Combo FPGA/DSP Moves into Multi-Protocol Design

Designers can offload a number of DSP operations onto FPGAs for highly efficient, cost savings designs.

by Arun Iyengar, and Deepak Boppana, Altera Corporation

The need for higher data rates is driving the evolution of wireless cellular systems from narrowband 2G Global System for Mobile Communications (GSM), IS-95

[1]

systems to current-generation Wideband Code Division Multiple Access (W-CDMA)-based 3G and 3.5G systems supporting peak data rates up to 10 Mb/s. Future third-generation partnership project long-term evolution (3GPP LTE) specifications point to complex signal processing techniques, such as multiple-input multiple-output (MIMO) along with new radio technologies like orthogonal frequency-division multiple access (OFDMA) and multi-carrier Code Division Multiple Access (MC-CDMA).

These approaches are key to achieving target throughputs in excess of 100 Mb/s. Alternate OFDM-based broadband wireless technologies such as WiMAX are now achieving transmission speeds in excess of 70 Mb/s. Higher order modulation and variable rate channel coding have enabled this improvement in data rates. Complex spatial signal processing schemes, including beam forming and MIMO antenna techniques, are also paths to increasing data rates at the expense of additional hardware.

These technologies create challenges for base station designers requiring scalability and cost-effectiveness as well as flexibility across multiple evolving standards. Designers must comply with such critical demands as increased processing speed, design flexibility, and a cost-reduction path to design in multi-protocol wireless base station capability.

Defining Multi-Protocol

However, it is important to first clearly define what a multi-protocol wireless base station is. It does not mean one that includes GSM, CDMA, W-CDMA, WiMAX and other protocols in the same base station, although the term could be perceived as

that. Rather, multi-protocol means that a wireless base station has the capability of supporting a W-CDMA to LTE and versions beyond LTE. Or a WIMAX BTS is capable of supporting future changes to the nascent standard.

To achieve that goal, designers must provision their designs to meet current and future requirements. For example, W-CDMA turbo de-coding requirements are up to two megabits per second (Mbps), which can easily be done by a coprocessor within a digital signal processor. But as designers move forward with LTE, the same turbo decoding must be performed at a rate of 50 or more Mb/s and the hard coprocessor in a DSP cannot support this requirement. Similarly, for functions such as crest factor reduction (CFR) which require recursive FFT and iFFT processing with low latency means the designer has to move away from cycle count and to a parallel approach. Again, the designer cannot rely solely on digital signal processing (DSP).

FPGAs are used in both examples above to offload DSPs. By deploying intelligent partitioning between FPGAs and DSPs, designers achieve an optimal combination of features and design savings. As a result, component count on an OFDM baseband board, for example, is significantly reduced from a large number of DSPs to a few DSPs and an FPGA. By taking this approach, designers perform cost-effective implementations and at the same time, provision for sufficient throughput to satisfy future changes required by multi-protocol wireless base station designs.

Mobile and wireless service providers and operators want wireless base station OEMs to assure them that the W-CDMA base stations they're placing in the field have the capability of supporting LTE. Hence, OEMs need to future proof their base stations with multi-protocol designs. This means a base station family will have the capability to move virtually seamlessly from one 3GPP release, for example, to a newer one without major costly re-designs, and ripping and replacing already deployed hardware.

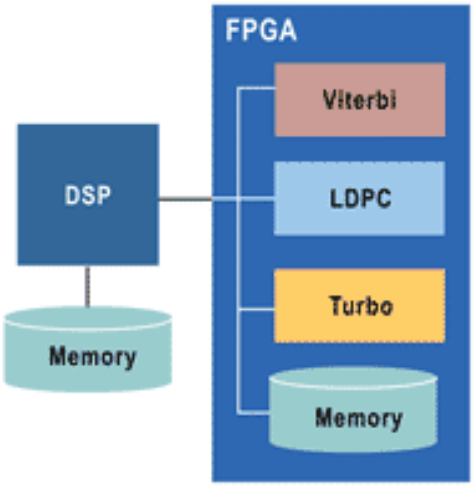

With LTE, wireless base station designers need to be aware there will be significant changes on the radio side. W-CDMA signal modulation will move to OFDM, as part of the migration to LTE, which has different characteristics. OFDM is more robust for delivering high throughput, but at the same time, stresses the throughput

capabilities of base stations.

OFDM also changes the peak to average characteristics of the modulated signal requiring new techniques to achieve CFR. In addition, there are more stringent requirements on the error vector magnitude (EVM) requiring the designer to pay particular attention to not just the type of algorithm used, but also the type of device to implement the algorithm.

On the baseband side, designers have to consider different data rates moving from W-CDMA to OFDM since the required throughput is considerably higher. Also, WiMAX has so far been based on data and is not voice oriented. But when voice is introduced, designers must start provisioning similarly to that of wire line systems since the quality of service (QoS) for voice is different than for data.

Partitioning Strategy

Partitioning strategy between FPGAs and DSPs depends on processing requirements, system bandwidth, configuration, and the number of transmit and receive antennas. Figure 1 shows a typical DSP/FPGA partitioning for baseband physical layer (PHY) functions in an OFDMA-based system such as WiMAX or LTE. System throughput is expected to be in excess of 70 Mb/s by incorporating advanced multiple antenna technologies.

Baseband Physical Layer can be categorized into bit- and symbol-level processing. Bit-level blocks include randomization, forward error correction (FEC), interleaving, and mapping to quadrature phase shift keying (QPSK) and quadrature amplitude modulation (QAM) functions on the transmit side. The corresponding receive processing bit-level blocks are symbol de-mapping, de-interleaving, FEC decoding and de-randomization.

All bit-level functions except FEC decoding are straightforward and not computationally intensive. For example, randomization involves modulo-2 addition of the data bits with the output of a simple pseudo-random binary sequence generator. While FPGAs offer more flexibility for bit-level manipulations than DSPs with fixed bus widths, the low computational complexity allows DSPs to manage these functions.

On the other hand, FEC decoding includes Viterbi decoding, turbo convolutional decoding, turbo product decoding and low-density parity check (LDPC) decoding. All are computationally intensive and consume significant bandwidth when they are performed with DSPs.

However, by using an FPGA coprocessor based on Stratix III or Cyclone III devices, system performance is boosted by an order of magnitude compared to conventional DSP processor-only-based implementations. FPGA coprocessors offload DSPs and efficiently execute computationally intensive blocks of a DSP algorithm due to inherent parallelism. For example, the FEC baseband processing operation of a WiMAX system can be offloaded to an FPGA coprocessor, as shown in Figure 2.

Symbol-level functions in OFDMA systems include sub-channelization and de-sub-

channelization, channel estimation, equalization and cyclic prefix insertion and removal functions. The time-to-

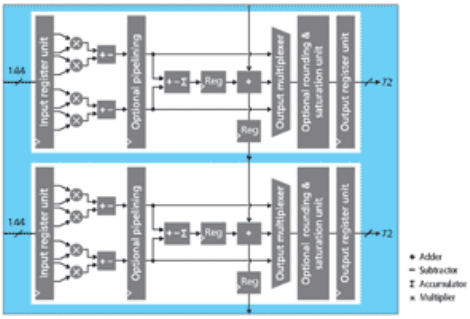

[2]

frequency domain conversion and vice-versa are implemented using fast fourier transform (FFT) and inverse fast fourier transform (IFFT), respectively. Channel estimation and equalization can be performed offline and involve more control-oriented algorithms better suited for DSPs. Conversely, FFT and IFFT are regular data path functions involving complex multiplications at very high speeds and are better suited for FPGA implementation.

Figure 3 shows embedded DSP blocks contained in Altera’s high-end Stratix III FPGA. DSPs usually have up to eight dedicated multipliers, whereas a Stratix III device, offers up to 896 18 &#215 18 dedicated multipliers. These dedicated multipliers translate into throughputs of up to 492 GMACs, about two orders of magnitude over DSPs with about 8 GMACs performance.

Such a big signal processing difference between FPGAs and DSPs is further accentuated when dealing with base stations using advanced, multiple antenna techniques such as space time coding (STC), beam forming, and MIMO schemes. The combination of OFDM-MIMO is widely regarded as a key enabler of higher data rates in current and future WiMAX and LTE wireless systems.

Figure 1 shows multiple transmit and receive antennas used at a base station. Symbol processing is performed separately for each antenna stream before MIMO decoding is performed, producing a single bit-level data stream. The symbol-level complexity grows linearly with the number of antennas when implemented on DSPs that serially perform operations.

For example, when two transmit and two receive antennas are used; FFT and IFFT consume about 60% of a 1 GHz DSP when the transform size is assumed to be 2048 points. In contrast, a multiple antenna-based design scales very efficiently when implemented with FPGAs. In this case, FPGAs provide parallel processing and time-multiplexing between the data from multiple antennas. The same 2 &#215 2 antenna FFP/IFFT configuration can be implemented using less than 5% of a mid-sized FPGA.

Digital IF Processing

Baseband channel card data is sent to an RF card for subsequent digital intermediate frequency (IF) processing. This processing includes digital up

conversion, CFR and DPD. Digital IF extends the scope of digital signal processing beyond the baseband domain to the antenna &#151 to the RF domain. This increases system flexibility while reducing manufacturing costs. Moreover, digital frequency conversion provides greater flexibility and higher attenuation and selectivity performance than traditional analog techniques.

CFR and DPD are required to improve the efficiency of base station power amplifiers. These functions also help to significantly reduce the total cost of an RF card. Both CFR and DPD involve complex multiplications at samples rates as high as 100+ Mb/s. Similar to DUC, digital down conversion (DDC) is required on the receive side to bring the IF frequency down to baseband. Both DUC and DDC use complex filter architectures including finite impulse response (FIR) and cascaded integrator-comb (CIC) filters.

Advanced FPGAs provide hundreds of 18 &#215 18 multipliers operating at speeds as high as 550 MHz. This design approach provides a platform capable of processing multiple channels in parallel and also gives designers a cost-effective integrated single-chip solution.

About the Authors

Arun Iyengar is senior director , Communications Business Unit and Deepak Boppana is technical marketing engineer for Altera Corporation, (408) 544-7000; www.altera.com.

Source URL (retrieved on *03/09/2014 - 6:14pm*):
http://www.wirelessdesignmag.com/product-releases/2007/03/combo-fpga/dsp-moves-multi-protocol-design

Links:
[1] http://www.wirelessdesignmag.com/sites/wirelessdesignmag.com/files/legacyimages/0703/altera1_lrg.gif
[2] http://www.wirelessdesignmag.com/sites/wirelessdesignmag.com/files/legacyimages/0703/altera3_lrg.gif