



XAPP726 (v1.0) July 25, 2005

Benefits of FPGAs in Wireless Base Station Baseband Processing Applications

Author: Hong-Swee Lim

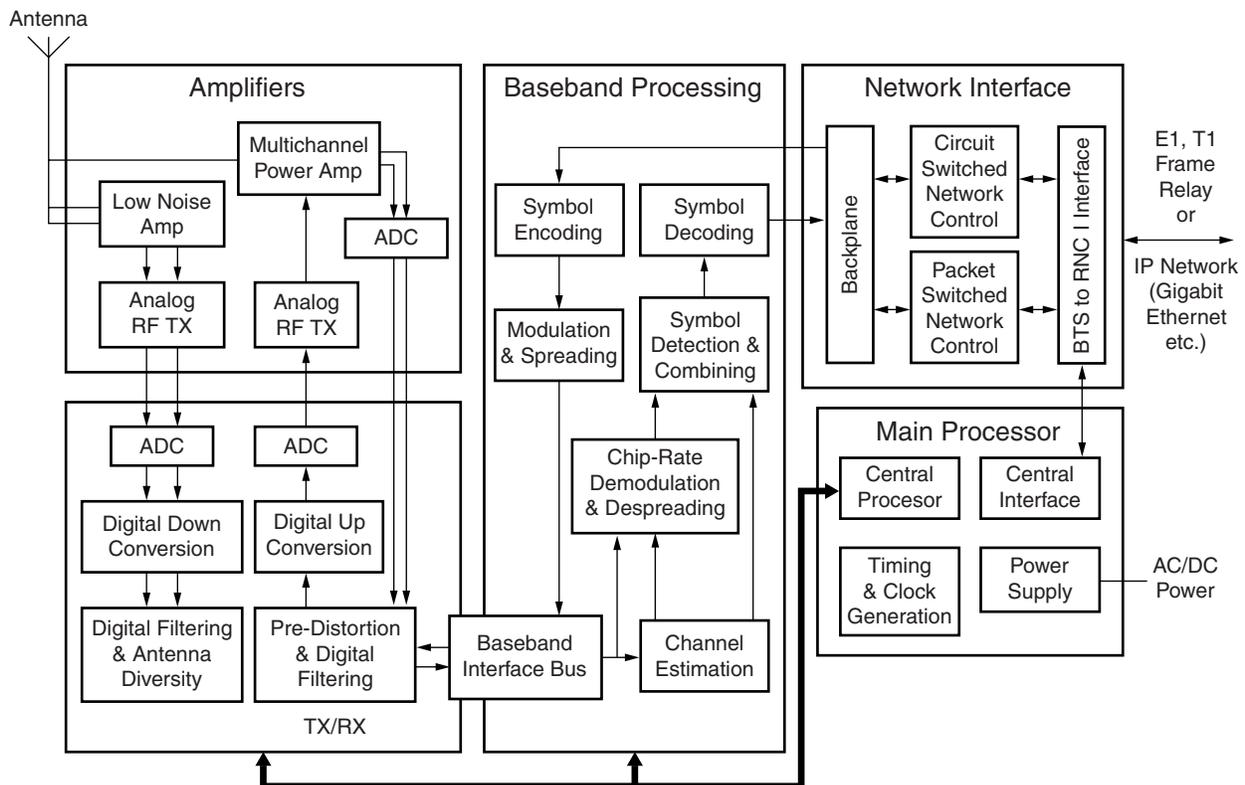
Summary

With the deployment of the 3G-wireless infrastructure gaining momentum, equipment manufacturers and network operators are searching for a highly optimized baseband processing solution with the greatest freedom in the areas of cost, flexibility, and scalability to meet increasing data service and time-to-market requirements.

This application note provides an overview of the baseband processing of a typical W-CDMA base station, along with the associated implementation challenges faced by W-CDMA equipment manufacturers, including the silicon cost, flexibility, and scalability trade-offs. It also provides a summary of an optimal solution to meet the extensive processing requirements of the W-CDMA specifications while retaining flexibility and minimizing the overall cost.

Introduction

The W-CDMA system is one of the leading wideband cellular technologies used in the 3G cellular market. At the heart of the W-CDMA system is the 3G base station, shown in Figure 1.



xapp726_01_062005

Figure 1: Block Diagram of the 3G Base Station

The early W-CDMA trial system, initially proposed and developed by NTT DoCoMo, and the UMTS Terrestrial Radio Access (UTRA), developed by Siemens and others, served as the basis for consideration by the International Telecommunications Unions IMT2000 initiative. The initiative used these two efforts as a baseline to eventually merge the development efforts into

one standard (W-CDMA) with two implementations: Frequency Division Duplex (FDD) and Time Division Duplex (TDD), governed by the Third Generation Partnership Project (3GPP) organization. The system is sometimes referred to as 3GPP W-CDMA or 3GPP Release 99 (Rel'99) to differentiate it from the earlier wideband CDMA versions.

In April of 2001, 3GPP Release 4 (Rel'4) of W-CDMA was completed and provided minor improvements to the 3GPP Rel'99 transport and radio interface and architecture. In March 2002, 3GPP Release 5 (Rel'5) was released. This new release defined features such as:

- the High Speed Downlink Packet Access (HSDPA) channel
- the IP Multimedia Subsystem (IMS)
- the IP Universal Terrestrial Radio Access Network (UTRAN).

These features provided significant spectral and network efficiency improvements and performance and functionality advantages over the 3GPP Rel'99 and 3GPP Rel'4 standards.

The 3GPP Rel'5 standard was developed so that its enhancements can coexist on the same RF carrier as the currently deployed 3GPP Rel'99. Thus, a current 3GPP Rel'99 carrier can be upgraded to support legacy 3GPP Rel'99 and new 3GPP Rel'5 terminals in the same 5-MHz band. HSDPA is one of the key 3GPP Rel'5 features that offers significantly higher data capacity and data user speeds on the downlink through the use of very dynamic adaptive modulation, coding, and scheduling with Hybrid Automatic Retransmission Request (H-ARQ) processing. Through HSDPA, operators benefit from a technology that provides improved end-user experience for web access, file download, and streaming services.

Future enhancements to 3GPP Rel'5 are mainly driven by the need for an improved user experience enabled by packet-based services combining non-real-time and real-time components available in both stationary and mobile environments. Consequently, the next planned release of 3GPP, Release 6 (Rel'6), is focused on improving quality of service, capacity, service enablers, and delivery for multimedia packet-based services. One of the key features targeted for 3GPP Rel'6 is the Enhanced Uplink for Dedicated Channels (EUDCH) feature. As the importance of IP-based services increases, there are growing demands to improve coverage and throughput and reduce the delay of the uplink. Applications that can benefit from an enhanced uplink include services such as video clips, multimedia, e-mail, telematics, gaming, and video streaming.

This application note focuses on the FDD mode of the W-CDMA specification, although most of the content is also applicable to TDD. Whenever the term W-CDMA is used, it is referring to the 3GPP Rel'5 specifications for W-CDMA FDD mode.

Implementation Challenges

One of the most critical sub-systems in a 3G base station is the baseband processing card. This card takes digitized baseband radio signals and base station control and signaling as input, and it produces multiple simultaneous user data channels, including voice and data streams.

The baseband processing can be broadly divided into symbol-rate processing and chip-rate processing. The symbol-rate processing for voice users is somewhat different from that of data users. Symbol-rate processing for voice users encompasses the Viterbi decoder, deinterleaver, and rate matching. Symbol-rate processing for data additionally includes turbo encoding and decoding. Chip-rate processing involves the processing of the input data at chip-rate or a multiple of the chip-rate. In 3GPP, chip-rate processing requires significant parallel processing and, therefore, is ideally suited for Xilinx FPGA implementation. The inherent flexibility of the FPGA provides a platform that can be easily migrated to support multiple cell sizes, deployment scenarios, and multiple standards.

3GPP Rel'5 has generated tremendous interest in the 3G base station industry for the wide-ranging economic and deployment benefits it offers. Below are some of the crucial requirements encountered by the baseband processing designers that ultimately drive the choice of the hardware platform.

Processing Speed

Several advanced signal processing techniques such as FIR filters, FFT/IFFT, and turbo convolution coding/decoding are being used in baseband signal processing. These are very computationally intensive and require several billion multiply and accumulate (MAC) operations per second.

Flexibility

With the continuing changes in the 3GPP specification, a flexible baseband processing solution is extremely important to enable equipment manufacturers to upgrade their base station to include new feature sets from future 3GPP specifications. Flexible solutions allow equipment manufacturers to add their own IP to differentiate their product and to scale their design based on cell size, latency, number of antennas and sectors, input sample bandwidth, and detection methods.

Low Power/Cost

To reduce cost, the design should be low power to allow the use of smaller and lower cost cooling systems and less expensive battery backup systems. Because current base station designs require significant space, the size of the baseband design directly impacts the deployment and rental costs for the space used by the base station.

Time to Market

Because 3GPP Rel'5 is relatively new, the availability of software, boards, IP, and reference designs enables the designer to shorten the development cycle. This has a direct impact on the time to market of the product and the early success in gaining market share.

FPGA-Based 3GPP Rel'5 System Design

This section outlines how to address each of the implementation challenges.

Addressing Baseband Processing Challenges Using FPGAs

The prospect of a software and/or hardware programmable base station meets manufacturing cost targets and has obvious benefits for both equipment manufacturers and network operators.

The ability to change hardware functionality all the way through the deployment cycle, into production, and ultimately after deployment, is an attractive proposition. It reduces the risk of an equipment manufacturer missing a delivery date to a customer due to a variety of issues including:

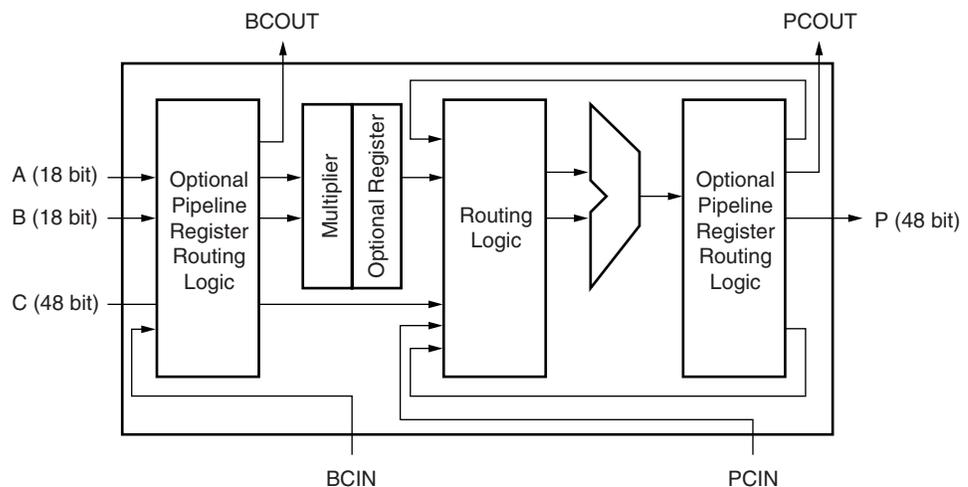
- a change in the specification
- a bug being found in an Application Specific Integrated Circuit (ASIC)
- an Application Specific Standard Product (ASSP) manufacturer slipping the schedule for their chipset

A hardware and software programmable base station allows for new design solutions or bug fixes to be implemented after field deployment using the existing deployed hardware. Such fixes are implemented by downloading either new software and/or new hardware configuration files, thereby mitigating possible field re-work or return costs. Further, localization, tuning, and other system enhancements can be deployed long after a product has shipped, with little restriction as to the type of change being implemented.

Processing Speed

To meet the increasing processing requirements of the baseband module, the most advanced FPGA, the Virtex™-4 device, has dedicated signal processing resources such as the integrated

XtremeDSP™ slices, sometimes referred to as DSP48s, shown in [Figure 2](#). Each slice contains a dedicated two's complement signed, 18 x 18 bit multiplier, and a three-input adder/subtractor with feedback for accumulation modes. The addition of a seven-bit op mode multiplexer allows the dynamic configuration of the XtremeDSP slice for one of more than 40 operating modes, such as addition, multiplication, accumulation, MAC functions, MAC cascading, wide (48-bit) addition, and wide multiplexing. This level of functionality built into the XtremeDSP slice obviates the need to build conventional adder trees that are required in other FPGA solutions. Eliminating the need for the adder trees gives the Virtex-4 FPGA a tremendous price, performance, and power advantage over competing solutions. Depending on the Virtex-4 family member, as many as 512 XtremeDSP slices can be utilized, each capable of providing 500-MHz throughput.



xapp726_02_071205

Figure 2: Virtex-4 XtremeDSP Slices Feature 18 x 18 Multiply, 48-Bit Accumulator

Low Power

Using integrated features, such as the Virtex-4 XtremeDSP slices which consume only 2.3 mW/100 MHz, eliminates the need to use logic slices for many signal processing and arithmetic tasks, reducing the requirement for power-consuming routing resources. Another way to reduce system power consumption is to use FPGA devices with the embedded PowerPC™ processor found in the Virtex-4 FX platform [[Ref 1](#)]. This device allows the option to trade gates for processor cycles for sequential control tasks.

Multichannel Implementations

Shift Register Logic (SRL16) is a valuable feature for increasing compute density in multichannel implementations. The benefits of using SRL16s can be seen by using a simple Reed-Solomon encoder example. Implementing a single-channel Reed-Solomon encoder in a Virtex-4 device can consume 56 logic slices. For a 16-channel implementation, one approach is to replicate this 16 times, resulting in a consumption of 16 x 56 (or 896) slices. [Figure 3](#) shows another implementation of the 16-channel solution using the integrated SRL16s inside the Virtex-4 device. This solution consumes only 86 logic slices, representing only 10% of the 16X replicated version. SRL16s can substantially pack more signal processing into a smaller area, allowing the potential to target a much smaller device than is possible with other FPGA architectures.

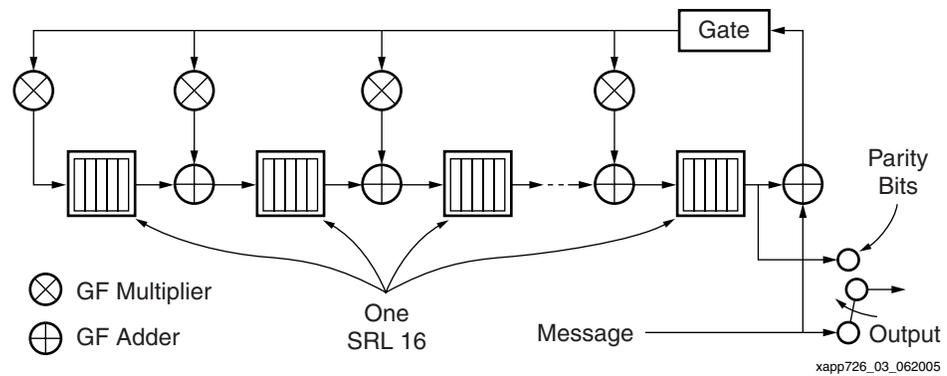


Figure 3: Efficient 16-Channel Reed-Solomon Encoder Using SRL16

Flexible Connectivity

In addition to embedded processors, the Virtex-4 FX family of devices includes multi-gigabit transceivers that are well-suited for interfacing the radio module to the baseband module cards using the industry standard OBSAI or CPRI interfaces. Also, the flexible interfaces, such as the PCI Express™ and serial RapidIO™, can be used to connect to other processors. The LVDS interface from the Virtex-4 device can also be used to connect to high-speed A/D converters to form part of the system data flow.

IP and Reference Designs

This section describes some of the IP and reference designs for baseband applications.

Forward Error Correction (FEC)

The dynamic market of fast evolving error-correction standards calls for a flexible FEC solution. The built-in programmability of FPGAs and the extremely high DSP performance enabled by parallelism are required to cope with the processing bandwidth of faster FEC schemes, demanding more processing bandwidth. Flexibility and programmability is also required to cope with evaluating the Bit Error Rate (BER) performance for various new schemes. Programmability allows the developer to download various FEC schemes to the same hardware platform to test and evaluate the effectiveness of the FEC scheme in an actual system environment. This allows the evaluation of the IPs before the actual system implementation.

The Turbo Convolution Code (TCC) provides an extremely effective way of reliably transmitting data over noisy channels. The TCC decoder operates very well under low signal-to-noise conditions and provides a performance close to the theoretical optimal performance as defined by the Shannon limit [Ref 2]. Turbo Convolutional Codec (Encoders and Decoders) IP that supports changing dynamic block size is crucial to enabling flexible, high-performance error correction. To facilitate accelerated BER plot generation, Xilinx has developed an integrated test harness inside the System Generator™ for DSP development tool, which operates within The MathWorks Simulink simulation environment. Using the hardware coprocessing capability of the System Generator for DSP allows the test harness to accelerate the generation of BER plots by orders of magnitude while operating within the Simulink simulation environment. The test harness gives full control over system-level parameters that can be modified from a MATLAB script to automate the generation of BER plots for different block sizes, coding rates, and iterations.

In 4G systems, Multiple Input Multiple Output (MIMO) processing greatly increases the communication bandwidth between transmitter and receiver by using multiple antennas to create parallel wireless channels over the air interface. This also adds to the complexity of modeling the channel effect in hardware. The System Generator tool and the hardware coprocessing functionality (commonly known as Hardware-in-the-Loop or HWITL) address this challenge.

High Speed Downlink Packet Access (HSDPA)

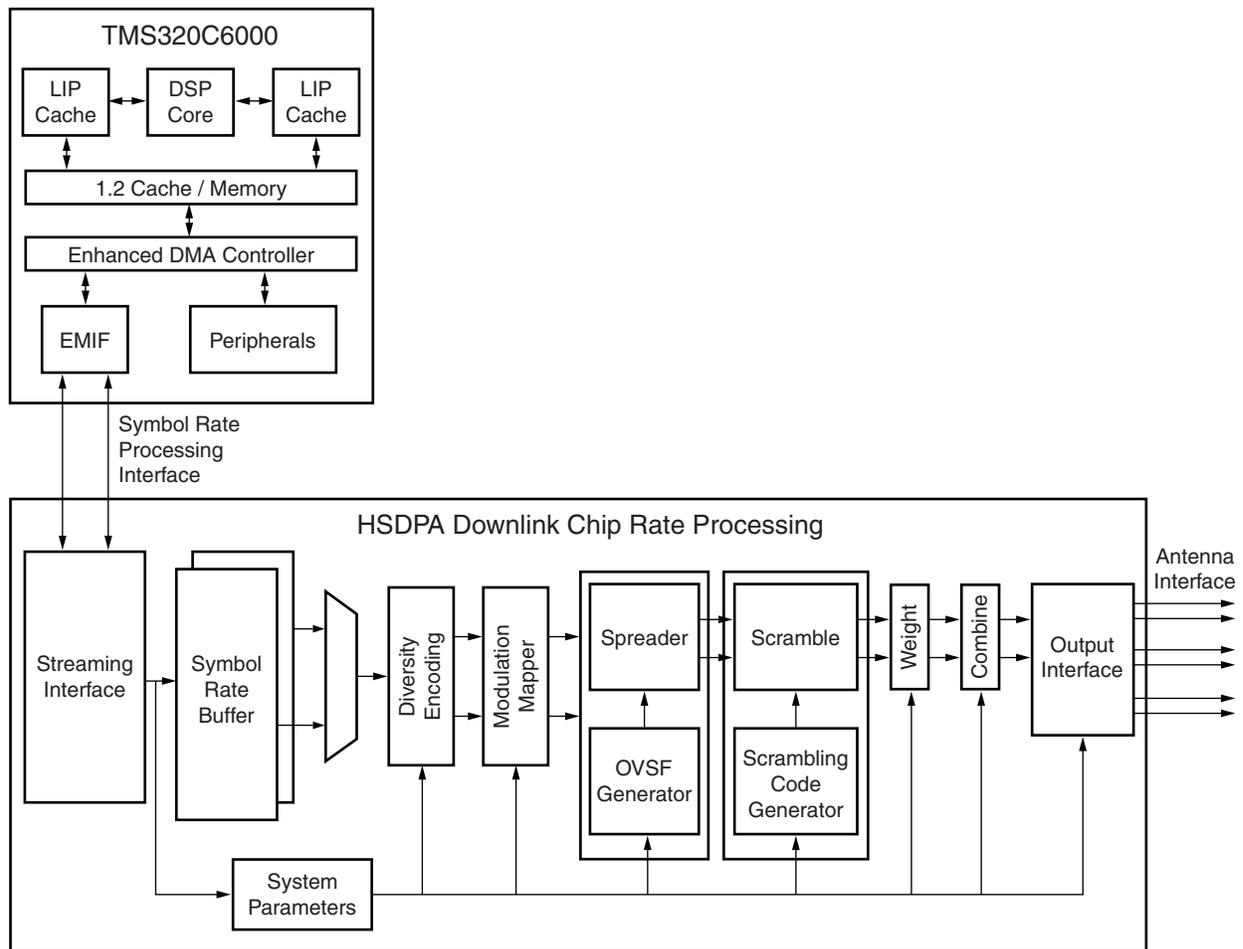
HSDPA involves cyclic redundancy check (CRC), turbo encoding, rate matching, interleaving, and quadrature amplitude modulation. These specifications demand significant levels of signal processing and are ideal for implementation in an FPGA, leveraging parallel processing techniques to increase design performance and improve throughput. The inherent flexibility of the FPGA allows the design to be easily scalable. This flexibility allows for picocell, microcell, and macrocell implementations.

The Xilinx HSDPA reference design includes the following key benefits:

- The HSDPA implementation is extremely small and fits nicely into a low-cost Spartan™-3 device. This is an efficient way to add HSDPA to an existing 3GPP Rel'4 compliant solution.
- The System Generator source design is available, allowing engineers to leverage the techniques used and quickly integrate the techniques into their own existing baseband processing design.
- The HSDPA reference design is used as a coprocessor to upgrade existing systems and process three sectors to simultaneously deliver 14.4 Mb/s bandwidth to each sector.

The Xilinx solution connects to the Texas Instruments (TI) EMIF or other proprietary interfaces and is used as a coprocessor to lower overall system cost and increase system performance.

Figure 4 shows an example of connecting the HSDPA reference design to the TI EMIF.



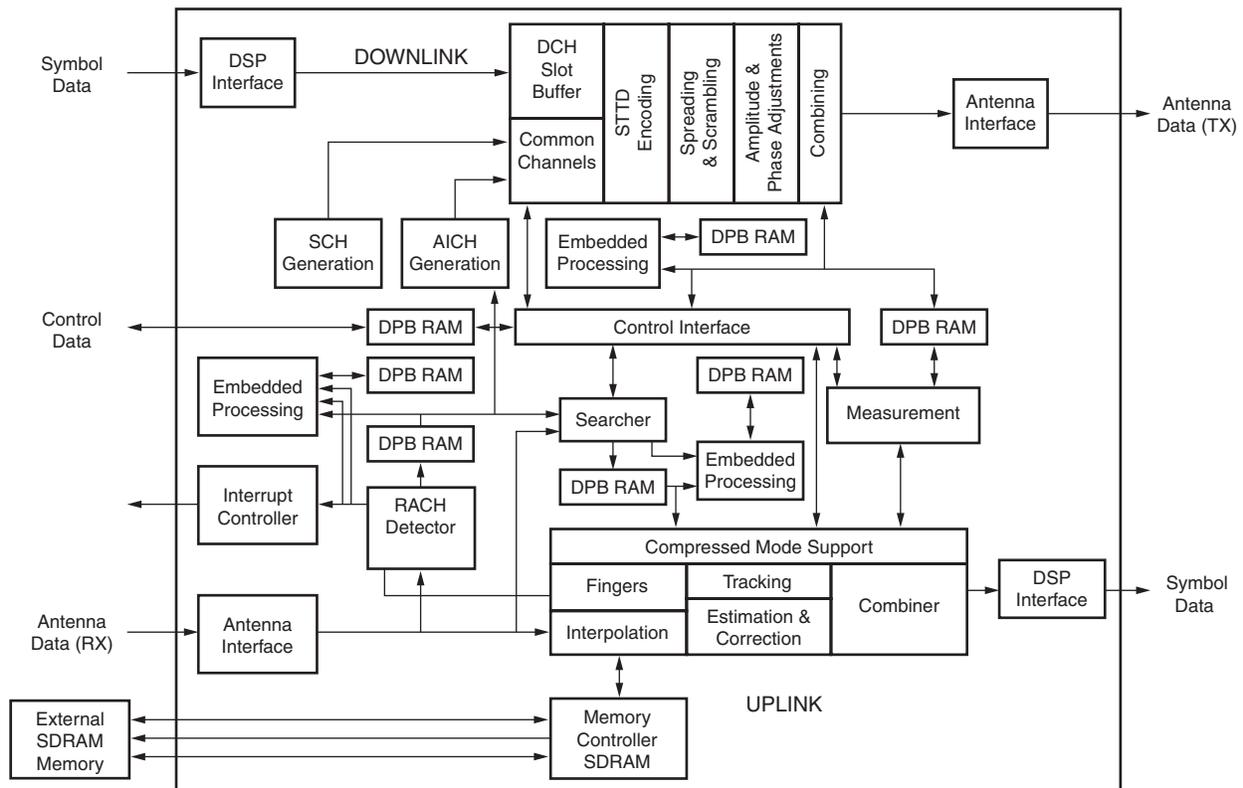
xapp726_04_062005

Figure 4: Coprocessor Interface

Random Access Channel (RACH)

For scarce wireless spectrum efficiency, packet transmission over the uplink shared channels is used. The RACH is an uplink shared channel for initial channel access and short data burst transmissions to the network. The RACH is one of three channels shared by all mobiles in the cell.

The RACH detector is located in the Uplink Chip Rate section of the Base Station Baseband processing card as shown [Figure 5](#).



xapp726_05_062005

Figure 5: Uplink Chip Rate Processing Reference Designs

The RACH detector finds any preamble sent from anywhere in the cell by correlating the code against a copy of itself, where a strong correlation only occurs when the codes are aligned in time. This is done with an adder/subtractor/accumulator module.

Xilinx has developed the W-CDMA RACH Preamble Detection reference design based on a novel implementation of the Fast Hadamard transform. The resulting solution meets stringent performance demands while yielding lower cost than traditional implementations based upon matched-filter architectures. The key parameters that affect the design are cell size, latency, antenna flexibility, input sample width, and detection methods. The proposed design provides a solution that can be easily modified to support variations in all these parameters.

Key benefits of the Xilinx W-CDMA RACH Preamble Detection reference design include:

- A flexible and scalable solution that can be easily implemented in a range of Xilinx FPGAs, including Spartan-3 and Virtex-4 devices, based on cell size, latency, number of antennas and sectors, input sample bandwidth, and detection methods.
- Area-efficient implementation that easily supports multiple cell sizes, deployment topologies, and 3G standards enabling system tuning even after deployment, to reflect the current needs of the network. This is a significant benefit over traditional fixed

implementations in which many of the parameters must be predefined and cannot be modified after deployment.

- Each access slot is independent from another, allowing sector hopping if required.
- Open VHDL source enables easy and efficient integration into existing baseband processing design.

Searcher

When a radio signal travels between the handset and the base station, it bounces off buildings and other obstacles within the environment. This signal scattering causes multiple images of the same signal to occur with slight time differences at the receiver. These scattered signals are called multipaths. The multipath propagation delays vary in relation to each other as the user moves through the cell. It is vital for the base station to track the top four or five multipaths from each user in order to recover as much energy of the original signal before trying to recover the data carried by the signal.

The basic function of the Searcher is to estimate the impulse response of the propagation channel between a known cell and a mobile [Ref 3]. It profiles the timing relationship for the top four or five multipaths for every user.

An important parameter of the Searcher is the Update Rate. This is the rate at which the Searcher returns and takes a new profile of the multipaths for each user. The Update Rate is also the rate at which the Searcher provides course resolution timing information of the multipaths to the Rake Receiver. Xilinx implemented a Searcher reference design to enable the customer to interface their proprietary post-processing algorithm to a high-density correlator function. This reference design demonstrates the benefits of an FPGA's inherently parallel architecture. It enables a step and repeat method that is highly valuable for addressing computation requirements of MIMO and Multi-User Detection (MUD) processing in future 3GPP releases.

DSP Development Tools

Baseband design development tools allow developers to move DSP algorithms into FPGAs quickly and easily. These development tools save valuable time in the design process, allowing the designer to concentrate on differentiating the product in the marketplace.

System Generator for DSP

The Xilinx System Generator for DSP is a high-level system design tool that allows development and verification of optimized DSP algorithms for Xilinx FPGAs. The System Generator for DSP operates within The MathWorks Simulink modeling and simulation environment. It enhances Simulink through powerful capabilities like hardware coprocessing (HWITL), verification, and debug, all within a common environment.

The System Generator includes a Simulink library of functional blocks for building DSP, arithmetic, and digital logic circuits. These polymorphic blocks either automatically compute their output types based on their inputs or have their quantized output types specified explicitly. Developers combine Xilinx blocks with MATLAB and Simulink blocks to create a realistic test bench and analyze data computed by the model. The high level of abstraction provided by the System Generator greatly simplifies algorithm development and verification. The System Generator also includes a new block, allowing the instantiation of an XtremeDSP slice and the ability to configure it for one of its many operating modes. The System Generator can also be used to build extremely sophisticated multirate systems.

The System Generator for DSP seamlessly integrates with many common simulators and development solutions. Figure 6 shows how the System Generator for DSP integrates into a system development environment.

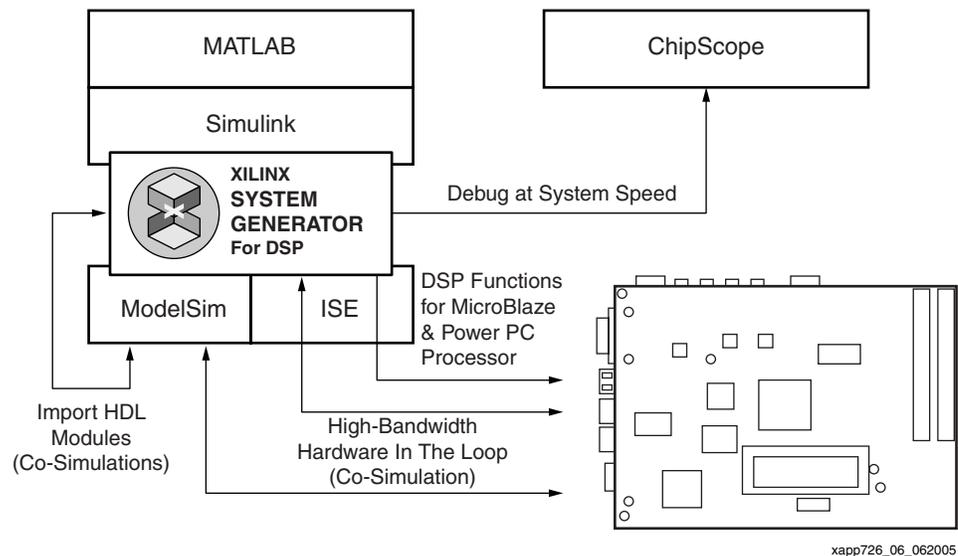


Figure 6: System Generator for DSP

In addition to a system-level modeling library, the System Generator includes a code generator that automatically generates a synthesizable VHDL netlist from Simulink model. This netlist includes IP blocks that have been carefully designed for high performance and density in Xilinx FPGAs.

HWITL (Hardware Coprocessing)

The baseband processing reference design is designed to work with the HWITL simulation platform as part of the Xilinx System Generator design flow. Employing HWITL reduces software simulation time by orders of magnitude, freeing time to focus on the important tasks of debugging and full system verification. There are two main benefits of HWITL:

- Designers verify their designs in hardware without leaving Simulink and feed the results from hardware back into Simulink. This mirrors traditional programmable DSP design flows such as those offered by TI.
- Designers accelerate the simulation by several orders of magnitude over software-based simulation without the need for expensive emulation hardware.

Rapid prototyping, debugging, and system verification are increasingly becoming the gating factors to shortening design cycles. The baseband processing reference designs and design flow address these gating factors by providing the necessary design and test files for easy integration and modification of the design to meet the unique processing requirements of the 3GPP baseband specification.

Reference Design

The reference design files are available upon request by sending e-mail to: rach@xilinx.com.

Conclusion

The solutions described in this document can be extended to support different permutations of uplink and downlink channel capacity by allowing proportionally more logic capacity for one than the other.

Because the hardware is flexible, the ratios can be changed throughout the lifetime of the base station in order to meet changing customer use models and demands. Such base station

architectures allow operators and manufacturers to balance the requirements of their customers, helping to reduce deployment costs while delivering the required quality of service demanded by the early adopters. And as demands change, it is also possible to reconfigure the hardware to meet changing traffic patterns throughout the day, allowing extra capacity at peak times and better data rates and service at other times.

To meet the baseband requirements of the W-CDMA specification, Xilinx developed the HSDPA coprocessor, RACH, and Searcher reference designs. This allows customers to develop an FPGA-based drop-in solution to add the mandatory 3GPP Rel'5 feature of HSDPA without having to totally redesign existing ASIC and control software.

References

1. Virtex-4 FPGA Handbook.
2. C Berrou, A. Glavieux, and P. Thitimajshima, "Near Shannon Limit Error-Correcting Coding and Decoding Turbo Codes," IEEE Proc 1993 Int Conf. Comm., pp1064-1070.
3. WCDMA Requirements and Practical Design, Sect 3.4.6.1 (page 92), ISBN 0-470-86177-0.

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
07/25/05	1.0	Initial Xilinx release.