

Boosting Wireless Subsystem Performance with FPGA Co-Processing

Analyzing different hardware/software partitioning schemes.

by Dave Nicklin
Vertical Marketing and Partnerships, Wireless
Xilinx, Inc.
dave.nicklin@xilinx.com

Tom Hill
System Generator Product Manager
Xilinx, Inc.
tom.hill@xilinx.com

You can realize significant improvements in the performance of signal processing functions in wireless systems. How? By taking advantage of the flexibility of FPGA fabric and the embedded DSP blocks in current FPGA architectures for operations that can benefit from parallelism.

Common examples of operations found in wireless applications include finite impulse response (FIR) filtering, fast Fourier transforms (FFTs), digital down and up conversion, and forward error correction (FEC) blocks. Xilinx® Virtex™-4 and Virtex-5 architectures provide as many as 512 parallel embedded DSP multipliers, capable of running in excess of 500 MHz, to provide a peak DSP performance of 256 GMACs.

By offloading operations that require high-speed parallel processing onto the FPGA and leaving operations that require high-speed serial processing on the processor, you can optimize overall system performance and cost while lowering system requirements.

Sub-System Partitioning Choices

The FPGA can be used with a DSP processor, serving either as an independent pre-processor (or sometimes post-processor) device, or as a co-processor. In a pre-processing architecture, the FPGA sits directly in the data path and is responsible for processing the signals to a point when they can be efficiently and cost-effectively handed off to a DSP processor for further lower-rate processing.

In co-processing architectures, the FPGA sits alongside the DSP, which offloads specific algorithmic functions to the FPGA to be processed at significantly higher speeds than what is possible in a DSP processor alone. The results are passed back to the

DSP or sent to other devices for further processing, transmission, or storage (Figure 1).

The choice of pre-processing, post-processing, or co-processing is often governed by the timing margins needed to move data between the processor and FPGA and how that impinges on the overall latency. Although a co-processing solution is the topology most often considered by designers – primarily because the DSP is in more direct control of the data hand-off process – this may not always be the best overall strategy.

Consider, for example, the latest specifications for 3G LTE, in which the transmission time interval (TTI) has been reduced to 1 ms, down from 2 ms for HSDPA and 10 ms for WCDMA. This essentially requires that data be processed from the receiver and through to the output of the MAC layer in less than 1,000 μ sec.

Figure 2 shows that using an SRIO port on the DSP running at 3.125 Gbps, with 8b/10b encoding and a 200-bit overhead for the Turbo decode function, results in a DSP-to-FPGA transfer delay of 230 μ sec (that is, nearly a quarter of the TTI period just to transfer the data). Taking into account other expected delays, the Turbo codec performance required to meet these system timings is a very demanding 75.8 Mbps for 50 users.

Using an FPGA to process the Turbo codecs as a largely independent post-processor not only removes DSP latency but saves time because there's no need to transfer the data at a high bandwidth between the DSP and FPGA. This reduces the throughput rate of the Turbo decoder down to 47 Mbps, a decrease that allows you to use more cost-effective devices and have reduced system power dissipation.

Another consideration is whether to use soft- or hard-embedded processor IP on the Xilinx FPGA to offload some of the system processing tasks, which in turn offers the possibility of additional cost, power, and footprint reduction benefits. Given such a wide range of signal processing resources, complex functions such as those found in baseband processing can be more optimally partitioned between the DSP processor, the FPGA-configurable logic blocks (CLBs), embedded FPGA DSP blocks, and

FPGA embedded processor. Xilinx offers two types of embedded processors: the MicroBlaze™ soft-core processor – often used for system control – and the higher performance PowerPC hard-core embedded processor for more complex tasks.

FPGA-embedded processors provide an

opportunity to produce efficient FPGA implementations from a higher level of abstraction than what HDL tools such as MATLAB models and C code can provide. With development tools like Xilinx System Generator for DSP and the AccelDSP™ synthesis tool, you can go from algorithm

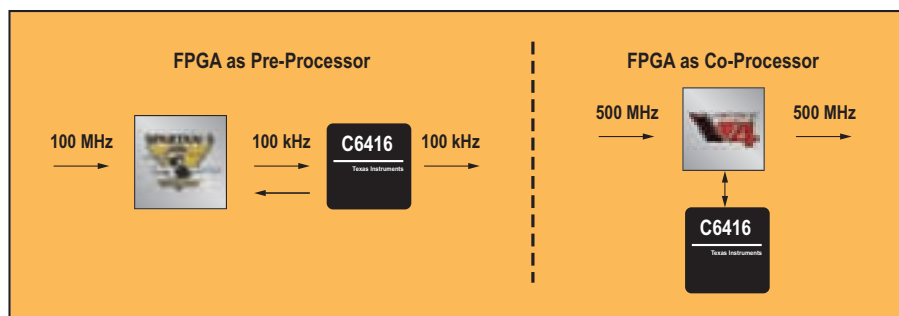



Figure 1 – FPGA as a pre-processor and co-processor solution



Implications for Turbo Decode (50-user example):

Processing bandwidth available for Turbo Decode = 380 μ s
 Throughput rate necessary = (576 bits \times 50 users) / 380 μ s = **75.8 Mbps**
 Throughput rate *without* DSP transfer latency = (576 bits \times 50 users) / 610 μ s = 47 Mbps

Data-transfer latency directly impacts cost and power dissipation.

Figure 2 – LTE example of co-processing data-transfer latency issues

opportunity to consolidate all non-critical operations into software running on the embedded processors, minimizing the total amount of hardware resources required for the overall system.

The Importance of Software and IP

A critical issue is how to unlock all of this potential capability. You need to consider both the software needed to abstract the complexity of the problem and the availability of IP, focused on key areas where FPGAs can provide an optimal solution.

Xilinx is dedicated to developing industry-leading tools and ecosystems that allow

to silicon as seamlessly as possible.

There is also an increasingly important ecosystem of tool providers whose products take development up to the electronic system level (ESL) through C/C++-to-gate design flows. ESL design tools are aimed at providing an integrated system-level approach to the production and integration of hardware-accelerated functions and the control code for processors controlling these functions.

No single high-level language or software tool is suitable for all of the different elements found in today's complex systems. The choice of language and design flow is

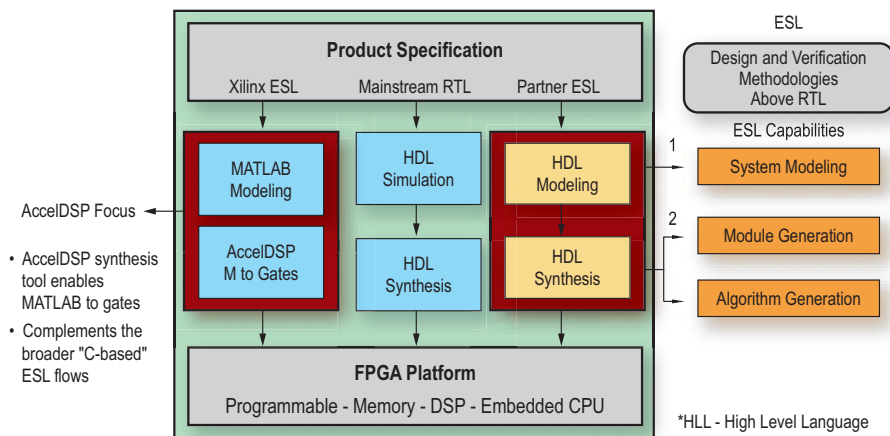


Figure 3 – System level-to-FPGA design flows

governed by the customer and sometimes the individual engineer. Xilinx has therefore developed a comprehensive set of integration features in order to meet customer needs and provide the most optimal design environments (see Figure 3).

Conclusion

Xilinx is also making significant investments to provide a comprehensive suite of high-value IP, boards, and reference designs that cover many critical areas in radio card and baseband applications. These include FFT/iFFT, modulation, digital up and down conversion, and crest factor reduction.

An example of this focused approach has been the development of industry-leading high-performance FEC functions, such as Turbo encoders and decoders, optimized for specific wireless standards and FPGA architectures. As we demonstrated in our analysis of 3G LTE latency and Turbo decoder throughput requirements, hardware acceleration of FEC functions and their impact on system architecture is an increasingly important necessity in modern wireless equipment design.

Although some specialist DSP processors integrate such functions as embedded blocks over time, it can take many months from when the parameters of an FEC function for a new wireless standard are set before the resulting embedded acceleration block appears in silicon. Once embedded, there are still challenges remaining, and occasions when not all of the functionality in the embedded blocks work as required. In the meantime, standards have quickly evolved to include new requirements that cannot be supported by the fixed embedded blocks.

Because of these situations, designers require flexibility. They are looking for the ability to swiftly develop and deploy complex baseband functions such as FEC and then adapt them based on feedback from field trials and evolving standardization efforts. Perhaps they wish to add their own proprietary IP to differentiate their solution in the marketplace. It is in these circumstances that designers should not simply consider the current solution portfolio from a supplier, but also investigate how easily the solutions can be adapted and what level of support and tools the supplier can provide. ●●

Performance. Delivered.

X5400m
PCI Express DSP Module

Download Data Sheets NOW!

Features

- Two 400 MSPS, 14-bit A/D Channels
- Two 500 MSPS, 16-bit D/A Channels
- ±1V, 50 Ohm, SMA Inputs and Outputs
- Xilinx Virtex5, SX95 FPGA
- 1GB DDR2 DRAM, 4MB QDR-II SRAM
- 8 Rocket IO Private Links, 2.5 Gbps each
- >1 GB/s, 8-lane PCI Express Host Interface
- Power Management Features

Perfect for

- Wireless Receiver and Transmitter
- WLAN, WCDMA, WiMAX front end
- RADAR
- Electronic Warfare
- High-Speed Data Recording and Playback
- High-Speed Servo Controls
- IP Development

Innovative Integration
... real time solutions!

805.520.4260 phone
www.innovative-dsp.com

TAKE THE NEXT STEP (Digital Edition: www.xcellpublications.com/subscribe/)

- Get more details on Xilinx DSP tool flows and our ecosystem partners.
- See the complete range of solutions available to wireless systems design engineers.