

NEW FPGAs REVOLUTIONIZE DIGITAL DOWN CONVERTERS

Rodger Hosking (Pentek, Inc., Upper Saddle River, New Jersey, USA,
rodger@pentek.com)

ABSTRACT

Digital Down Converters (DDCs) represent a cornerstone technology in SDR communication system. Over the past several years, the tuning, data reduction and filtering functions associated with DDCs have shifted from ASIC implementations towards IP cores in FPGAs. This shift brings many critical advantages including architectural flexibility, higher precision processing, higher channel density, and lower power and cost per channel.

With the advent of each new higher performance FPGA family, these benefits grow. This article, through practical examples, explores some to the key advantages of implementing DDC designs in FPGAs and describes some of the situations when ASICs can still offer the best solution.

1. INTRODUCTION

In the 1980s, when the first programmable DSP chips began appearing on the market, engineers quickly began designing boards and systems around these parts to replace the discrete registers, adders, state machines and logic of traditional digital signal processing hardware.

About ten years later, programmable RISC processors for workstations like the Intel i860 and Motorola PowerPC were found to perform quite well on DSP algorithms. Even though these chips were not aimed at the DSP embedded computing market, because their built-in floating-point ALUs and hardware multipliers, they were recruited as alternatives to native DSP chips and began appearing on embedded system boards.

Then at the turn of the century, the relatively sophisticated FPGA (field programmable gate array) emerged to replace more primitive programmable devices used for glue logic, state machines, and interface engines. With a reasonable density of resources, FPGAs could now implement the ALUs, multipliers and logic functions essential for DSP. Once FPGA vendors became aware of this trend, they invested heavily in boosting DSP capabilities. As a result, FPGAs can now easily outperform DSP and RISC processors in many key benchmark algorithms.

But this performance comes at a price: FPGA algorithm development requires a completely different skill set and is

much more challenging than programming a DSP or RISC processor. To help overcome this, optimized IP cores with fully characterized performance are now available from many sources.

First, a review of current FPGA technology is in order, followed by a few IP core examples including design strategy, tradeoffs and benefits.

2. FPGAs: HORSEPOWER FOR SDR

Without exception, the latest device offerings from major FPGA vendors offer third generation DSP blocks. They include extended precision multiplier/accumulators, advanced arithmetic units, logic engines, and flexible memory structures that can be tailored into block memory, dual-port RAM, FIFO memory and shift registers.

Competition among FPGA vendors is stronger than ever, leading to an exciting race for features that deliver maximum performance and specific benefits. Winning this race, however, is a complex and elusive goal. With so many different types of resources including block RAM, distributed RAM, DSP blocks, logic blocks, micro-controllers, gigabit ports, I/O drivers and pins, etc., determining a single optimum ratio is futile because each application requires a different blend.

For example, the design engineer selecting the best part for a logic-intensive application will avoid an FPGA heavily burdened in cost and power with a wealth of powerful DSP blocks. As a compromise, vendors have developed multi-pronged product offerings, each targeting different classes of applications.

Two examples are the Xilinx Virtex-4 and Virtex-5 FPGAs with characteristics shown in the table in Figure 1. Unlike earlier device families, Xilinx introduced FPGA sub-families, each one emphasizing distinct strengths. For the recently announced Virtex-5 family, there are a total of four distinct sub-families, all using a 65 nm process that reduces core voltage down to 1.0 volt. This allows an improvement in maximum clock speed to 500 and 550 MHz, respectively, while reducing power consumption.

Logic cells are the basic elements used for implementing state machines, combinatorial logic, controllers, and sequential circuits. They are composed of logic "slices" with flip-flops, look-up-tables (LUTs), multiplexers, Boolean logic blocks, and adder/subtractors

with carry-look-ahead functions. The Virtex-5 uses 6-input LUTs instead of the 4-input LUTs in the Virtex-4, providing additional logic functions, fewer levels of logic for faster speed and less power due to simpler routing.

	V-4 SX SX55	V-4 FX FX100	V-5 LX LX330	V-5 LXT LX330T	V-5 SXT SX240T	V-5 FXT FX200T
Logic Cells	55,296	94,896	331,776	331,776	239,616	196,608
Block RAM (bits)	5,760k	6,768k	10,368	11,664	18,576	16,416
Max I/O Pins	640	768	1,200	960	960	960
DSP Multipliers	512	160	192	192	1,056	384
Power PC Cores	-	2	-	-	-	2
3 GHz Serial Ports	-	20	-	24	24	-
6 GHz Serial Ports	-	-	-	-	-	24
Gbit ENET MACs	-	4	-	4	4	8
PCIe End Points	-	-	-	1	1	1

Figure 1. Evolution of FPGA Resources for SDR

Another essential resource for DSP is memory, which has become much more flexible in these latest generation FPGAs and comes in different forms. Distributed memory is used for LUTs, FIFOs, single- and dual-port RAMs, and shift registers. For larger memory structures, 18-kilobit block RAMs can be used for deep FIFOs, large circular delay memory buffers, deep caches, as well as bigger single- and dual port RAMs. The Virtex-5 offers both 18- and 36-kilobit block RAMs to support wider memory structures of up to 72 bits within a single block.

One of the more significant advances in the Virtex-4 family was the new XtremeDSP slice. Following the market demand for more powerful signal processing structures, Xilinx surrounded the popular 18x18 hardware multipliers first introduced in the Virtex-II series with a 48-bit adder/subtractor capable of acting as a registered accumulator. Due to tight, dedicated logic, this facility can operate at clock speeds up to 500 MHz and can propagate the results between XtremeDSP slices with 48-bit precision at the same rate.

The 48-bit path allows this fast, fixed-point hardware to rival the precision of floating-point engines by preserving the 36-bit multiplier outputs with plenty of overhead for bit growth as results propagate through cascaded slices.

Each XtremeDSP slice features 40 dynamically controlled logical and arithmetic modes and supports mode changes during runtime without the need to recompile the FPGA. In this way, each XtremeDSP slice behaves like a miniature DSP processor, and there are as many as 512 of these in a single FPGA.

With so much demand for DSP capability, the Virtex-5 family introduced the DSP48E slice that boosts 18x18

hardware multiplier to 18x25 so that single-precision floating-point arithmetic can be implemented within two DSP48E slices instead of the four slices required with the Virtex-4 XtremeDSP. Also, the adder has been enhanced with a logic stage to save the need for an external logic block.

The four Virtex-5 sub-families include the LX, LXT, SXT and FXT. The LX devices offer maximum resources for logic intensive applications, with the LX330 providing over 331,776 logic cells, far more than any predecessor, but offer no gigabit serial ports. This important resource has been added to next three sub-families, denoted by the “T” suffix. The LXT maintains the logic resources, sacrifices some of the block RAM, adding not only the gigabit serial ports but also Ethernet MAC engines, and a new resource for Xilinx FPGAs, the PCI Express endpoint engine.

The SXT FPGAs deliver maximum resources for DSP, with 1,056 DSP48E engines in the largest device. The FXT sub-family adds useful resources for embedded computing applications, including IBM 440 microcontroller cores and additional Ethernet MACs. One major new resource, unique to the FXT devices, are the 6.5 GHz GTX gigabit serial ports, twice as fast as other Xilinx device families.

To take advantage of the tremendous DSP horsepower available, some example IP cores for key DSP algorithms are now presented.

2. DIGITAL DOWN CONVERTERS

DDCs (digital down converters and often called digital receivers) perform the two essential software radio functions: frequency translation and channel filtering. In a basic narrowband CIC filter DDC shown in Figure 2, a mixer and local oscillator perform the frequency translation.

The local oscillator consists of a digital phase accumulator that advances each clock by a programmable increment equal to the tuning frequency. The phase accumulator is a register whose full-scale value represents 360 degrees of a sinusoid. A sine/cosine lookup table converts the phase angle of the accumulator to the digital voltage value of the sinusoid.

The mixer consists of two digital multipliers that accept complex sine/cosine outputs from the local oscillator with digital samples of the receiver input signal produced by an A/D converter. Multiplication in the time domain produces a sum and difference signal in the frequency domain. If the local oscillator is set to the frequency of the input signal of

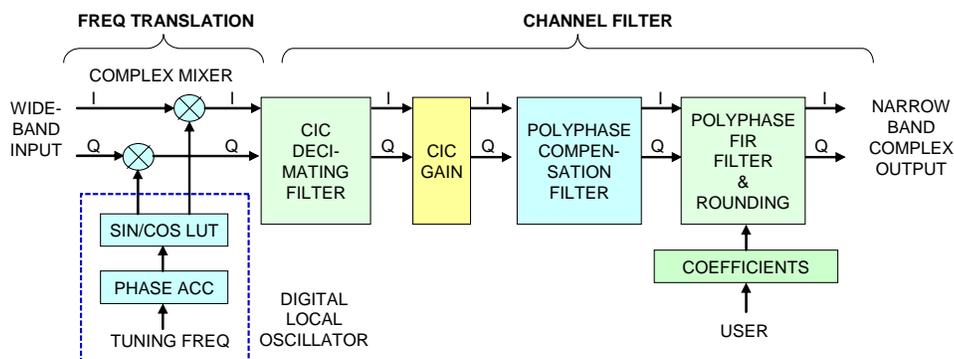


Figure 2. Typical Single Channel Narrowband DDC IP Core

interest, the difference term will be that input signal translated down to 0 Hz.

The CIC decimating filter does not require multipliers and efficiently achieves high orders of bandlimiting and decimation at the expense of pass band flatness. The compensation filter (CFIR) corrects the pass band response with a polyphase FIR design. A second FIR with user programmable coefficients (PFIR) to establish channel frequency characteristics follows this. CIC-type DDCs are available both as ASIC chips and as FPGA IP (intellectual property) cores. Commercial ASICs feature as many as four DDC channels per chip, like the popular Texas Instruments/Graychip GC4016.

IP core DDCs, like the LogiCore DDC from Xilinx can be scaled for various levels of SFDR (spurious free dynamic range) performance to use more or less of the available resources. For example, a complex DDC with 84 dB SFDR consumes approximately 1,700 slices. In a mid-sized FPGA, such as the Virtex-4 SX55 with 34,000 available slices, only about 20 of these DDCs can be accommodated.

For applications requiring several dozen or even hundreds of channels, this approach can become impractical. This is because each conventional DDC requires its own local oscillator (phase accumulator and sine table) and complex mixer (two multipliers) and these blocks must operate at the full input sample clock rate. Clock rates for A/Ds commonly used in software radios range between 100 and 200 MHz. Since this is the same clock range rating for the multipliers inside the DSP blocks, the critical hardware resources required for one channel cannot be shared across other channels.

However, imagine that the input data sample rate is reduced by a factor N. By operating the critical DDC hardware resources required for one channel at the full clock rate, those same resources can then be multiplexed (time shared) across N channels. Of course, provisions must be made for buffering the data for all channels while multiplexing. This is usually done in RAM or in delay memory, a common feature in FPGAs.

One way to achieve this input rate reduction is to split the input signal into a bank of N adjacent frequency bands using a channelizer. Then, the output sample rate for each band can be reduced by a factor of N. The output from the band containing the signal of interest can be selected as the input to any given DDC to fine tune within that band.

For this strategy to work, of course, the channelizer design must use fewer hardware resources than the DDC structures it replaces.

2. SINGLE INPUT 256 CHANNEL DDC

Figure 3 shows an FPGA-based 256 channel DDC IP core that combines a proprietary channelizer stage followed by a multiplexed DDC stage. The channelizer accepts a single wideband input stream and delivers a channel bank of 1024 output bands equally spaced in frequency, but with significant overlap between adjacent bands. Because of the reduced output bandwidth of each filter, the output sample rate of each band is the input sample rate (F_s) divided by 256.

Inside this DDC Core 430, a crossbar switch matrix accepts 1024 inputs from the channelizer and delivers 256 outputs, one for each DDC channel. The switch is non-blocking so that any of the 256 outputs can be independently sourced from any of the 1024 channelizer bands with no restrictions.

Channel tuning is performed with a separate 32-bit frequency word for each of the 256 channels. The most significant bits are sent to the switch matrix for coarse tuning to select the correct channel band for each channel. The least significant bits of the frequency word are used by the DDC stage for fine tuning within the selected band.

Since the channelizer outputs exhibit frequency droop at the band edges, a fixed FIR filter flattens the pass band to within 1 dB across a span equal to twice the band-to-band spacing.

A total of 256 compensated switch matrix outputs feed a bank of 256 independently tuned DDC sections, each with its own local oscillator mixer and FIR filter. Because the channelizer has dramatically reduced the input sampling rate to each DDC section by a factor of 256, the DDCs are implemented using highly multiplexed hardware resources and block RAM to preserve the data for each channel. A gain stage, output multiplexer and data formatter complete the design.

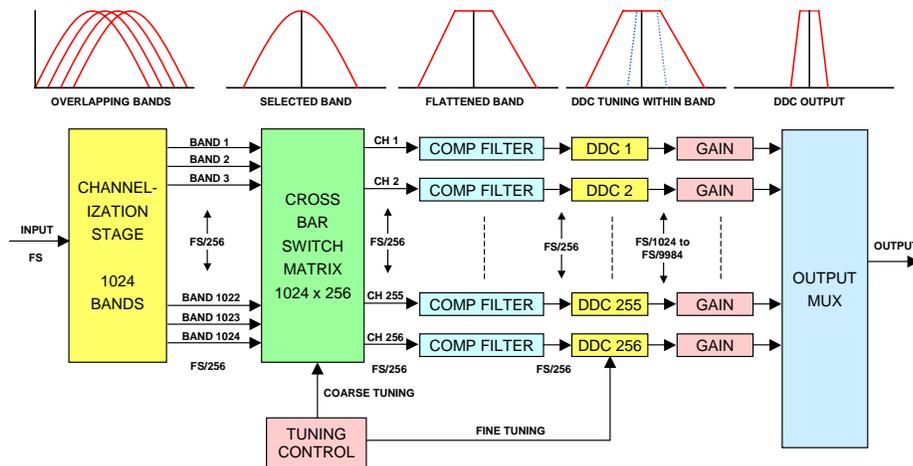


Figure 3. 256 Channel Digital Down Converter IP Core 430

The channelizer band-to-band spacing of $F_s/1024$ equals the maximum output bandwidth of this design. For an input sample rate of 100 MHz, this spacing is about 100 kHz. Thanks to its broadened response, each channelizer output has a clean pass band equal to twice the band spacing, or about 200 MHz.

This allows the DDC to perform fine tuning by sliding its local oscillator frequency across the selected 200 kHz channelizer band ± 100 kHz to precisely center the DDC output. Choosing a wider DDC output bandwidth would restrict the DDC tuning range, since the edge of that wider bandwidth would cross the edge of the flat, spurious-free region of the channelizer output.

Translated signal samples from the mixer feed the decimating FIR at the channelizer output sample rate of $F_s/256$. Since the maximum available DDC output bandwidth is $F_s/1024$, the lowest decimation factor allowed in the FIR is 4.

For narrower output bandwidths, the maximum decimation factor is determined by the complexity (number of taps) of the FIR filter, which must perform at least as well as the channelizer filter in order to maintain that dynamic range of 75 dB. Choosing a reasonable number of multiplier/accumulator stages yields an FIR filter suitable for decimation factors from 4 to 39 in steps of 1.

Because the stages are cascaded, the channelizer decimation factor (256) and FIR filter decimation factor (4 to 39) multiply. Thus, the overall range of decimation for the entire core is 1024 to 9984 in steps of 256. Each of these 36 available decimation factors requires its own set of filter coefficients, which are stored in a table within the FPGA. For an input clock of $F_s = 100$ MHz, the range of output bandwidths using the default 80% filter characteristic is approximately 8 kHz to 80 kHz.

Because of the multiplexed DDC hardware, all 256 channels must have the same decimation factor setting. For high channel count systems, this limitation is usually not an issue since it is quite common for all such channels to have the same bandwidth.

Overall performance of the complete 256-channel FPGA-based DDC IP Core 430 includes a spurious free dynamic range of 75 dB, a pass band ripple of 0.4 dB, a pass band edge droop of 1.0 dB, and frequency tuning resolution of $F_s/2^{32}$. The maximum clock

frequency depends on implementation details, but can be as high as 185 MHz in a Virtex-4 FPGA with speed grade 12.

The core consumes approximately 18,000 logic slices of a Virtex-4 device, compared to 1,700 slices for a single channel DDC LogiCore reference design. Although there are some limitations in decimation factors and dynamic range, this new core represents an improvement in the channel-per-slice ratio by a factor of more than 20!

3. QUAD INPUT 256 CHANNEL DDC

Following a slightly different architecture, the 256 channel DDC IP Core 7151 features 64 channel DDC stages using advanced signal processing techniques like the previous example. This allows four separate input signals, typically sourced from four independent A/D converters.

Figure 4 shows a PMC (PCI mezzanine card) module based on this architecture featuring four 200 MHz 16-bit A/D converters, the 256 channel DDC core, and a PCI-X interface for control and data delivery to the PMC carrier board.

In this design each of the four 64-channel banks can have a different decimation setting that ranges from 128 to 1024, in steps of 64, for a total of 15 different values. For a 200 MHz A/D sample clock, this means a range in DDC channel bandwidth from about 156 kHz to 1.25 MHz, matching GSM channel requirements very nicely.

Each 64-channel bank is fed by its own 4-input multiplexer so that any one of the four A/Ds can source each bank, with selection independent from the other banks. This means that each bank can be dedicated to a different A/D or all four banks can share a single A/D. All other combinations are also possible.

4. WEIGHING THE BENEFITS

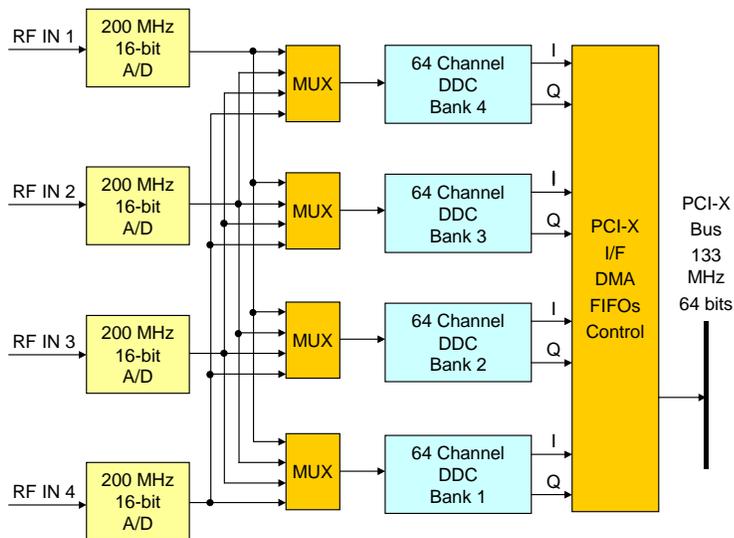


Figure 4. Model 7151 PMC 256 Channel Digital Down Converter

Each of the 256 DDC channels can be independently tuned using a 32-bit tuning word. The tuning words are stored in a double-rank RAM table accessible from the PCI bus for easy programming. This allows each of the 256 frequencies to be entered separately and then synchronously transferred to operating DDCs with a hardware trigger or software command. Switching between frequencies is phase continuous.

The decimating low pass filters for all DDCs within a 64-channel bank share a common set of coefficients, also stored in user accessible RAM. Each of the 15 decimation settings is supported by its own unique set of coefficients, automatically selected by changing the decimation setting.

Any number of DDC channels can be enabled within each bank. All enabled channel outputs are delivered in complex format (16I + 16Q or 24I + 24Q) in a channel interleaved sequence into four FIFOs.

A four-channel DMA controller delivers output data from each of the FIFOs to four different PCI bus memory destinations. Full chaining operation is supported and the DMA controller can interrupt the system processor or start another operation when each transfer is complete.

This 256 channel DDC fits inside a single Virtex-5 XC5VSX95T FPGA, a midrange part in the SXT family. It consumes virtually all of the resources and is fully characterized to operate at 200 MHz.

Comparing this DDC IP core to ASIC alternatives reveals the impressive benefits offered by FPGA implementations for high channel count applications.

Figure 5 shows resources required for a 174-channel GSM-E DDC system. With an ASIC-based implementation, eleven 16-channel boards would be required driving the system power to 88 watts and resulting in a cost of \$563 per channel.

FPGAs draw more power and are more expensive than ASIC devices. However, by using a single 256 channel Model 7151 module, system power goes down to 30 watts and the cost per channel is about \$69. It is important to realize that as the number of DDC channels increases, FPGA implementations become increasingly attractive.

Obviously, with such space, weight and power savings, new applications become feasible such as very compact, high-density communication systems and small, unmanned surveillance systems for fixed, mobile and airborne installations.

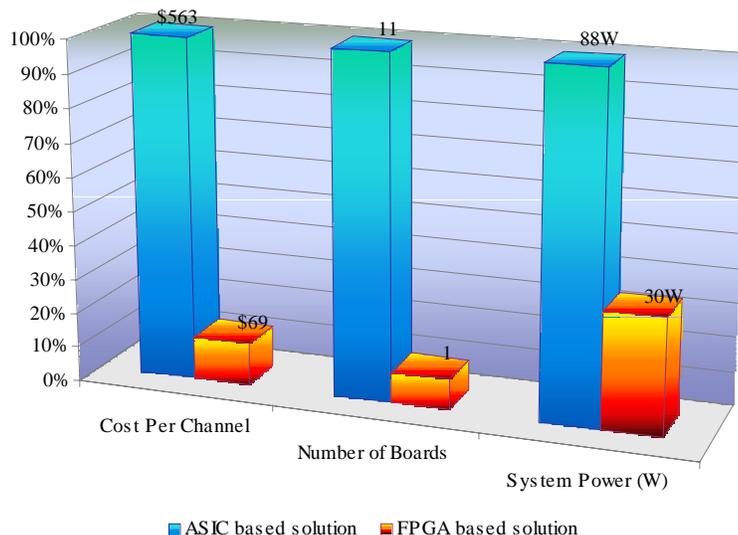


Figure 5. 174 Channel E-GSM DDC System – ASIC vs. FPGA

Driven by both military and commercial markets, FPGAs will continue to evolve to meet the insatiable demands of SDR based systems

Copyright Transfer Agreement: The following Copyright Transfer Agreement must be included on the cover sheet for the paper (either email or fax)—not on the paper itself.

“The authors represent that the work is original and they are the author or authors of the work, except for material quoted and referenced as text passages. Authors acknowledge that they are willing to transfer the copyright of the abstract and the completed paper to the SDR Forum for purposes of publication in the SDR Forum Conference Proceedings, on associated CD ROMS, on SDR Forum Web pages, and compilations and derivative works related to this conference, should the paper be accepted for the conference. Authors are permitted to reproduce their work, and to reuse material in whole or in part from their work; for derivative works, however, such authors may not grant third party requests for reprints or republishing.”

Government employees whose work is not subject to copyright should so certify. For work performed under a U.S. Government contract, the U.S. Government has royalty-free permission to reproduce the author's work for official U.S. Government purposes.

