARM processors, or the ARM processors in conjunction with the FPGA can be programmed to operate autonomously.

CONCLUSION

The purpose of this article was to familiarize a traditional radio engineer about the latest hardware, firmware, software, and design tools available from COTS vendors to create an SDR system that can be used for a 5G development platform. These SDR platforms provide superior signal integrity performance, high test repeatability and modular assemblies that adjust to constantly changing 5G design requirements. Future 5G implementations will require many development platforms for experimentation and the use of a COTS system as a starting point will ensure an accelerated time-to-market.

REFERENCE MATERIAL

1. Title: Putting FPGAs to Work in Software Radio Systems (11th Edition)
 - a. Handbook
 - b. Author: Rodger Hosking, Director of Sales, Mercury Systems
2. Title: Taking the Mystery of the Infamous Formula “SNR = 6.02+1.76dB,” and Why You Should Care
 - a. Analog Devices, MT-001 Tutorial, Rev. A 10/08
 - b. Author: Walt Kester
3. Title: Strategies for Deploying Xilinx’s Zynq Ultrascale+ RFSoc
 - a. Author: Robert Sgandurra, Senior Product Manager of Mercury Systems
4. Title: Towards 5G; Xilinx Solutions and Enablers for Next-Generation Wireless Systems
 - a. WP476 (v1.0) June 13, 2016
 - b. Author: Sassan Ahmadi, PhD
5. Title: A Wide Dynamic Range Radar Digitizer
 - a. High Frequency Electronics, September 2008
 - b. Authors: Ed Crean and Paul Hiller of Symtx Inc.
6. Title: 5G New Radio Architecture and Challenges
 - a. CommAgility White Paper
 - b. Author: Dr. Paul Moakes, CTO of CommAgility



Corporate Headquarters

50 Minuteman Road
Andover, MA 01810 USA
+1 978.967.1401 tel
+1 866.627.6951 tel
+1 978.256.3599 fax

International Headquarters
Mercury International

Avenue Eugène-Lance, 38
PO Box 584
CH-1212 Grand-Lancy 1
Geneva, Switzerland
+41 22 884 5100 tel

Learn more

Visit: pentek.com/go/5gswr

For technical questions,
contact dl-sdl-techsales@mrchy.com



The Mercury Systems logo and the following are trademarks or registered trademarks of Mercury Systems, Inc.: Mercury Systems, Innovation That Matters, and BuiltSECURE. Other marks used herein may be trademarks or registered trademarks of their respective holders. Mercury believes this information is accurate as of its publication date and is not responsible for any inadvertent errors. The information contained herein is subject to change without notice.

