



CHAMELEON
S Y S T E M S , I N C.

Wireless Base Station Design Using Reconfigurable Communications Processors

Introduction

As the communications market continues its explosive growth and rapid rate of change, equipment vendors struggle with the conflicting goals of performance, flexibility, low cost and fast time-to-market. Traditional processing approaches such as DSPs, ASICs, ASSPs and FPGAs all force the designer to sacrifice at least one of these key parameters.

A new class of processor from Chameleon Systems, the Reconfigurable Communications Processor (RCP), enables designers to meet all these goals simultaneously for multi-channel, data-processing intensive applications.

Design Challenges

Insatiable Demand for Performance

The explosive demand for increased communications bandwidth is forcing communication equipment vendors to focus on performance. There are several driving factors behind this demand for performance. First, there is the dramatic increase in the number of Internet users. The second driving factor is the shift in Internet usage to more bandwidth intensive applications, such voice over IP (VoIP) and streaming video.

Traditionally, these impacts were limited to fiber optic and copper networks. Today, performance demands are also impacting the wireless market as users start to implement wireless LANs and access the Internet over cellular connections. Analysts forecast the number of worldwide wireless Internet subscribers to grow from 2 million in 1999 to over 93 million in 2004, and these new users will demand the same features and bandwidth they get on their desktop computers.

The current Wireless Application Protocol (WAP) supports limited email and text-based web browsing. Emerging wireless protocols that support the features demanded by next-generation users require more processing power. A single 30kHz TDMA channel, for example, requires about 40Mips for channel filtering, equalization and modulation/demodulation. In comparison, a 1.2288 Mcps CDMA correlator serving perhaps 20 users requires about 10 GOp/sec for just the rake receiver processing.

Need for Flexibility

In addition to performance considerations, equipment vendors are forced to build in flexibility to adapt to rapidly changing market requirements. Convergence of voice, data and video, changing standards, and a high demand for evolving features require the equipment vendor to build systems that are flexible and field upgradeable.

Today, there is a strong demand for multi-protocol systems that can adapt to changing traffic patterns or support multiple markets. U.S. wireless infrastructure for example, must handle analog traffic, cdmaOne and TDMA digital traffic as well as the emerging cdma2000 standard. Since no one can accurately predict the volume of each type of traffic over the next few years, vendors strive to create flexible systems that can instantaneously adapt to changing patterns. Flexibility also allows vendors to differentiate their products and create higher value using proprietary algorithms.

Time-to-market Requirements

In an environment where requirements are rapidly changing, the time required to bring new products to market can spell the difference between success and failure. Equipment vendors need an approach that allows them to quickly design, debug and verify their systems.

Cost Pressures

Despite the increased demand for new features, the per-user price of communications equipment and services will decline over the next few years. Price reductions are, in fact, necessary to bring emerging technologies into the mainstream market.

Choosing a Processor Technology

In the past, equipment suppliers based their designs on Application Specific Standard Products (ASSPs) with programmable logic acting as glue. Alternatively, they employed programmable DSPs or FPGAs during early design and field trials, then ported the design to an ASIC implementation to reduce cost for high volume production.

Chameleon Systems' RCP provides suppliers with a new category of processor with the most favorable characteristics for multi-channel, data-processing intensive communications applications.

	RCP	ASIC	FPGA	DSP	ASSP
Flexibility	High	Low	High	High	None
Cost	Low	Low	High	Medium	Low
Performance	High	High	Medium	Low	High
Time-to-market	Medium	Long	Medium	Medium	Short

Table 1 - Comparison of Available Technologies

DSP

The biggest limitation of DSPs for multi-channel applications is performance. A single DSP does not have the bandwidth to process multiple wide data streams at speed. As a result, designers are forced to partition the system using multiple DSPs, which significantly increases design complexity and cost per channel.

ASIC

ASICs offer high performance, but take longer to design and lack the programmability required to provide adequate flexibility. Once deployed, systems built using ASICs suffer long delays and high costs for even minor changes.

FPGA

FPGAs are flexible, but cannot provide a complete solution for signal and protocol processing. To implement a complete system, they must be combined with a processor through a specially designed interface. The high-density FPGAs required to implement such a system also carry a significant cost premium.

From a performance standpoint, FPGAs are unpredictable. Carefully optimized designs are faster than DSPs. However, minor changes to the design can result in long optimization cycles to get back

to the required performance. In addition, FPGAs have bit-orientated structures that incur serious overhead in speed and gate resources when applied to wide datastream applications.

While FPGA solutions are flexible, they require hundreds of milliseconds to reprogram due to large configuration files. This prevents the designer from using the reconfigurability to increase performance and reduce cost by applying multiple algorithms to a data stream in a single chip.

ASSP

Application Specific Standard Products (ASSPs) are typically not available for emerging standards. Designers may have to wait a year or more after a standard is frozen before they can find an ASSP that fits their needs.

When ASSPs are available, they enable fast time-to-market and cost effective implementations, but offer little flexibility to the designer. With an ASSP-based design, vendors cannot use their intellectual property to differentiate their system, and have no flexibility to add features or adapt to changing standards.

RCP

Chameleon Systems' Reconfigurable Communications Processor, enables wireless base station designers to achieve a combination of low cost, fast time-to-market, high performance and complete flexibility.

The RCP provides a platform-based approach that incorporates three core architectural technologies: a complete 32-bit embedded processor subsystem, a high-performance 32-bit reconfigurable processing fabric, and eConfigurable™ Technology, Chameleon's patented instantaneous reconfigurability.

RCP Architecture

Chameleon Systems' CS2112 RCP, built using a 0.25-micron CMOS process, provides 24,000 MOPs and 3,000 MMACS processing power – about ten times that of a high-performance DSP. This is enough to implement 50 channels of chip-rate processing for cdma2000 in a single device.

The following architectural features, shown in Figure 1, enable this tremendously high performance.

Reconfigurable Processing Fabric

The Reconfigurable Processing Fabric (RPF) is organized in slices, each of which can be independently reconfigured. The CS2112 includes four slices consisting of three tiles each. Each tile comprises seven 32-bit Datapath Units (DPUs), two 16x24-bit single-cycle multipliers, four Local Store Memories (LSMs) and a Control Logic Unit (CLU). A dynamic interconnect connects the modules within the Fabric.

High Bandwidth Programmable I/O

Unlike DSPs, the RCP includes four banks of programmable I/O pins that provide tremendous bandwidth. The CS2112 includes four banks of 40 I/O pins per bank, providing an aggregate I/O bandwidth of 2 GByte/sec., enabling high-performance data streaming for signal processing and protocol processing applications.

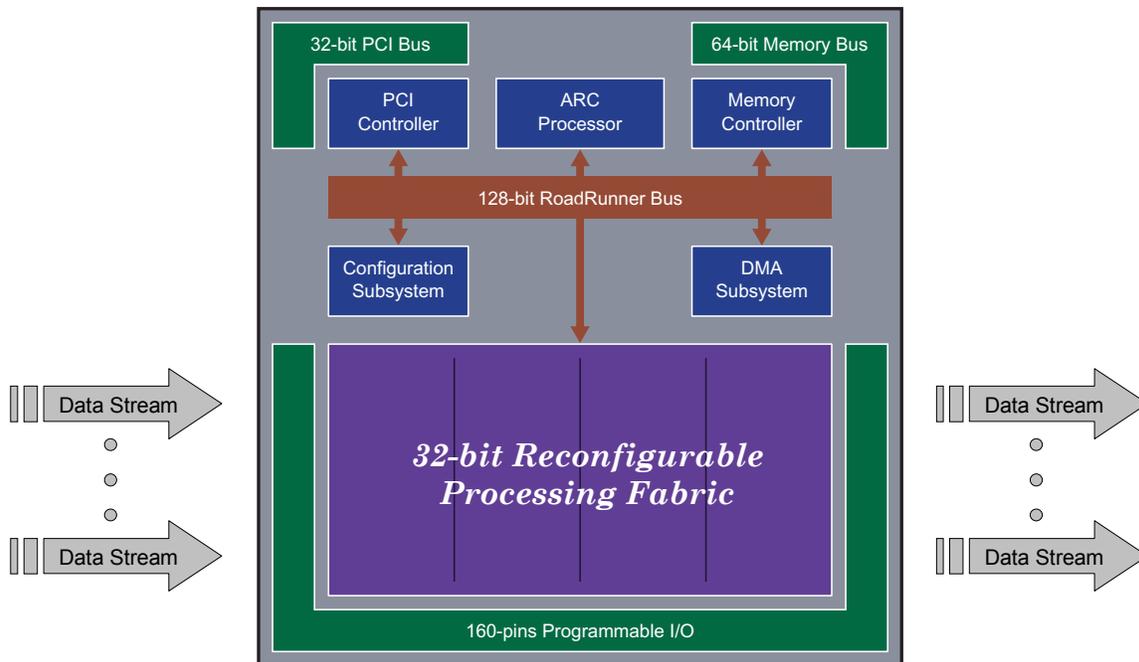


Figure 1 - Chameleon Systems' RCP

32-bit ARC Processor

Optimized by Chameleon for the RCP, the 32-bit ARC Processor delivers 120 MIPS at 125 MHz. Licensed from ARC Cores Ltd., the processor employs a four-stage pipeline, 64 general-purpose 32-bit registers and a 32-bit address space. The execution unit provides fast barrel shift, fast multiply, swap, min/max and normalize operations. The processor includes a 4-KByte instruction cache and a 4-KByte data memory. A fully integrated JTAG port is also included.

64-bit Memory Controller

The 64-bit Memory Controller provides a complete high-performance solution for off-chip memory. The SSRAM Controller and the SDRAM Controller both supports a 1-GByte/sec transfer rate. The Flash EEPROM Controller supports a wide variety of devices in x8 and x16 configurations with capacity from 8 to 32 Mbits.

DMA Subsystem

The DMA Subsystem supports 16 DMA channels, transferring data between the modules in the Embedded Processor System and to/from the Local Store Memories in the RPF. Each DMA channel can be set up as a continuously streaming buffer.

32-bit PCI Controller

The 32-bit PCI (Peripheral Component Interface) Controller provides a complete interface solution to the PCI bus, supporting Master/Slave operation.

Flexibility Through eConfigurable Technology

Chameleon Systems' proprietary eConfigurable technology enables the entire processing fabric of the RCP to be reconfigured instantly for the ultimate in flexibility. Utilizing a background configuration plane to store the next set of configuration bits, the next configuration can be loaded

from external memory in just 3 μ secs per slice, without interfering with active processing in the fabric.

Once loaded, the configuration in the background plane can be swapped into the active plane in just one clock cycle, allowing the RCP to instantaneously adapt to changing traffic patterns or signal quality. Contents of on-chip memory are maintained during reconfiguration, allowing the user to apply multiple algorithms to the same data without using off-chip buffers.

This technology also provides the benefits of traditional reconfigurable devices, allowing systems to be upgraded in the field to enable new features or to accommodate changes in emerging protocols.

Time-to-Market

Chameleon Systems' RCP utilizes a platform-based approach to insure the fastest possible time-to-market. Every sub-system in the device is fully integrated and pre-verified. This saves the designer the time required to implement a memory sub-system and controller, DMA engine, PCI interface, and micro-processor interfaces.

The RCP's co-design environment allows high-level algorithms to be implemented quickly using standard C and HDL languages. Complete observability and fully synchronous timing enables processor-style debug and verification of the reconfigurable processing fabric.

A break-thru eBIOS™ interface manages the interaction with the RCP, overcoming the challenges associated with custom SOCs.

eConfigurable technology also provides all the time-to-market advantages of a traditional reconfigurable device by enabling rapid in-circuit design iterations.

Cost

For multi-channel applications, the processor cost is typically measured by dividing the overall processor cost by the number of channels supported. The resulting cost-per-channel metric allows vendors to easily compare costs of approaches supporting different numbers of channels.

Chameleon's RCP provides a very cost-effective solution for wireless base station applications. High-volume pricing is projected to reach less than \$1.00 per channel of chip-rate processing by 2H 2001 for the North American cdma2000 standard.

Design of a cdma2000 Base Station

Wireless Base Station Infrastructure

The need for performance, flexibility, fast time-to-market, and low cost are especially critical in wireless base stations. Performance demands for next-generation systems are radically increased by the greater signal processing requirements of new standards, and the new features required by users. In the physical layer, for example, aggressive signal processing, such as beamforming and multi-user detection techniques are required to increase capacity and coverage.

Flexibility is required to handle the varying levels and quality of traffic from old and new protocols simultaneously. In CDMA base receivers, processing resources are allocated to received signals

depending upon their signal quality. Additional rake fingers, for example, are directed towards those channels suffering severe multipath interference.

In the access layer, service providers overlay different 2G and 3G protocols. Voice, data and video are expected to co-exist. This means that hardware and software resources must be dynamically allocated to users depending upon their bandwidth requirement.

Time-to-market pressures, already intense in the competitive communications market, are increased by the demand for deployment of early trial equipment even as the standards continue to evolve. Relentless cost pressures are driven by the demand for continual reduction in service pricing while the myriad of new features demanded by users continues to grow.

Finally, the growth of new features, shown in Table 2, drives the need for increased processing capability. Some of these applications will be installed after deployment of the hardware, further highlighting the need for flexibility.

OSI layer	Features	Enabling technology
Upper Layers	Voice, data, multimedia QoS Spoken email Emergency location Enhanced security	Adaptive compression algorithms optimized to link Voice recognition, synthesis Position finding algorithms Flexible encryption
Access	Interoperability	Multimode capability 1G/2G/3G overlay Future protocols Future connectivity with other wireless devices
Physical	Traffic resource management Multi-user detection Improve C/I	Dynamic resectoring High bandwidth signal processing Beamforming and noise rejection

Table 2 – New Features Driving the Need for High Performance

Figure 2 shows a block diagram of a typical cdma2000 base station. On the receive side, the antenna receives data which is sampled as parallel words. These samples, or chips, represent the basic unit of data in the wireless symbol domain.

Once the signal is received, the rake receiver searches through the signal's sample-time window and looks for sets of transmitted original and delayed versions. Each 'finger' of the rake searches for a given delayed version of the transmitted signal.

The demodulator processes the data and recovers the transmitted signal. Each finger of the rake receiver multiplies the received sample by a delayed version of the specific pseudo random number that was used to encode the data. This delay factor compensates for the multipath effects of the wireless channel.

In Chameleon Systems' CS2112 device, the Pseudo-Random Number Sequence Generator (PNGEN) is implemented using a pre-computed polynomial look-up-table and delay-line technique that achieves a throughput of 64 chips per clock. The signal is decoded using a match-filter technique that interpolates the received data. The derived data is then passed to a set of filter stages whose outputs are used to locate the best match based on a PN sequence.

The chip-rate and symbol-rate processors for a system with 50 user channels can be implemented in two Chameleon CS2112 RCP devices. As shown in the following figure, the chip-rate processor is implemented in one device and the symbol-rate processor is implemented in the second device.

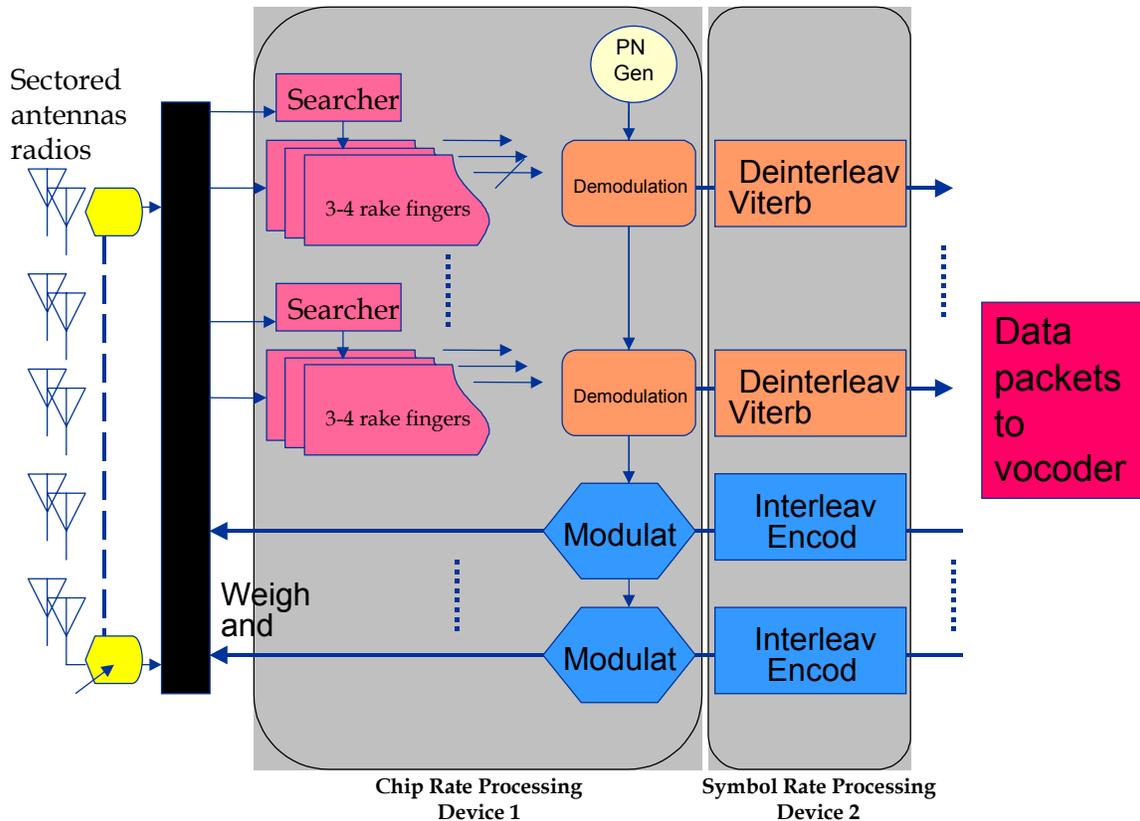


Figure 2 - cdma2000 Base Station

In Chameleon Systems' implementation of the wireless base station, eConfigurable technology is used to dramatically increase performance of the Reconfigurable Communications Processor. In this implementation, a frame of data is stored in the reconfigurable processing fabric's Local Store Memory and the device is instantaneously reconfigured to apply different algorithms to the data.

As shown in Figure 3, each frame of data, called a power control group (PCG) is 1250 μ sec long. The four algorithms that are applied to the data are loaded into the reconfigurable processing Fabric one at a time.

First, the entire Fabric is dedicated to PNGEN for 77 μ sec. While PNGEN is processing, the DMOD algorithm is loaded into the background configuration plane. In a single clock cycle, the

entire fabric is swapped to the DMOD algorithm. While the DMOD algorithm is active, the Finger Search algorithm is loaded into the background plane.

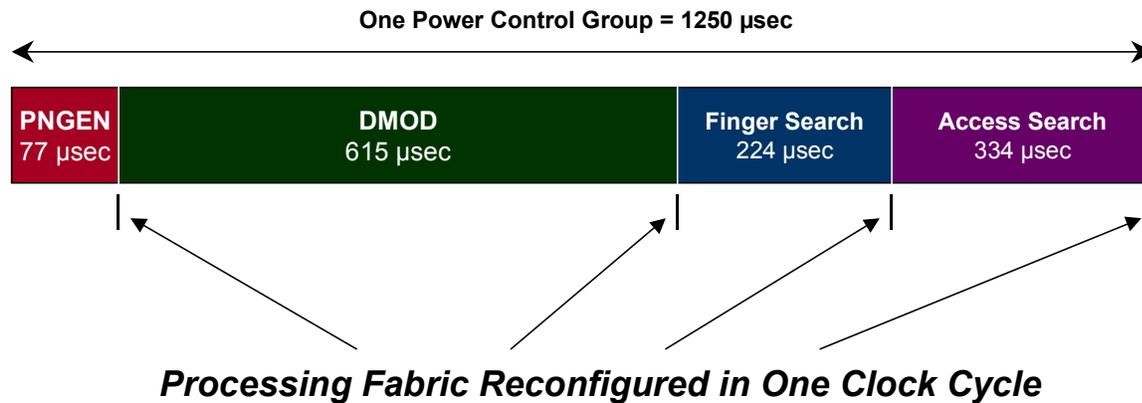


Figure 3 – Chip-rate Processing Using eConfigurable Technology

This continues until all four algorithms have been applied to the data. Since the entire RPF is dedicated to just one algorithm at a time, much higher performance, lower cost and lower power are achieved. In addition, there is no need to ‘move’ the data to the physical logic that implements the next algorithm, eliminating typical performance bottlenecks found in ASICs and FPGAs.

eConfigurable technology enables the entire chip-rate processing for a 50 channel system to be implemented in one Chameleon CS2112 device. Traditional approaches implement each of the four chip-rate processing algorithms as separate hardware modules in ASICs or FPGAs.

Summary

In the race to deploy next-generation wireless protocol systems, equipment vendors must overcome the challenges of performance, flexibility, cost and time-to-market. Chameleon Systems’ Reconfigurable Communications Processor meets all four of these goals enabling vendors to be first to market with the highest performance, most flexible products at the lowest cost.