
Architecture and Component Selection for SDR Applications

Introduction

In wireless communications, particularly the military space, software-defined radio (SDR) is the goal. The basic concept of SDR is to position the digital-to-analog separation as close as possible to the antenna. This is accomplished by implementing many functions traditionally performed by analog circuits using reconfigurable digital circuits. Successfully implemented, this can deliver an obsolescence-proof radio product which can support many existing and future air-interfaces and modulation formats.

In practice, these objectives can be difficult to achieve, particularly for smaller, battery-powered applications. Requirements related to cost, power consumption, size, and environmental concerns will usually dictate design compromises. Although these requirements vary by application, there is still substantial common ground. In this paper, we will explore an optimized architecture for a typical SDR application.

Generic Requirements

- Low power consumption, particularly while in standby mode
- Reconfigurable baseband, with sufficient processing resources to implement many waveform types, such as frequency shift keying (FSK), quadrature amplitude modulation (QAM), Code Division Multiple Access (CDMA), and orthogonal frequency division multiplexing (OFDM).
- Multiple antennas
- Small form factor
- Video display of information
- Common external interfaces to other equipment, such as Ethernet and USB
- Wide operating temperature range without active cooling
- Wide operating frequency band
- Multiple frequency band support
- Moderate cost
- High manufacturing yield

A critical design parameter, which has a major impact in meeting any set of requirements, is the location of the digital-analog separation in the radio chain, both for receive and transmit. The SDR methodology dictates that this conversion should take place as close to the antenna as possible.

Recently, intermediate frequency (IF)-subsampling analog-to-digital converters (ADCs) operating at over 100 MSPS at 12+ bits with analog input bandwidths up to 1 GHz have become available. The current offerings by leading vendors require as little as 200 mW of power consumption per channel.

This makes IF frequencies above 500 MHz practical. It still will not allow direct sampling of a 2-GHz signal, but it will allow for a single conversion radio to sample multiple signals over a wide bandwidth. Increasing the sampling rate also increases the effective signal-to-noise ratio (SNR) as the signal of interest is decimated to baseband processing rates, improving the sensitivity and dynamic range of the radio receiver.

Interfacing at these levels of IF frequencies is a bit more challenging on the transmit side. Most baseband circuits have difficulty producing transmit samples at rates much higher than 200 MHz, which limits the transmit IF to about 40 percent of the rate, or about 80 MHz. But some digital-to-analog converter (DAC) vendors have addressed this issue by incorporating interpolation into their parts, and adding circuitry to allow the DAC to selectively output in higher Nyquist zones (thereby eliminating the 40 percent requirement). Another option is to integrate interpolation, digital mixer, and the DAC. This also can permit high transmit IF frequencies, on the order of 400 MHz. However, these approaches can be expensive, both in terms of cost and power consumption. If lower transmit IFs are acceptable, then a much less costly and lower power solution can be implemented.

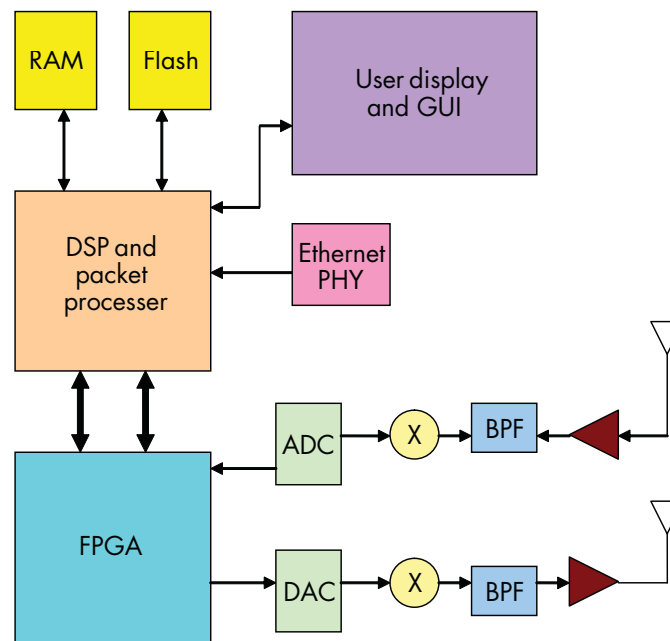
For baseband processing, there is a plethora of possible solutions and architectures. For SDR applications, one of the key requirements is configurability and flexibility. This eliminates the lowest cost and power consumption solution, which is purpose-built baseband ASICs such as those used in commercial cell phones. The most flexible digital signal processing (DSP) solution is to use both a DSP processor and a FPGA. This approach has the following advantages:

- The DSP processor provides the ideal platform for implementation of very complex algorithms
- The FPGA provides very high processing rates and bit resolution required for certain functions such as finite impulse response (FIR) filters, fast Fourier transforms (FFTs), CDMA RAKE receiver, Turbo decoders, and numerous other functions
- FPGAs can act as an effective co-processor for DSP processors, provided suitable high-bandwidth connections are available; this approach allows algorithms to be split across DSP devices and FPGAs in an optimal fashion
- Both DSP and FPGA vendors offer extensive vendor or third-party intellectual property (IP) for many complex but standardized functions
- Mature and well known development tool environments are available
- Can later migrate latest devices within DSP/FPGA product families, which will continue to add capability while preserving initial investment of in-house development, IP, tools, and design methodology
- DSP processors can incorporate protocol and packet processing functions, either with dual processors, or by using a single convergent processor architecture
- DSP processors can incorporate the BT656 video interface to the user display

The choice of vendor and product family will have a major impact on power consumption, cost, radio data rates, and system upgradeability. These choices, once made, are difficult to change later without a major engineering redesign cycle.

Figure 1 shows a generic block diagram of a typical SDR radio implementation. The rest of this paper discusses critical components that have a major impact on radio performance and capability, and provides comparisons between different vendors' product offerings. The key components considered are the ADC, DAC, DSP, and FPGA devices. Cost is not part of this comparison for two reasons. First, pricing numbers are often difficult to come by, and are always subject to negotiation. But more importantly, due to market pressures, manufacturers are usually willing to match their competitor's prices for comparable devices, making any cost advantage difficult to establish and quantify.

Figure 1. High-Level SDR Block Diagram



ADC Selection

The receive technique of IF sub-sampling is commonly used in SDR designs. This can allow for very high receive IF. In theory, the receive IF is only limited by the analog bandwidth of the ADC. In practice, a very high receive IF can degrade ADC performance, and also impose tight requirements on the jitter of the ADC clock. For these reasons, it is important to compare both the analog bandwidth of the ADC as well as the maximum characterized receive signal in the manufacturer data sheets. ADC performance at higher IF rates than characterized in the data sheet should be verified with the device manufacturer. In [Table 1](#), only ADCs with > 100 million samples per second (MSPS) capability and < 500 mW per channel were considered. A 12-bit resolution is usually sufficient to meet the requirements of most modern radio requirements, and 12- to 14-bit ADCs were considered.

Dual ADCs are attractive, as most SDR receivers will require diversity.

Table 1. Comparison of Competing ADCs

ADC Manufacturer	Part Number	Max F_{CONV}	Full analog BW	Bits	Power at Max F_{CONV}	Number of Channels	SNR at Max Characterized IF Frequency
Analog Devices, Inc.	AD9254	150 MHz	650 MHz	14	430 mW	Single	71 dB at 170 MHz IF
Analog Devices, Inc.	AD9230	250 MHz	700 MHz	12	430 mW (1)	Single	65 dB at 250 MHz IF
National Semiconductor	ADC14DS105	105 MHz	1 GHz	14	450 mW	Dual	72 dB at 240 MHz IF
National Semiconductor	ADC12DS105	105 MHz	1 GHz	12	450 mW	Dual	68 dB at 240 MHz IF
Texas Instruments	ADS6245	125 MHz	500 MHz	14	1 W	Dual	69 dB at 230 MHz
Texas Instruments	ADS6225	125 MHz	500 MHz	12	1 W	Dual	68 dB at 230 MHz

Note:

- (1) AD9230 is sampling at much higher rate than the other parts in the table. ADC power consumption generally increases linearly with ADC conversion rates, so it likely requires ~200 mW at around 100 MSPS. This part is a good choice if a much higher sampling rate is required, or if a single ADC per package is desired.

The National Semiconductor ADCs offer very high analog bandwidths with low power consumption (~200 mW per channel). For nearly all SDR applications, 100 MSPS is an adequate sampling rate.

DAC Selection

The power and cost penalties make a very high transmit IF unattractive, unless there is a compelling reason for this approach. Since a 400-MHz transmit IF is still far too low for most frequency bands, an analog upconversion stage will be required in nearly all cases. For typical SDR applications where the signal bandwidth is less than 25 MHz, a sub 100-MHz transmit IF frequency coupled with a single mixer stage to RF is more practical. [Table 2](#) summarizes the comparison results among several competitive DACs. A 14-bit resolution is usually sufficient to meet the performance requirements of modern radio modulations.

Table 2. Comparison of Competing DACs

DAC manufacturer	Part number	Max F_{CONV}	Max IF	Bits	Power at Max F_{CONV}	Number of Channels	Comments
Analog Devices, Inc.	AD9778	1 GHz	480 MHz	14	~1 W	Single	Integrates up to 8x interpolation with Nyquist Zone selection, requires external analog quadrature modulator
Analog Devices, Inc.	AD9957	1 GHz	400 MHz	14	1.3 W	Single	Integrates up to 255x interpolation, complex numerically controlled oscillator (NCO)
Analog Devices, Inc.	AD9744	210 MHz	85 MHz	14	150 mW	Single	Low-cost, pin-compatible family with 8,10,12,14 bits
Analog Devices, Inc.	AD9755	300 MHz	140 MHz	14	155 mW	Dual	Pin-compatible, 10, 12-bit version available
Texas Instruments	DAC5672	275 MHz	110 MHz	14	330 mW	Dual	Low-cost, pin-compatible, 12-bit version available

For a single transmit chain, the AD9744 is a good choice. However, based upon common requirements of many modern communication systems for diversity transmit, the AD9755, offering over 100-MHz transmit IF, may be a better choice. The interpolation filters and complex NCO functions can be easily and efficiently implemented in the FPGA.

DSP Selection

As mentioned above, the choice of DSP processor is the one to get right the first time, as it is difficult to change later. The following list of requirements was used to narrow the field to two vendor families.

- Established product family, well received in marketplace, expected to be well supported and increase in capability in future
- Good options for operating system: vendor supplied, Linux, third party
- Able to support packet processing tasks (media access control (MAC) layer)
- Able to support radio upper layer protocol layers
- Mature development environment
- 500-MHz core performance
- Low-power architecture (less than 1 W)
- Support for Ethernet and/or high-speed USB
- External DDR (or SDRAM) memory interface
- High-bandwidth interfaces suitable to interface to FPGA
- Able to drive video display interface

After review of several vendors, two DSP product families stand out: Blackfin from Analog Devices and DaVinci from Texas Instruments. As shown in Table 3, both of these DSP families meet all of the requirements discussed above. In addition, each vendor has many family members, with different peripherals and features. Both families have major backing from their manufacturers as well as from third-party software IP developers. Both have large-scale market acceptance, and will be improving features and performance for years to come.

Table 3. DSP Comparison Table

Requirements	Analog Devices	Texas Instruments	Comments
Part Number	ADSP-BF527	TMS320DM6441	
Maximum DSP core clock Maximum RISC core clock	600 MHz	513 MHz 257 MHz	Other DaVinci parts have 600 MHz, but are higher power or don't have ARM RISC core
Embedded L1 RAM	132 Kbytes	112 Kbytes + 40K for RISC	Both devices also have additional L2 RAM
16-bit giga multiply-accumulate operations per second (GMACS)	1.2	2.05	Compare to FPGA GMACS
Core power consumption	270 mW at 600 MHz 120 mW at 400 MHz	800 mW at 513 MHz 550 mW at 405 MHz	Major Analog Devices advantage (using BF537 data)
Integrated ARM RISC core	No	Yes	Major Texas Instruments advantage
Flash controller	Yes	Yes	
Video BT656 IF/video accelerator	Yes/No	Yes/Yes	
Ethernet support	10/100 MAC	10/100 MAC	
USB support	Yes	Yes	
OS support	Yes	Yes	Vendor-supplied BIOS, Linux, third party
Integrated development environment (IDE)	Visual DSP	Code Composer	
Host port interface	16 bit	16 bit	Can be used by FPGA to access DSP internal memory
On-time programmable memory section	Yes	No	Can be used for security features
External memory interface	16-bit SDRAM	32-bit DDR	
Serial ports	2	2	
Serial peripheral interface (SPI) I/F	Yes	Yes	
On-the-fly phase-locked loop (PLL) adjustments	Yes	Yes	Dynamically reduce clock to save power
Package	208, 289 BGA	361 BGA	

Both Blackfin and DaVinci have many family members. This makes a selection difficult, as different applications may emphasize different processor features.

For Blackfin, the requirement of Ethernet MAC narrows the field considerably. The best fit seems to be BF527. If USB 2.0 can be substituted for the Ethernet MAC, then consider the BF548, as this part has an extra parallel port, which is very useful for high-bandwidth connection to an FPGA. The BF548 also packs improved video processing instructions.

For DaVinci, only family members with integrated RISC core were considered, such as the 6441. This is a major advantage, as it allows packet processing software to reside in a separate CPU core. In contrast, Analog Devices uses a “convergent” architecture, such that its Blackfin core is designed to simultaneously support packet processing and DSP applications.

To summarize, the Blackfin has a large advantage in power consumption, size, and possibly cost. The DaVinci, with its DSP and RISC core, has more processing power. On DaVinci, separate cores for packet processing/protocol and DSP can ease software development. While both Analog Devices and Texas Instruments processors are designed to be media processors, the DaVinci has more support for video processing.

For SDR handsets, which require much lower power consumption, the Blackfin is probably the better choice. For vehicle-mounted systems, the DaVinci offers more processing power and would have the advantage.

When interfacing a DSP processor to an FPGA, high bandwidth and multiple interfaces can make a huge difference in performance. Many radio algorithms are best implemented split across the DSP processor and the FPGA. The algorithmically complex but lower rate-processing tasks can be implemented in the DSP. Algorithmically simpler

tasks with processing requirements should be implemented in the FPGA. This requires flexible and high-bandwidth interfaces.

A common example is initial acquisition of the receive signal. Often an iterative approach is taken to find the frequency offset, timing recovery, and synchronization simultaneously. Usually the algorithms employed are complex, and the best approach is often to split across the DSP processor and the FPGA, with the FPGA doing most of the processing while the DSP calculates the algorithm parameters. The inter-device communication requires the ability to transfer blocks of data with low latency using direct memory access (DMA) to and from DSP memory, as well as rapid DSP access to control registers in the FPGA. Interrupts, often DMA triggered, are used to synchronize processing between the DSP processor and the FPGA.

FPGA Selection

The FPGA represents another important architectural component choice. FPGAs have evolved beyond just programmable hardware, as the FPGA manufacturers offer extensive IP cores, microcontroller soft cores, reference designs, system-on-chip design methodology, and comprehensive design environments.

Due to power and cost constraints, only the Altera® Cyclone® III FPGA and the Xilinx Spartan-3 device families were considered. Both manufacturers offer extensive DSP capability in these families, dwarfing the processing rate achievable in a DSP processor. The digital up and down conversion chains, pulse shape filtering, CDMA bit level processing, OFDM FFT processing, Transmit Crest Factor Reduction, and Turbo, Viterbi, and Reed Solomon decoding are all very good candidates for FPGA implementation. Both manufacturers offer IP cores and reference designs for most if not all of these functions.

FPGAs also come in various densities, which are often pin compatible. This allows larger devices to be used, and allows for more processing margin for future radio protocols.

Table 4 shows a comparison of two recently announced Xilinx Spartan-3A DSP devices to the Altera Cyclone III family. These two Xilinx devices have been optimized for DSP applications, as has the Altera Cyclone III family. Not all Cyclone III family devices are shown.

Table 4. Comparison of Spartan-3A and Cyclone III Devices

FPGA Device	Xilinx XC3SD1800A	Xilinx XC3SD3400A	Altera EP3C25	Altera EP3C40	Altera EP3C55	Altera EP3C120
Logic elements (LEs)	37.4K	53.7K	24.6K	39.6K	55.9K	119.1K
18 x 18 multipliers	84	126	66	126	156	288
9 x 9 multipliers	84	126	132	252	312	576
Block RAM	1510 Kbits	2270 Kbits	610 Kbits	1160 Kbits	2400 Kbits	4000 Kbits
Max MAC rate	250 MHz	250 MHz	260 MHz	260 MHz	260 MHz	233 MHz
Total GMACS (18 x 18)	21	31.5	17.1	32.8	40.6	67.1
Max diff I/O	227	213	83	227	163	233
Technology	90 nm	90 nm	65 nm	65 nm	65 nm	65 nm
Dynamic power at 100 MHz	No data available	No data available	0.48 W	0.73 W	1.0 W	1.8 W
Static power at 25°C	198 mW	272 mW	60 mW	63 mW	65 mW	70 mW

Multiplier resources include only hard multipliers, not soft multipliers built from logic resources. Similarly, only block RAM was considered, as distributed RAM consumes LEs, which are then unavailable for logic use. Multiplier counts show configuration with a maximum of 18 x 18 or 9 x 9 multipliers.

While FPGAs have large DSP capabilities (compare GMACS in table above with DSP values), they can also have high power requirements. So a critical parameter to consider is power consumption. This is not always an easy number to evaluate (the Xilinx data sheet does not give any dynamic power number). Therefore, both manufacturers

have power estimators, which can be used early in the design cycle. An easy comparative number to use is the value for static power (device not clocked or programmed), which while not indicative of the FPGA's actual power usage, can be used for comparative purposes. The Altera parts are approximately one-third the static power requirement of the Xilinx devices.


Even though Xilinx has targeted these two parts for DSP applications, they fall short of the Altera Cyclone III FPGA capabilities. The recommended choice here is the Altera EP3C40, in the 484-pin BGA package. As the table indicates, the Cyclone III FPGA offers much more DSP capability than similarly sized Spartan-3A DSP parts. The Cyclone III EP3C40 device has much lower power consumption and a pin-compatible migration path to 3 larger parts, up to 120K LEs, 288 multipliers, and 4 Mbits RAM. For designs where power consumption is critical, Altera Cyclone III parts are available in 5K and 10K LE densities, which still offer 6-GMACS DSP performance, with power consumption on the order of 200 mW at 100-MHz clock rates.

Summary

The information presented here should provide a useful starting point for SDR design engineers in their efforts to design a very complex product with many conflicting requirements and features. Designing battery-powered SDR products will continue to be a challenge. These products cannot be expected to approach the power consumption levels of consumer products, which are highly specialized and have short life cycles. But with careful attention, and using some of the latest semiconductor product offerings, a reasonable talk and standby time is achievable in moderate-size form factors.

Acknowledgement

- Michael Parker, Sr. Technical Marketing Manager, IP and Technology Product Marketing, Altera Corporation



101 Innovation Drive
San Jose, CA 95134
www.altera.com

Copyright © 2007 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries. All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, maskwork rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera Corporation. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.