



WP450 (v1.2) March 22, 2018

The Application of FPGAs for Wireless Base-Station Connectivity

By: Paul Newson

Painstakingly architected to be a generation ahead, Xilinx® 7 series and UltraScale™ devices provide the unprecedented efficiency dictated by today's crowded wireless base station landscape.

ABSTRACT

This white paper is concerned with connectivity between functional units within conventional and evolved wireless base stations. In-phase/quadrature (I/Q) radio sample distribution between the baseband and radio modules is addressed, as well as internal transport of traffic and control data within the base station.

The white paper is organized as follows:

- The [Introduction](#) section provides a summary of how Xilinx FPGA products provide an ideal interconnection medium between functional units within wireless base stations.
- The [Background](#) section provides an overview of the architecture of the conventional wireless base station, a view on how this is likely to evolve to meet future system requirements, and a review of the connectivity solutions typically adopted.
- The [Xilinx Solutions for Base-Station Connectivity](#) section presents a summary of Xilinx technology solutions for wireless base station connectivity. Both silicon devices and soft IP are addressed.
- The [FPGA Connectivity Architectures for Base Stations](#) section provides an analysis of basic connectivity requirements and architectures for a range of base-station applications. Several functional enhancements to the basic architecture are then described.
- The [Conclusion](#) section summarizes the main points of the white paper.

Introduction

Connectivity between functional units is a key element in the design of modern wireless base stations. The connectivity solution must provide high levels of throughput with low latency and offer the flexibility required to operate over a diverse range of system configurations. Field Programmable Gate Array (FPGA) technology is ideally adapted to meet these challenges because it is based in a configurable fabric well suited to the implementation of standard telecom interfaces and switching functionality and offers a significant number of flexible high-speed transceiver and I/Os, which are necessary to provide the physical interfaces.

This white paper examines the connectivity requirements within wireless base stations and demonstrates how Xilinx FPGA technology can be utilized to perform each of the functions associated with signal distribution. The principal applications considered are multi-mode macrocells and evolved high-density platforms, because these present the most significant challenges for the connectivity technology. Nevertheless, the principles and design techniques described are applicable to the wide range of base station architectures associated with heterogeneous network deployments.

The primary focus of this document is on functional analysis of the main components associated with base station connectivity, the implementation of these functions in Xilinx FPGA technology, and the expected resource and device-mapping requirements for several example applications. The practical goal is to provide an overview of how wireless base station connectivity applications can be addressed using a combination of Xilinx FPGA devices and connectivity IP, the latter being either available from Xilinx or developed by the end user.

Background

This section provides an overview of the conventional macrocell base station architecture, the connectivity requirements and architectures typically adopted in current applications, the evolution of cellular network design trends, and how these could impact new base station design.

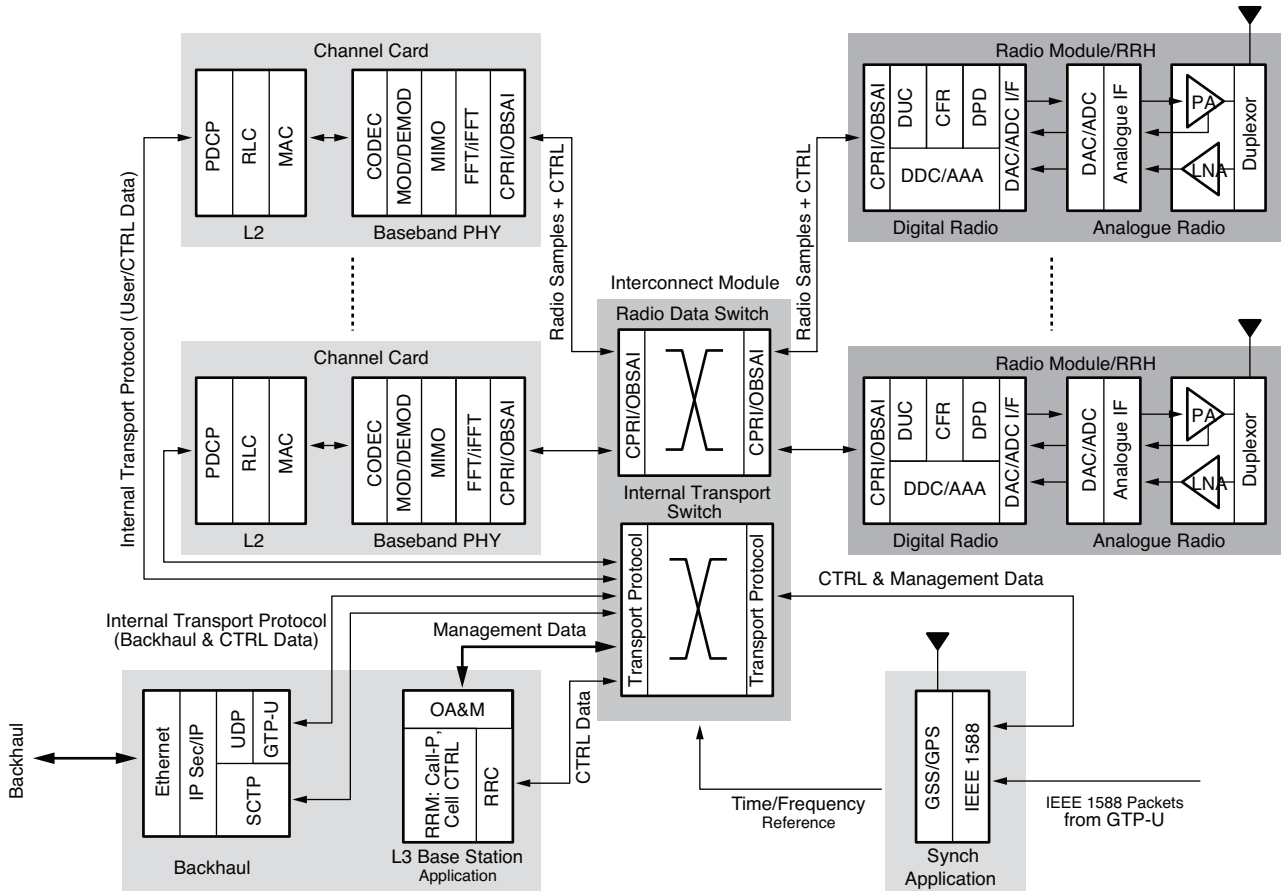
Conventional Wireless Base Station Overview

The conventional macrocell base station is used to address a wide variety of cellular air-interface standards and system deployments. Standards covered include: LTE, WCDMA, TD-SCDMA, GSM, and CDMA2000, with mixed mode operation being increasingly common. Deployments range from large multi-carrier, multi-sector implementations to smaller single sector microcells. The functionality performed by the base station includes analog and digital radio processing, baseband signal processing, and higher-layer and backhaul transport packet processing. Generally, a modular architecture is adopted in which each function is implemented in dedicated modules optimized for the specific functionality performed. Often, the architecture is common across a range of air-interface standards and deployment types; the application addressed by any particular base station is dictated by the number and types of modules deployed.

Interconnect between modules is required to transfer signal and control data between functional units in accordance with the requirements of the overall application. The connectivity solution must provide high levels of throughput with low and often deterministic latency, and offer significant flexibility such that many different configurations can be supported on the same platform. Consequently, base station connectivity is one of the key elements of system design.

In many architectures the interconnect functionality is centralized in a dedicated module. However, because this functionality is closely associated with the overall system control and baseband processing, it is often integrated within these modules. Even radio modules might provide some limited interconnect/switching functionality beyond a simple connection to the base station digital unit in order to support flexible radio signal distribution topologies such as tree, ring, and daisy chaining. The functionality typically performed by the interconnect unit includes physical layer interfacing, transport protocol processing, protocol conversion, data switching/routing, signal conditioning, and packet/signal synchronization.

A high-level functional block diagram of a typical macrocell base station configured to support LTE is shown in [Figure 1](#). The figure shows each of the common functional units, plus an interconnect module configured to provide both radio sample connectivity and internal transport for user and control data.



WP450_01_040114

Figure 1: Typical LTE Macrocell Base Station Functional Block Diagram

Evolutionary Trends

The fundamental requirement driving future development of cellular systems is to provide increased capacity and throughput with improved coverage at lower system cost [Ref 1]. In order to support this, a range of evolved base-station architectures are being introduced which differ from that of the conventional macrocell. Many of these architectures impose even more stringent requirements on the base-station connectivity network than those which exist today.

The principal factors driving architectural evolution can be broken down into three somewhat interrelated areas:

- The requirement to support multi-band and mixed-mode operation
- The evolution of cellular system topologies towards Heterogeneous Networks (HetNet)
- The evolution of cellular standards to support higher performance and greater flexibility with the introduction of LTE-A

The requirement to support multi-band and mixed mode operation comes from the drive to increase deployment flexibility and to reduce system cost. Today, support for such operating modes is not uncommon both at the base station and the module level. However, in the future, technology will allow manufacturers to address a wider range of requirements with support for ultra-wide band, with an increased number of carriers being possible. It is envisaged that in the future, the vast majority of base stations will provide integrated support for multi-band and mixed-mode operation.

The HetNet [Ref 1] vision is that cellular service is provided by a network of tightly coordinated base stations of various types, each of which is optimized to provide coverage for a specific environment. The objective is to provide higher levels of coverage and performance by tailoring deployments to suit environments. The base station types envisaged include conventional macrocells; distributed base stations; high-density base stations, deployed as a part of cloud radio-access networks (CRANs); adaptive antenna array (AAA) configurations, both centralized and distributed; small cells; and relay nodes.

Essentially, this range is addressed by two fundamental base-station architecture types optimized to meet quite different requirements:

- Architectures similar to that of the conventional macrocell that cover high-capacity, high-performance applications such as high-density cell sites, distributed base stations, macrocells, and centralized AAA systems.
- Highly integrated base stations that address small cell related applications.

Connectivity is a major factor driving the architecture of the high-capacity applications, but it is less important in the small-cell architecture; the low capacity and the low number of cells provisioned (typical one or two) allow for high levels of functional integration. This in turn obviates the requirement for data transfer between independent modules. Consequently, this white paper focuses on the multi-mode, high-density platform architectures.

LTE-A [Ref 2] [Ref 3] [Ref 4] is an evolution of the existing LTE standard with the objective being to increase throughput/capacity and network coverage while maintaining backward compatibility. LTE-A features were initially introduced in 3GPP Releases 10 and 11, but they will be further

enhanced in Release 12 and beyond. Key concepts introduced and developed in LTE-A are support for carrier aggregation, higher-order multiple-input/multiple-output (MIMO), cooperative multi-point (CoMP) operation, and relay node. Each of these features is designed to increase user throughput/capacity and to improve coverage — and, as a consequence, lead to increasing user data rates and the requirement to transfer more control data within the base station and across the network in general.

The principal impacts of network evolution to support the HetNet vision and next-generation standards on base-station connectivity are:

- Throughput requirements will increase significantly — at least linearly with the bandwidth and the number of antennas supported
- The number of data types to be supported will grow
- Latency requirements are likely to become even more critical
- Flexibility requirements will increase

Hence, the use of flexible high-speed interconnect is set to become an even greater challenge in high-capacity base-station design in the future.

Evolved Base-Station Architectures

This white paper focuses on the connectivity requirements and architectures used in conventional macro-cell and multi-mode high-density platforms. The principal evolved macro-cell architectures that fall into this category are distributed, AAA, and CRAN base stations.

Distributed Base Stations

The distributed base station is a variant of the macrocell in which the radio units are remotely located. The physical separation of the baseband unit (BBU) and remote radio unit (RRU) can range from hundreds of meters to several tens of kilometers. From the cellular system viewpoint, this architecture opens up the possibility of implementing more flexible network topologies than the conventional hexagonal grid based on centralized base stations. However, in terms of the base-station architecture itself, the only real difference is that the base station must have the capability to support RRUs located at a range of distances from the centralized BBUs. This functionality is most often provided by the interconnect function within the base station, which must support optical transmission over long distances, as well as support the capability to compute and compensate for the transmission delay to each RRU.

In general, the number of RRUs supported and the associated throughput and interconnect flexibility requirements within the distributed architecture are similar to those of the conventional macrocell. Indeed, most manufacturers base conventional and distributed base stations on the same basic product with specific variants to support both applications. Consequently, connectivity requirements and architectures are typically common; these are described in [Wireless Base-Station Connectivity](#).

Adaptive Antenna Array (AAA) Systems

AAA systems are based on the principle that electronic beamforming can create a large number of independent coverage regions within the geographical area served by the base station [Ref 4]. This requires a significantly increased number of antennas compared with conventional applications. The objective is to increase capacity and/or improve coverage within a specific local area. Such systems are normally deployed in urban and suburban environments and complement macrocell coverage in high-capacity areas.

Beamforming is performed on both the downlink (DL) and uplink (UL) and involves the intelligent combination of signals over N baseband sources and M antennas (where $N \leq M$) in order to form orthogonal beams carrying independent user data. This process comprises three basic functions:

- Calibration of the complete TX and RX radio paths within the base station
- Computation of the adaptive beamforming weights
- Implementation of the beamforming network itself

The beamforming network requires access to multiple baseband and radio data streams, which means that the system interconnect function within the BBU, radio, or stand-alone connectivity module is ideally located to perform this functionality.

Various AAA base-station configurations have been proposed. These broadly fall into two categories referred to in this white paper as *conventional* and *integrated* AAA systems.

Conventional AAA System

The conventional AAA system is based on an architecture similar to that of the macrocell, with independent modules performing baseband and radio processing. The system can be implemented either as a centralized or a distributed architecture and provides capacity and coverage commensurate with the macrocell. The conventional AAA architecture imposes several additional requirements on base-station interconnect over and above those associated with the macrocell:

- Due to the significant increase in the number of antennas that must be supported, data rates and number of connections increase significantly
- Beamforming technology (calibration, weight computation, and signal combining) must be implemented within the base station. If this is implemented as part of the interconnect function, then the conventional switching functionality must be augmented to support these additional processes on the I/Q radio samples received from the BBU and radio units.

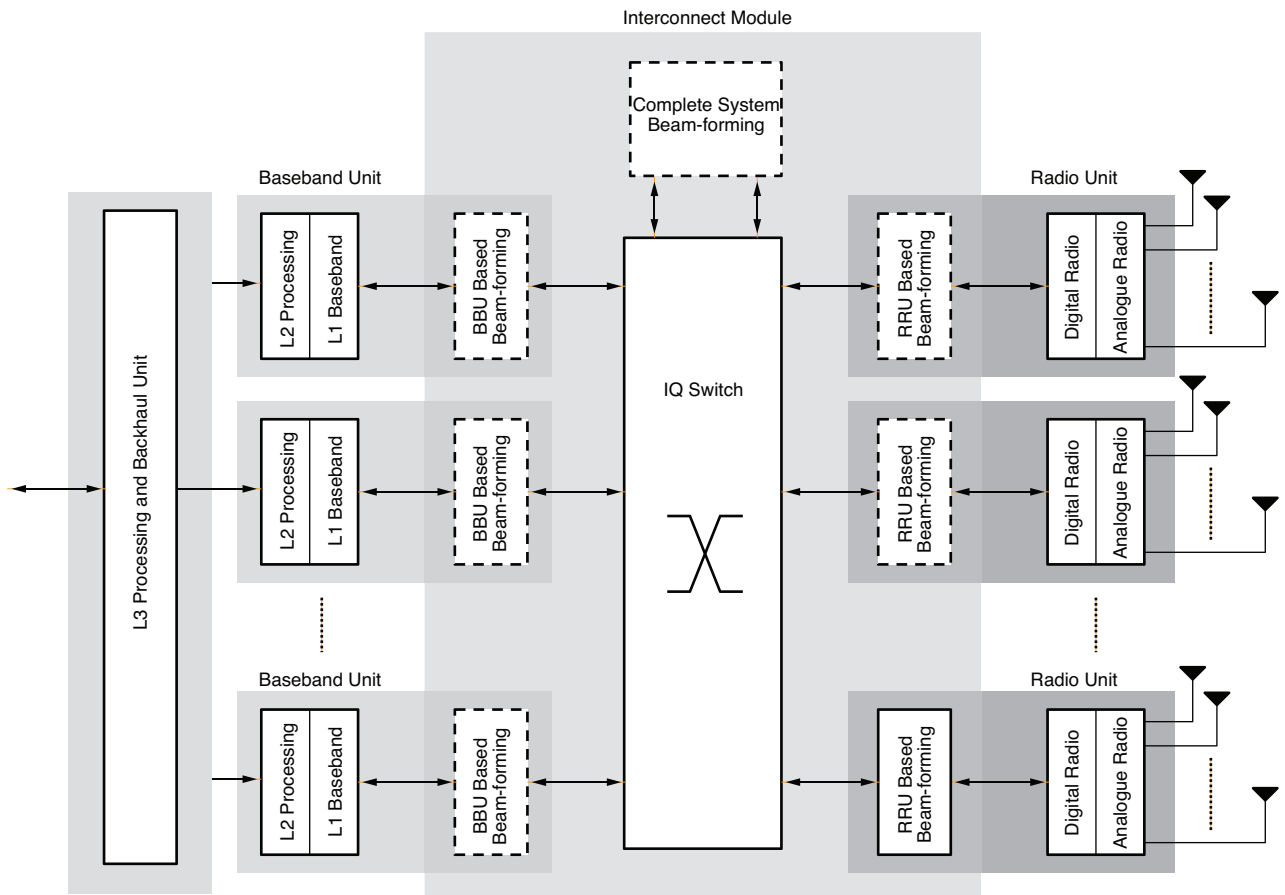
A high-level functional block diagram of a typical conventional AAA base station is given in [Figure 2, page 8](#). The figure concentrates on connectivity and beamforming functionality.

The I/Q switching function is similar to that used in a conventional macrocell (described in Section 4.1) and can be implemented either as part of the BBU or as a stand-alone interconnect module. Beamforming can be performed as part of the BBU, the central interconnect module, the RRU or potentially even distributed throughout several modules. The choice of location depends upon the system level architecture constraints and flexibility required for the application. In general if beamforming is limited to individual radio units then the AAA processing is most often performed within the radio itself, because this minimizes data transfer bandwidth between modules. Here the

beamforming functionality would normally be integrated into the digital radio FPGA. If, however, beamforming is performed across several radios then the AAA processing must be performed in a central location and hence is either implemented in the Interconnect Module or as part of the I/Q switching function within the BBU. The diagram shows all of the possible locations of the AAA processing functionality.

Integrated AAA System

In the integrated AAA system, all functionality — i.e., radio, baseband, higher layer, and backhaul processing — are integrated into the same physical unit. This unit is often co-located (or even integrated) into the antenna array itself. As a consequence, the system architecture is similar to that used in the small cell. In this configuration, the capacity and range supported are often lower than those of the conventional system, and deployment is normally targeted at dense urban areas. Due to the adoption of the small-cell-like architecture, internal connectivity is not a major consideration within the design of the integrated AAA system.



WP450_02_040814

Figure 2: Typical Conventional AAA Base-Station Architecture

CRAN Base Stations

The concept of CRAN is to centralize baseband, higher-layer, and backhaul processing in a single location that provides very high levels of processing capacity. A large number of radios are associated with the centralized processing unit and are distributed within the environment to provide widespread radio coverage.

Overview of CRAN Objectives and Architectures

The principal objectives of the CRAN architectures are:

- Reduce system costs and power consumption by exploiting the increased pooling efficiency that results from centralized core base-station processing
- Increase deployment flexibility using RRUs for macrocell and/or small-cell-like coverage in areas that are not well-served by a centralized macrocell approach — for example, in shadowed areas where in-fill coverage might be required
- Become an integral part of the HetNet strategy, providing additional deployment options well-suited to certain conditions

There are two somewhat different views on the implementation of CRAN architectures.

One approach is basically *evolutionary* in nature. CRAN is viewed as an extension of the distributed macrocell architecture, and the capacity of the centralized processing unit and the number of associated RRUs are both dramatically increased. However, the basic technologies and system architectures used in this approach are not radically different from those used in conventional base stations; the main difference is simply one of scale. The principal attraction of this approach is the *re-use* of existing system architecture and technology. These systems are sometimes referred to as *super macrocells* or *high-density base stations*.

The alternative approach is more revolutionary in nature. Here, in addition to the increase in capacity and coverage associated with CRAN, there is also a drive to use commodity technology to reduce system cost, increase flexibility, and provide a more open platform on which the base-station application can be built. The use of commodity technology targets the RRUs, the centralized processing unit, and the interconnect subsystem.

Consequently, even though both implementation approaches to the CRAN base-station architecture appear quite similar at the high level, the two architectures differ radically at the module level, because they are based on very different technologies.

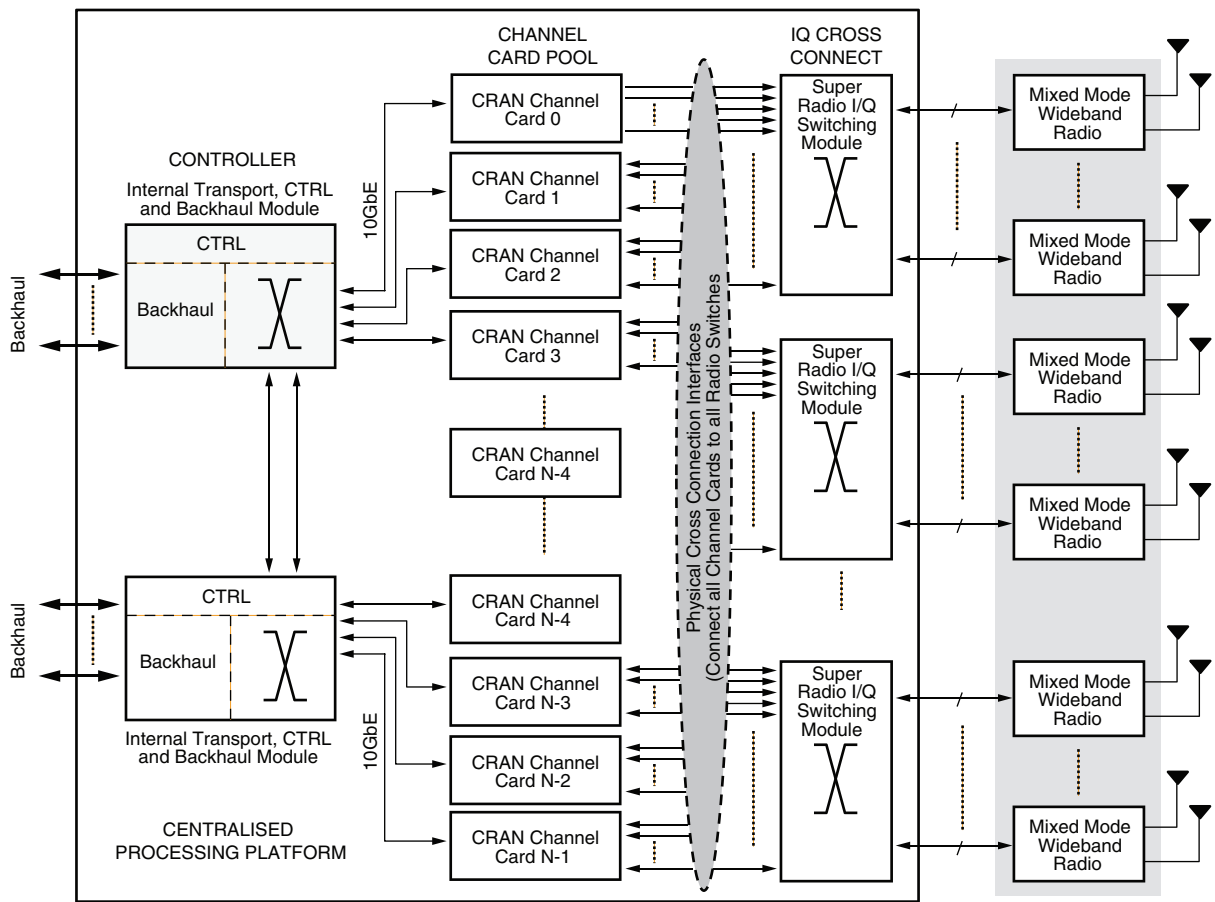
Overview of CRAN Applications

In the CRAN application, base-station connectivity is a major consideration, because the number of connections and the associated switching complexity are extremely high. Moreover, the system management functionality related to CRAN connectivity (i.e., delay compensation, signal routing, redundancy management, etc.) is also complex. Therefore, the principal challenges associated with CRAN connectivity are:

- CRAN requires an extremely high number of connections to many BBU processing elements and RRUs
- Distances between the RRUs and the centralized unit can vary significantly; in some instances, these distances could be tens of kilometers
- Challenges such as synchronization and delay compensation for a large number of radios must be addressed
- The centralized processing platform would, in all likelihood, be required to support both new and legacy radio and baseband equipment, and therefore be required to support various interconnect standards as well

A top-level functional block diagram of a possible CRAN base-station architecture is given below in Figure 3. The figure shows each of the functional units and highlights the importance of the I/Q interconnect module used to transfer the data between the baseband pool and the array of RRUs. In this instance, because the interconnect architecture is extensive, it might be implemented in several dedicated modules, distributing the functionality across the system.

The architecture of Figure 3 makes no specific assumptions about the technology around which the base station is designed; it is representative of *both* the evolutionary and revolutionary approaches.



WP450_03_040814

Figure 3: A Possible CRAN Base-Station Architecture

Wireless Base-Station Connectivity

The principal base-station interface types under consideration within this document are the radio I/Q sample interface and internal base-station data transport. The functional requirements of these interface types are somewhat different, and consequently different transport protocols are most often used.

Radio Sample Connectivity

The radio sample interface is used to transfer I/Q samples between the BBUs and radio units. Data transfer is full duplex with DL I/Q baseband samples being carried from the BBU to the radio transmitter and down-sampled UL I/Q samples from the radio receiver to the BBU. In general, the sample rate used is the modulated baseband data rate (or a low multiple of it); the resolution is in the range of 5-bit to 20-bit complex data. Raw data rates for various air-interface standards are shown in [Table 1](#).

Table 1: Raw Data Rates for Various Interface Standards

Air-Interface Standard	DL/UL	Sample Rate (MSPS)	Typical Resolution	Raw Data Rate per Antenna per Carrier (Mb/s)
WCDMA TX	DL	$f_s^{(1)}$	16 bits I/Q	122.88
	UL	$2f_s$	6 bits I/Q	92.16
LTE (5 MHz)	DL	$2f_s$	16 bits I/Q	245.76
	UL	$2f_s$	16 bits I/Q	245.76
LTE (20 MHz)	DL	$8f_s$	16 bits I/Q	983.04
	UL	$8f_s$	16 bits I/Q	983.04
CDMA2000	DL	1.2288	16 bits I/Q	39.3216
	UL	2×1.2288	6 bits I/Q	29.4912
GSM ⁽²⁾	DL	16×0.325	16 bits I/Q	166.4
	UL	2×0.325	16 bits I/Q	20.8

Notes:

- f_s is the WCDMA sampling rate and is defined as 3.84 MSPS.
- The “high” GSM sample rate (325 KSPS) is assumed in this example. GSM also supports the “normal” sample rate (270.833 KSPS).

Other important desired attributes of the radio sample interface:

- Stable, highly accurate clock supports local frequency and timing reference generation in remote radio applications
- Minimally low in latency, time-invariant, and deterministic for any configuration
- Delay measurement capability supports delay compensation in remote radio applications
- Control signaling support enables configuration and control of remote units

Standard interfaces have been defined to provide this functionality. The most commonly used today are *Common Public Radio Interface (CPRI)* [[Ref 5](#)] and *Open Base-Station Architecture Initiative (OBSAI)* [[Ref 6](#)] [[Ref 7](#)].

CPRI Overview

CPRI™ is specifically defined to carry radio I/Q data samples between baseband and local or remote radio units. It was initially specified to carry WCDMA data, but was later enhanced to support other standards, including LTE, GSM/EDGE, and WiMAX. CPRI uses a time-division multiplexing scheme to support independent data flows relating to individual carriers and antennas.

The frame structure also supports the transport of control data by reserving specific time slots; effectively, 1 in 16 CPRI data words is reserved for control information. Currently, CPRI supports data rates of between 614.4 Mb/s and 10.1376 Gb/s using the following line rates:

CPRI Line Rate 1:	614.4 Mb/s	CPRI Line Rate 2:	1.2288 Gb/s (2 x 614.4 Mb/s)
CPRI Line Rate 3:	2.4576 Gb/s (4 x 614.4 Mb/s)	CPRI Line Rate 4:	3.072 Gb/s (5 x 14.4 Mb/s)
CPRI Line Rate 5:	4.9152 Gb/s (8 x 614.4 Mb/s)	CPRI Line Rate 6:	6.144 Gb/s (10 x 614.4 Mb/s)
CPRI Line Rate 7:	9.8304 Gb/s (16 x 614.4 Mb/s)	CPRI Line Rate 8:	10.1376 Gb/s (16.5 x 614.4 Mb/s)

For line rates 1 through 7, CPRI employs 8B/10B line encoding to increase the robustness of the link. The proportion of the data rate dedicated to radio I/Q data samples, taking into account control channel and line encoding overhead, is $(15/16) \times (8/10) = 3/4$. For CPRI Rate 8, a line rate of 10.1376 Gb/s is achieved by a minor modification to the frame structure and the use of 64B/66B line coding in place of the 8B/10B scheme. This achieves an increase in I/Q sample throughput of approximately 25% over CPRI Rate 7.

Other features supported by CPRI:

- Frequency and time reference support:
 - Defined to enable high-quality frequency and timing references derived from CPRI link
- Accurate delay measurement capability:
 - Provides mechanisms for measuring and reporting round-trip delay (RTD) between transmitter and receiver to very high accuracies; used at the system level to compensate for apparent delays between BBUs and radio modules
- Transmission of various types of control and synchronization data:
 - Synchronization data for frame and time alignment (and to support RTD measurement)
 - L1 in-band signaling data for information that is directly related to the link (directly transported on the physical layer)
 - Control and management data⁽¹⁾ exchanged between the control and management entities, with two different Layer 2 protocols supported:
 - Subset of high-level data link control (HDLC)
 - Ethernet
 - Vendor-specific information to allow user-defined data transfer
 - Protocol extensions reserved for future specification

1. The definition used in CPRI is that control data is related to call processing functionality and management data is related to OA&M.

OBSAI Overview

The OBSAI protocol defines a set of common interfaces within the base station referred to as reference points (RPs). The RPs of relevance for transfer of radio I/Q samples are RP3 (connection to a local radio) and RP3-01 (connection to a remote radio). The RP3 protocol is designed to carry similar types of I/Q data to CPRI (i.e., WCDMA, LTE, GSM/EDGE and WiMAX); however, a fixed-length, packet-based structure is used. The message length is defined to be 19 bytes, 16 of which are dedicated to the payload. Messages are grouped into what are termed message groups (MGs); one message within the group is dedicated to control signaling. MGs are transmitted consecutively over the interface and grouped together to form a Master Frame. The Master Frame length is 10 ms, but the protocol allows for a wide range of MG sizes; thus, signaling overhead can be adapted to the system requirements. Like CPRI and OBSAI, RP3/RP3-01 employs 8B/10B line encoding. RP3/RP3-01 supports these line rates:

RP3x1:	768 Mb/s	RP3x2:	1.536 Gb/s
RP3x4:	3.072 Gb/s	RP3x8:	6.144 Gb/s

In general, OBSAI RP3/RP3-01 supports mechanisms similar to those of CPRI for synchronization, delay measurement, and control/management data transfer over the link.

Envisaged Future Directions

Currently, the maximum line rates supported by CPRI and OBSAI are 10.1376 Gb/s and 6.144 Gb/s, respectively. In the future, it is likely that increased rates will be required to support evolving system requirements. In addition, there is a drive for a greater level of standardization and for the use of lower-cost commodity technology. As a consequence, there are several possible paths that future interface evolution might take.

- Evolution to support higher throughputs is planned for the CPRI protocol. However, it is currently unclear which rates will be targeted, though 16.2 Gb/s, 19.6 Gb/s, and 25 Gb/s have been postulated as possibilities. Such rates can be achieved by minor modifications to the frame structure and appropriate choice of line-coding rates.
- A new standard for radio-sample interconnect, the open radio equipment interface (ORI), is also to be introduced. ORI provides greater interoperability between distributed elements within cellular base stations. It is designed to support GSM, UMTS, and LTE air-interface standards. The ORI specification covers the protocol layers that directly impact interoperability. Built upon the CPRI lower layers, ORI supports features similar to those of CPRI (I/Q transport, synchronization support, transport of C&M data, etc.). However, some options are removed while others are added to enable the interface to be fully interoperable.
- Another possibility under investigation is the use of 10 Gb Ethernet (10GbE) [Ref 10] to transport radio I/Q samples. This approach is attractive because Ethernet is widely perceived to be a commodity technology, offering the possibility of cost reduction. However, the standard 10GbE protocol does not support many of the features necessary for radio-sample interface applications. Consequently, the use of synchronous Ethernet (SyncE) [Ref 11] with IEEE Std 1588 [Ref 12] timing support is proposed to overcome these limitations.

Currently, there are no plans to evolve the OBSAI RP3/RP3-01 standard, and it is envisaged that its usage will decline in the coming years.

Internal Transport Connectivity

This interface (or set of interfaces) is used to transport user data and control and management information between base-station modules. Although information transfer could be implemented between any of the system modules, it is most often deployed between the baseband PHY, the L2 higher-layer processor, the L3/control processor, and backhaul units. Assuming an LTE base-station configuration, typical data types and characteristics are given in [Table 2](#).

It is important to note that the information in [Table 2](#) is intended only to provide a guide to the internal interface requirements. In practice, the actual data types and characteristics depend upon the detailed functional partition adopted — and in any particular base station, this might differ somewhat from the example presented here.

Table 2: Typical Internal Base Station Transport Data Types

Data Type	Data Transfer	Data Rates (Mb/s)	Latency Considerations	Comments
TrCH Packets (U-Plane & C-Plane)	MAC/PHY PHY/MAC	~350 Mb/s per 20 MHz LTE cell	Fast path, low latency	Approximate LTE DL user plus CTRL ⁽¹⁾ data rates, including protocol overhead
U-plane PDUs	Backhaul/L2 L2/backhaul	~350 Mb/s per 20 MHz LTE cell	Fast path, low latency	Approximate LTE DL U-plane data rates (including protocol overhead). In practice, this is somewhat less than TrCH packets carried over the MAC/PHY interface, because the L2 protocol adds some overhead
C-plane L2 CTRL data	L2/L3 L3/L2	~50 Mb/s per 20 MHz LTE cell	Slow path (latency less critical)	Approximate LTE C-plane data rates (including protocol overhead)
C-plane L3 CTRL data	RRM/L3 L3/RRM	~25 Mb/s per 20 MHz LTE cell	Slow path (latency less critical)	Approximate LTE C-plane data rates (including protocol overhead)
C-plane RRM CTRL data	Backhaul/RRM RRM/backhaul	~25 Mb/s per 20 MHz LTE cell	Slow path (latency less critical)	Approximate LTE C-plane data rates (including protocol overhead)
Backhaul management data	CTRL/backhaul backhaul/CTRL	~10 Mb/s per 20 MHz LTE cell	Management signaling (latency not critical)	Approximate data rate to support backhaul unit management
L3 management data	CTRL/L3 L3/CTRL	~10 Mb/s per 20 MHz LTE cell	Management signaling (latency not critical)	Approximate data rate to support L3 processor management
L2 management data	CTRL/L2 L2/CTRL	~10 Mb/s per 20 MHz LTE cell	Management signaling (latency not critical)	Approximate data rate to support L2 unit management
PHY management data	MAC/PHY PHY/MAC	~100 Mb/s per 20 MHz LTE cell	Fast path, low latency	Approximate data rate to support PHY management. It is assumed that management is performed by MAC, but could be performed by CTRL
Synchronization signaling	SYNC/other units other units/SYNC	~10 Mb/s per 20 MHz LTE cell	Sync signaling (latency less critical)	Transfer of timing reference data

Table 2: Typical Internal Base Station Transport Data Types (Cont'd)

Data Type	Data Transfer	Data Rates (Mb/s)	Latency Considerations	Comments
General OA&M signaling	CTRL/other units other units/CTRL	~ 10 Mb/s per 20 MHz LTE cell	Management signaling (latency not critical)	OA&M signaling for the system in general, including synchronization and other modules not specified in the system block diagram of Figure 1

Notes:

- CTRL denotes the system controller, which is assumed to perform overall system management and OA&M.

As shown in [Table 2](#), the requirements for internal data transport within the base station are varied and range from the critical fast-path transfers related to U-plane data to the less critical transfers related to OA&M signaling. For this reason, the architectures adopted for transport networks within base stations vary significantly. Moreover, it is common for more than one protocol to be used. These protocols find widespread use within wireless base stations:

- Gb Ethernet (GbE)** is often used for low-rate signaling paths for which latency is not critical. Many critical system components support this interface, and components such as GbE switches and PHYs are widely available at low cost. Moreover, it supports communication over a wide variety of physical media and transmission ranges, and is therefore well-suited for inter-module communication.
- 10 Gb Ethernet (10GbE)**, with or without KR physical layer support for transmission over a copper PCB backplane, is increasingly being used for critical low-latency, fast-path communication. This technology is well suited to this function, because the data rate supported is consistent with the throughput and latency requirements of typical macrocell base stations and its widespread adoption in communication systems means that it is well understood and widely available. Like all Ethernet variants 10GbE can support communication over a wide variety of physical media and transmission ranges and is well suited for inter module communication.
- Serial Rapid I/O (SRIO)** is also used for low-latency, fast-path communication, particularly within the BBU. SRIO is a flexible serial interface standard which is well suited to these requirements. It offers the possibility of very high throughput (5 Gb/s and 6.25 Gb/s per lane [v2] with low latency. The range of components which support SRIO is more limited than Ethernet but many components associated with wireless baseband processing have adopted it as their principal high-speed data interface. SRIO is less widely used on general-purpose processors (GPPs); hence, its utilization is often limited to the BBU application.
- Peripheral Component Interconnect Express (PCIe®)** is sometimes used for both slow- and fast-path communication within the base station. It is able to support high data rates (5 Gb/s per lane [v2] and 8 Gb/s [v3]) and low latency. Its use is less widespread in wireless base stations than either Ethernet or SRIO, and it is normally used to provide connectivity to a GPP.
- Processor local buses** have been used to provide intra-module communication between a host GPP and its peripherals. Such interfaces are used to support both slow- and fast-path communication. This type of architecture is now rarely used in new designs and has largely been superseded by serial interfaces and/or Ethernet-based interconnect systems.

- **Proprietary interfaces** can be used to provide slow- and fast-path communication within the base station when the system designer has control of both ends of the link. This is possible when the transport interface is implemented in an FPGA or ASIC. The proprietary interface can be implemented on either a parallel or serial link, and it can be designed specifically for the target application. This approach to interfacing can result in a very high degree of efficiency; design and validation costs, however, are typically quite high. For this reason, proprietary interfaces are normally not used for new designs, unless no standard interface can meet the application's technical requirements.

Even though it is common today to use multiple protocols to provide the range of interconnection required within the base station, the overall tendency is towards further harmonization. It is envisaged that the most likely interface to be used to cover the majority of internal connectivity requirements will be 10GbE, irrespective of the base station's architecture. It is also likely that other interfaces will be used, although reserved for highly specific functionalities.

Xilinx Solutions for Base-Station Connectivity

Xilinx has long recognized the importance of interconnect in wireless base-station design and has introduced silicon products and IP specifically to support this application. In this section, an overview of current generation silicon devices is provided with specific emphasis on devices suited to the base-station interconnect application. The section also provides a brief review of Xilinx and partner wireless connectivity-related IP, reference designs, and white papers.

7 Series Silicon Solutions

Xilinx offers a comprehensive range of All Programmable FPGA devices ideally suited to the wireless base-station application. The 7 series [Ref 13] comprises three broad classes of silicon solution: All Programmable FPGAs, 3D-ICs, and heterogeneous processing system-on-a-chip (SoC) technology. Traditionally, FPGA technology has been the basis for the base-station interconnect solution. However, FPGA-based SoCs and 3D-ICs offer additional features that can find application within evolved cellular base-station designs. In this white paper, examples of both *conventional FPGA-based* connectivity solutions and *evolved SoC-based* applications are explored.

Xilinx 7 series technology [Ref 14] offers a range of devices well suited to base-station connectivity and switching applications. The key attributes required for the application are:

- Availability of numerous serial transceivers that can be configured to support the wide range of physical interfaces typically adopted within wireless base stations
- High-performance programmable logic used to implement the interface protocol itself and any switching functionality associated with it
- Integrated block RAM used to support switching functionality, and delay compensation mechanisms that are often employed in high-capacity base stations
- High-performance SelectIO™ technology, supporting lower-speed interfaces potentially required by the application
- DSP capability supporting the signal-processing functionality that can be applied to the data to condition the signals prior to transfer between modules (e.g., bandwidth compression and beamforming).

The 7 series FPGAs comprise three complementary families that address a wide range of system requirements: the Artix®-7, Kintex®-7, and Virtex®-7 families. Various members of the Kintex-7 and Virtex-7 families are well adapted to wireless base-station interconnect. For low-end applications, the Kintex-7 XC7K325T and XC7K355T FPGAs offer (respectively) 16 and 24 serial transceivers supporting speeds up to 12.5 Gb/s with sufficient fabric resource to support the associated interfacing and switching functionality. For applications requiring support for up to 32 transceivers, other members of the Kintex-7 family (i.e., the XC7K420T and XC7K480T devices) are the more appropriate choice. The specific application examples presented in this white paper utilize the XC7VX415T, XC7VX690T, and XC7VX980T devices, which support (respectively) 48, 80, and 72 transceivers operating at data rates of up to 13.1 Gb/s.

The logic fabric in the Zynq®-7000 All Programmable SoCs [Ref 15] is essentially the same 28 nm FPGA programmable logic (PL) featured in the 7 series FPGAs. In addition, the Zynq-7000 All Programmable SoC products integrate a complete dual-core ARM® Cortex™-A9 processing system

(PS) with the 7 series PL fabric, enabling the implementation of *custom logic* in the PL and *custom software* in the PS. This integration within a single device makes possible the realization of unique and differentiated system functions that two-chip solutions (such as an ASSP and separate FPGA) simply *cannot match* due to the limited I/O bandwidth, latency, and power budget of the two-chip architecture.

In wireless intra-base-station applications that require the kind of processing functionality not normally associated with connectivity, Zynq-7000 AP SoC devices must be evaluated. For example, the application considered in this white paper, a low-end centralized AAA system, offers an ideal fit for a Xilinx Z-7045 AP SoC device.

As part of the 7 series, Xilinx has also introduced 3D IC technology [Ref 16]. This enables Xilinx to offer very high-capacity homogeneous 3D FPGA solutions and the possibility of heterogeneous products combining FPGA and complementary technologies. 3D IC technology is based on stacked silicon interconnect (SSI) technology, which uses a passive silicon interposer to enable high-bandwidth connectivity between multiple die by providing a large number of high-speed connections. Several devices of the Virtex-7 family are based on this technology, which enables the family to offer unprecedented FPGA capabilities.

In addition to integrating homogeneous FPGA solutions, SSI technology also enables the integration of multiple types of die. For example, the Virtex-7 XC7VH870T FPGA ties together three FPGA die as well as separate 28 Gb/s-capable transceiver circuits via the silicon interposer. Such technology has obvious applications in future high-capacity base-station connectivity, which typically requires numerous high-speed transceivers.

UltraScale Architecture Silicon Solutions

UltraScale FPGA devices [Ref 17][Ref 18] are based on 20 nm planar and 16 nm FinFET process technologies. They are designed to keep customers a generation ahead of the competition with an expansion of the Xilinx All Programmable offerings [Ref 19]. The UltraScale portfolio comprises All Programmable FPGAs, 2nd-generation SoCs, and 3D ICs. Each of these elements is designed to deliver an extra node worth of performance, power, and integration, and is ideally suited to meet the challenges associated with wireless base-station connectivity.

The UltraScale FPGA lays the foundation for the rest of the portfolio. The core FPGA combines the 20 nm/16 nm processes with a new set of design innovations. These next-generation devices give a 50 percent price/performance-per-watt improvement, twice the memory bandwidth, and the next generation of industry leading system optimized transceivers.

The 2nd generation All Programmable SoC features a multi-core heterogeneous processing architecture. Combined with increased bandwidth between the processing system and the programmable logic, this architecture delivers higher levels of programmable systems integration and performance at a fraction of the power consumption.

When FPGA technology is infused into All Programmable 3D ICs, the values are multiplied beyond what is typically expected from a generational improvement. By combining multiple die of different types, the 20/16 nm 3D IC provides integrated wide memory and up to 2X the capacity, system-level performance, and transceiver bandwidth when compared to the previous generation.

The key technological advances offered by Xilinx's UltraScale for base-station connectivity solutions are: increased gigabit transceiver capability throughout the product range; increased programmable performance to address the demands of evolving base-station architectures; and a significant reduction in power needs over previous generations.

More specifically, the UltraScale Kintex devices offer up to sixty-four 16 Gb/s serial transceivers and over 663k LUTs and 5,520 DSP slices — which means that many connectivity applications can be mapped to the Kintex family. For more complex applications, UltraScale Virtex devices offer an increased number of transceivers (up to 104), and logic densities of over 2.5M LUTs are supported. Within this white paper, all the UltraScale Kintex devices are considered as possible implementation targets for the specific applications examples investigated.

Connectivity IP

Xilinx and its partners offer a wide range of connectivity-related IP, reference designs, application notes, and white papers. These cover aspects from basic telecom interfaces, through switching, to baseband and higher layer acceleration. Much of the IP and associated documentation is available as part of the standard Xilinx IP catalogue; however, within this white paper reference is also made to application notes and reference designs which are not currently released; certain of these can be made available on request. Xilinx is committed to this market space and plans to enhance this IP product portfolio in the future.

Within this white paper, the architectures proposed and their associated resource requirements are based, wherever possible, on Xilinx and partner IP. However, some applications require the use of functionality not currently available as IP blocks. In such cases, Xilinx has estimated the required resources based on preliminary architecture proposals. Throughout the document, indications are given as to the source of each of the functional blocks.

Xilinx has a rich portfolio of IP that supports standard telecom interfaces. These include 10GbE MAC, tri-mode Ethernet MAC (10 Mb/s, 100 Mb/s, and 1 Gb/s), PCIe Gen 2 (integrated block in Kintex-7 and Virtex-7 FPGAs), CPRI v6.0, OBSAI RP3 v4.2, SRIO Gen 1.3 and SRIO Gen 2.1. These IP blocks, which are included within the standard Xilinx IP catalog, are fully supported, standards compliant, highly configurable, and designed to meet the requirements of a range of base-station connectivity applications. For details on these and other interfaces supported by Xilinx IP, go to www.xilinx.com.

Xilinx also has significant experience in the design of routing solutions for both packet processing and CPRI switching applications. Specific designs of relevance to this white paper are a lightweight Ethernet switch for CPRI control data and a CPRI I/Q data switch. These designs are not available as part of the IP catalogue, but further information can be provided on request.

Several of the advanced architecture examples examined within this white paper also make use of other types of IP available from Xilinx and partners. These generally perform the functions of hardware acceleration for baseband signal processing and/or higher-layer packet processing. Xilinx proposes a range of baseband IP for cellular applications. The examples used in this white paper are generally based on the LTE standard, for which Xilinx offers a wide range of IP blocks including [inverse] fast Fourier transform (iFFT/FFT), MIMO encoding/decoding, physical uplink shared channel (PUSCH) estimation, discrete Fourier transform (DFT), physical random access channel (PRACH) correlator, physical uplink control channel (PUCCH) receiver, and several forward error

correction (FEC) cores. Xilinx also offers a more limited set of IP for other air-interface standards, which are in general dedicated to FEC applications. All of this IP is available as part of the standard Xilinx IP catalog.

Higher-layer acceleration IP is offered by Xilinx partners. Specific examples used within this white paper include: IEEE Std 1588 v2, available from IPClock, Ltd.; IPsec; and other cryptographic components, including SNOW 3G, AES, ZUC, and Kasumi; and Robust Header Compression (RoHC). All are available from Elliptic Technologies Inc. Both IPClock, Ltd. and Elliptic Technologies, Inc. are part of the Xilinx Alliance Partner Program.

In addition to application-related IP, Xilinx also offers a range of infrastructure IP used extensively within the architecture proposals of this document. This includes AXI4, AXI4-Lite, and AXI4-Stream interfaces and interconnect, and the MicroBlaze™ embedded processor, which is optionally used for control and configuration of the overall application.

FPGA Connectivity Architectures for Base Stations

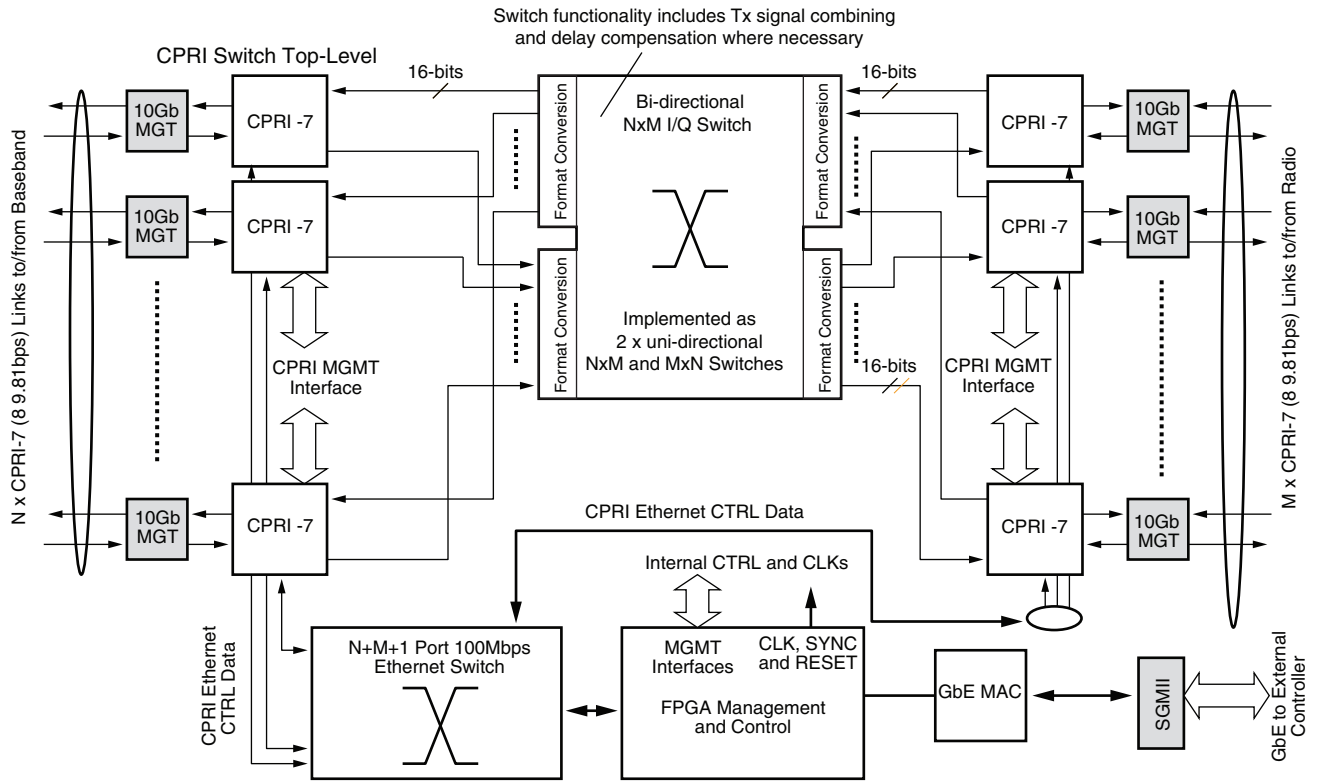
FPGA technology is ideally suited to meet the wide range of connectivity requirements associated with wireless base stations. In this section, several practical applications are considered. Those under consideration include basic radio I/Q sample connectivity, advanced I/Q sample connectivity, base-station internal transport, and integrated I/Q and internal transport networks.

Basic Radio Sample Connectivity

Radio sample connectivity is one of the principal connectivity problems to be addressed within wireless base stations. The problem is that the dimensionality and physical location of the radio and baseband units is often quite different. This gives rise to the general requirement to connect any radio to any baseband processor. This requirement is well established in WCDMA, because softer hand-off must often be supported. However, in other forward-looking applications such as multi-mode base stations that might support both new and legacy equipment, high-density cell sites, and LTE-A systems capable of supporting carrier aggregation and CoMP, the requirement becomes even more critical. The connectivity problem is compounded still further by the fact that the configuration generally varies from one cell site to the next, and might even change dynamically over time. This implies that not only must it be possible to interconnect a large number of radio and baseband modules, but also that the interconnect must be highly flexible.

The solution generally adopted is that radios and baseband processors are connected to a centralized switching hub, which performs the necessary signal routing; this is in fact one of the primary functions of the interconnect module depicted in [Figure 1](#). Of course, the switching function need not be implemented only in a stand-alone module; it could be integrated within other module types as well (e.g., system controller or baseband units).

Assuming that the radio sample interface is implemented using the CPRI protocol, the functionality typically performed by an FPGA dedicated to radio sample interfacing is shown in [Figure 4](#).



WP450_04_040814

Figure 4: Typical Radio Sample Interconnect Functionality Implemented with FPGA Technology

The functionality can be summarized as follows:

Radio Sample Interconnect

- N x CPRI interfaces to baseband modules (or components)
 - CPRI Rates 2 to 7
 - CPRI master or slave, Ethernet support, delay measurement support and AXI Lite CTRL interface
- M x CPRI interfaces to radio modules
 - CPRI Rates 2 to 7
 - CPRI master, Ethernet support, delay measurement support and AXI Lite CTRL interface
- Maximum I/Q data rate per CPRI sample path = $3.84 \text{ (CPRI MSPS)} \times 15 \text{ (CPRI I/Q words)} \times 128 \text{ (bits per word)} = 7.3728 \text{ Gb/s}$
- Embedded Ethernet connection from each CPRI end-point
 - Maximum of $N+M$ Ethernet streams at a data rate of typically $<100 \text{ Mb/s}$
- Resource estimates are derived from Xilinx CPRI IP core v7.0

I/Q Switching Functionality

- TDM switching between N baseband and M radio units providing bidirectional connection between baseband and radio, and potentially baseband loopback functionality.
- I/Q data format conversion from CPRI to the format used in the switches.
- TX signal-combining (baseband to radio).
- TX/RX delay compensation support through all switches. The amount of delay compensation to be implemented is dependent upon distance between the baseband and radio modules. The delay (in numbers of CPRI frames per kilometer) is approximately 16. Hence, per kilometer, approximately $16 \times 128 \times 15$ bits must be stored for each CPRI link (or approximately $1 \times 36K$ block RAM per kilometer per CPRI link, for each direction).
- Resource estimates are derived from Xilinx Internal CPRI I/Q Switch reference design.

Embedded Ethernet Connection Switching Functionality

- Ethernet switching support for CPRI radio and baseband interfaces plus one local master port for external connection or local termination (embedded or external controller, if required)
- A maximum of $N+M+1$ ports at < 100 Mb/s per port (maximum aggregated throughput of < 1 Gb/s over four ports)
- Resource estimates are derived from the Xilinx Internal Lightweight CPRI Embedded Ethernet Switch reference design

Internal Control and Interconnect Functionality

- Internal control used for configuration and management of the switching and interface blocks. This can be implemented using a hardware state machine or a MicroBlaze processor (resource estimates assume use of MicroBlaze processor IP core)
- Optional 1GbE connection to external controller (GbE MAC plus SGMII PHY)
- Internal interconnect based on AXI4 crossbar providing connection from the embedded controller to the control interface of each CPRI interface and switching block
- Resource estimates are derived from the MicroBlaze processor and AXI4 IP infrastructure

A top-level resource utilization breakdown is shown in [Table 3](#) for an 8×8 CPRI switch processor, which can be used as part of a baseband or interconnect module.

Table 3: Resource Utilization for an 8x8 CPRI Processor

Logical Block	FF	LUT	18K Block RAM	DSP	Primary Clock (MHz)	Block Functionality Assumed
CPRI BBU and RRU interfaces (IP v7.0)	53,776	41,744	64	0	245.76	9.83 Gb/s, master functionality, Ethernet support, delay measurement support, AXI Lite CTRL interface
CPRI switch and associated functionality	36,608	19,886	144	0	245.76	BBU-to-RRU and RRU-to-BBU switching plus signal addition on BBU-to-RRU interface. 32-bit data interface. Total throughput per port supported 15.728 Gb/s (2x basic data rate). FPGA operating frequency: 245.76 MHz.
Delay compensation for TX and RX path for eight radio connections	2,240	2,240	256	0	245.76	Delay compensation for up to 8 km for both TX and RX paths
CPRI Ethernet Functionality	20,550	14,236	87	0	245.76	Seventeen-port light-weight Ethernet switch, GbE MAC, and 1000Base-X SGMII
FPGA Management and CTRL	7,700	7,700	40	0	156	Implementation in the MicroBlaze processor with PLB, DMA, and AXI-Lite
AXI Interconnect	1,200	3,700	0	0	156	AXI-Lite interconnect to all managed components
Total:	122,074	89,506	591	0		

The clocking scheme proposed is that all CPRI-related functionality (interfaces, I/Q switch, and CPRI Ethernet switch) operates at 245.76 MHz, control and internal interconnect at 156 MHz, and GbE functionality at 125 MHz. It should be noted, however, that for high-complexity applications, logic and block RAM resource reduction can be achieved if the CPRI I/Q switching functionality is configured to operate at 491.52 MHz. This is highlighted in Table 4. The table shows device mapping for both 7 series and UltraScale devices for a range of applications, from a low-capacity internal RRU CPRI processor to a full CRAN I/Q cross-connect switch.

Resource utilization for the high-complexity 32x32 CPRI processor application is shown for switch clock speeds of both 245.76 MHz and 491.52 MHz. As can be seen, the 7 series XC7VX690T device is over-mapped at a switch clock of 245.76 MHz, but does have reasonable implementation margin if a clock of 491.52 MHz is used. In UltraScale devices, the use of the 491.76 MHz switch clock enables the use of the XCKU095 device in place of the XCKU115 device with a significant increase in implementation margin.

Table 4: Device Mapping for a Range of CPRI Switch Applications

CPRI Processor Application	Device Mapping	Resource Utilization					Comments
		FF	LUT	18K Block RAM	DSP	Transceivers	
2x2 CPRI processor in radio module	Radio Processing FPGA	20k	16.5k	54	0%	3	Up to 3 external CPRI Rate-7 connections and one internal connection to the local radio. 4 CPRI ports, 2x2 full duplex CPRI switch, 4-Port 100 Mb/s Ethernet Switch. Internal controller and external Ethernet port not included. Signal addition functionality not required.
8x8 CPRI processor baseband module	XC7K325T	30%	44%	66%	0%	100%	16 CPRI Rate-7 ports, 8 x 8 full-duplex CPRI switch, 17-Port 100 Mb/s Ethernet switch, embedded soft controller, and external 1GbE connection supported.
	XCKU035	30%	44%	55%	0%	100%	
12x12 CPRI processor baseband/interconnect module	XC7K355T	50%	68%	72%	0%	100%	24 CPRI Rate-7 ports, 12 x 12 full-duplex CPRI switch, 25-port 100 Mb/s Ethernet switch, embedded soft controller, and external 1GbE connection supported.
	XCKU060	33%	46%	48%	0%	75%	
16x16 CPRI processor baseband/interconnect module	XC7K420T	55%	76%	83%	0%	100%	32 CPRI Rate-7 ports, 16x16 full duplex CPRI switch, 33-Port 100 Mb/s Ethernet Switch, embedded soft controller, and external 1GbE connection supported.
	XCKU060	43%	60%	64%	0%	100%	
32x32 CPRI processor CRAN I/Q cross-connect module. Switch clock 245.76 MHz	XC7VX690T	90%	116%	127%	0%	80%	64 CPRI Rate-7 ports, 32x32 full duplex CPRI switch, 65-Port 100 Mb/s Ethernet Switch, embedded soft controller and external 1GbE connection supported
	XCKU115	59%	76%	86%	0%	100%	
32x32 CPRI Processor CRAN I/Q cross-connect module. Switch clock 491.52 MHz	XC7VX690T	56%	84%	73%	0%	80%	64 CPRI Rate-7 ports, 32x32 full duplex CPRI switch, 65-Port 100 Mb/s Ethernet Switch, embedded soft controller, and external 1GbE connection supported.
	XCKU095	44%	66%	64%	0%	100%	

For the radio module 2x2 CPRI processor, no specific device is identified. This is because the CPRI functionality is assumed to be integrated within the digital radio FPGA solution. Consequently, the target device is primarily dependent upon the characteristics of the radio application rather than of the CPRI interface. For this reason, absolute resource requirements for any specific device are given rather than percentage of utilization.

Advanced Radio Sample Connectivity

In this section, the basic application of radio I/Q sample connectivity is extended to address additional functionality and advanced base-station architectures. The specific applications considered comprise I/Q data compression, the integration of baseband and higher-layer acceleration into the connectivity processor, AAA processing, and connectivity for CRAN base stations.

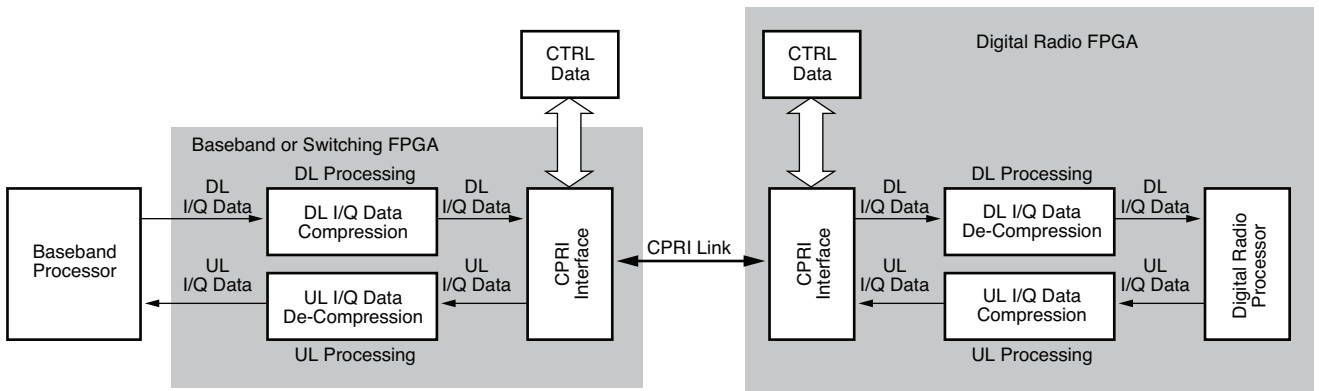
I/Q Data Compression

The fundamental motivation for I/Q data compression is to increase the number of antenna-carrier data streams that can be supported over the radio/baseband link. This is important in order to support the wider bandwidths and the increased number of antennas associated with advanced cellular applications. Numerous compression schemes have been proposed [Ref 21] [Ref 22] [Ref 23] [Ref 24]. A small number of representative techniques are used to demonstrate the application in this white paper.

Overview

In general, data rate reduction can be achieved both by reducing the transport protocol overhead and by applying compression schemes to the user data. Within this document, I/Q user data compression is considered. The algorithms applied are quite different than those commonly associated with data transport, and are often based on signal processing rather than on data or packet processing techniques. Data compression can be considered as an optional enhancement to the transport protocol and can be applied to any transport mechanism. In this white paper, CPRI is used as a concrete example, but the methodology can be extended to other protocols.

Data compression can be applied to UL and DL signals on all common air-interface standards. For the DL, data compression is applied at the BBU prior to formatting for the transmission protocol. The data is then formatted and transmitted over the link. At the RRU, the data is de-formatted and then decompressed to regenerate the original signal. The compression process performed on the UL is in general the reverse of that of the DL. Each of these processes is depicted for a single CPRI link in Figure 5.



WP450_05_040814

Figure 5: Baseband and Radio Module Supporting CPRI I/Q Data Compression

Even though compression can be performed for all air-interface standards on both UL and DL, the characteristics of the signal types are in general quite different — e.g., noise levels, dynamic range, over-sampling rates, etc. This means that compression techniques that might work well for certain signals could be less effective for others. In this white paper, I/Q compression for LTE is considered, and techniques specifically adapted to this waveform are presented.

Algorithm Description

Data compression can be subdivided into lossless and lossy schemes. By definition, lossless schemes conserve the characteristics of the source data so that signal fidelity is not degraded. However, the compression ratio increase achieved could be quite modest and it is often difficult to ensure a constant bit rate and get a meaningful deterministic latency measurement. In contrast, lossy schemes are able to achieve increased compression ratios, but at the cost of signal quality degradation. Many of the practical schemes proposed comprise a combination of lossless and lossy algorithms, the objective being to find an optimum compromise between acceptable signal quality and the compression ratio achieved.

In this white paper, the implementation of two LTE specific schemes is considered:

- Time-domain resampling plus quantization resolution reduction
- Frequency-domain data transport with block floating point (BFP) signal representation

These schemes achieve compression by exploiting the oversampling implicit within the LTE protocol and by optimizing the quantization resolution.

Time Domain Resampling plus Quantization Resolution Reduction

In this technique, resampling filters are added to the input/output of the BBU and RRU signal chains in order to condition the signals prior to multiplexing onto (and demultiplexing from) the CPRI frame. On the RRU, the overall architecture changes little from conventional implementations; the only difference is that radio up- and down-sampling rates must be adjusted to comply with the used bandwidth rather than with the oversampled bandwidth. On the BBU, TX and RX resampling filters must be implemented to convert from the reduced sample rates used over CPRI to the conventional oversampled rates defined in LTE and required by the baseband processor. At the BBU, this is new functionality which could be implemented within the CPRI switch FPGA.

The best-case compression ratio that can be achieved is 75/128 for LTE bandwidths of between 5 MHz and 20 MHz. This assumes “ideal” resampling filters — which cannot, of course, be implemented in practice. In effect, practical resampling filters cause increasing signal quality degradation as the filter bandwidths approach the bandwidth of the signal. Consequently, even though time-domain resampling can be considered a lossless scheme in theory, significant signal distortion is apparent in practice for high compression ratios. Within this work, it is assumed that resampling filters are implemented to achieve a compression ratio of 3/4, which provides a reasonable compromise between signal fidelity and compression performance. The resampling filters are 123-tap FIR structures with a normalized passband cut-off of 0.1 and stopband at 0.167. In addition to 3/4 resampling, a reduction of quantization resolution is also applied from a nominal 16-bit I/Q resolution to 12 bits. This results in a combined compression ratio of 9/16.

Because the sampling rate is no longer an integer number of CPRI frames, the resulting data might not pack conveniently into the CPRI format. To overcome this, it is necessary for the user to have low-level control over the way in which I/Q words are packed into the frame. This is a requirement common to many of the compression techniques under consideration, and in practice the implementation adopted must provide significant flexibility in how data is formatted.

Frequency Domain Data Transport with Block Floating Point Signal Representation

In OFDM systems such as LTE, the basic data to be transmitted is initially QAM-modulated onto a set of orthogonal subcarriers. This frequency-domain data is then converted to the time domain by an iFFT and the cyclic prefix (CP) added. At the receiver, the reverse of this process is performed — i.e., the CP is removed, the data converted back to the frequency domain by an FFT, and then demodulated on a subcarrier basis. Because only data related to the used subcarriers is required, the frequency-domain representation requires significantly fewer samples than the time-domain representation. Hence, if data is transferred between the RRU and BBU using the frequency domain, then significant reductions in data rates can be achieved. For LTE bandwidths of between 5 MHz and 20 MHz, an ideal compression ratio of 35/64 can be achieved for an equivalent number of quantization bits.

Frequency domain data transport requires that the conventional system functional partition between BBU and RRU be modified such that iFFT and FFT processing are performed in the RRU rather than in the BBU. The iFFT/FFT could be conveniently implemented in the digital radio FPGA. With the exception of the removal of the iFFT/FFT, the BBU processing remains unchanged.

Unfortunately, in LTE, the repartition of iFFT/FFT functionality between BBU and RRU is not necessarily straightforward. Three issues that do *not* arise in the conventional partition must be addressed here:

- First, because the RRU must perform the iFFT/FFT processing to construct the time/frequency domain waveforms, it must be aware of the TX and RX symbol timing to a high degree of precision. This is, however, not a significant problem, because CPRI provides highly accurate round-trip delay (RTD) measurement capability, which is used to compensate for delays within the system. This mechanism can also be used to compute TX/RX symbol timing at the radio.
- Second, in LTE, the format of the PRACH is quite different from that of other DL and UL channels; the time/frequency resource division and basic slot structure are unique to the PRACH. This means that the FFT preprocessing used for other UL channels is not appropriate for the PRACH; consequently, the PRACH must be treated as a special case. One solution is to apply time-domain resampling to the PRACH (which occupies only a fraction of the band of the other channels) and send this in parallel with the FFT data. Here a small proportion of the CPRI I/Q transport capacity would be dedicated to PRACH transmission. This solution is relatively efficient, particularly if the subcarriers dedicated to the PRACH are not sent as part of the FFT data, and overall efficiency is degraded only marginally from ideal values.
- Third, on the DL, user channels, control channels, and reference signals can be allocated different power levels. These power differentials must be taken into account when performing the iFFT. Consequently, such information must be made available at the radio.

Within this work, the target compression ratio for frequency-domain data transfer is 5/8, which is relaxed from the ideal value of 35/64 principally in order to simplify data packing onto the CPRI frame. In addition, a BFP signal representation is proposed in which all data related to a single LTE physical layer resource block (RB) is encoded using a BFP representation with either 12 or 8 bits I/Q resolution. This results in a best-case combined compression ratio of 5/16 (not taking into account the PRACH).

Algorithm Performance

A summary of compression algorithm performance is provided in Table 5. The table summarizes compression ratio (CR), signal-to-noise ratio (SNR), error vector magnitude (EVM), and latency performance of the two schemes under consideration.

EVM is a measure of the difference between an ideal waveform and the actual waveform, and is defined for LTE in Equation 1:

$$EVM\% = \sqrt{\frac{E[|\bar{x} - x|^2]}{E[|x|^2]}} \times 100\% \tag{Equation 1}$$

where \bar{x} (x-bar) is the actual measured waveform and x is the ideal signal. Strictly speaking, EVM is defined only for DL signals in 3GPP, but a similar measure can be used for UL signals where the ideal signal is defined as the received signal prior to quantization and compression.

SNR is not defined in 3GPP as a specific measurement that must be applied to base-station or UE signals. It is, however, a useful measure of signal quality because it can be directly related to asymptotic demodulation performance and, therefore, gives a good intuitive understanding of overall system performance. Signal-to-noise ratio (SNR) in dB is defined in Equation 2:

$$SNR(dB) = 10 \times \log_{10} \left[\frac{E[|x|^2]}{E[|\bar{x} - x|^2]} \right] \tag{Equation 2}$$

Results are given for a 20 MHz LTE system in which all subcarriers are occupied. The results relate to *best-case* and *worst-case* conditions. *Best case* is defined to be that case where all users are received at the same power level; *worst case* is defined to be that case where a single user operates at the minimum acceptable power level *and all other users* operate at equal power, assuming a received signal dynamic range of 30 dB.

Table 5: Compression Scheme Performance

Compression Scheme Description	CR	Best-Case Performance		Worst-Case Performance		Latency	Comments
		SNR	EVM	SNR	EVM		
Rate 3/4; 12-bit time-domain resampling	0.56	52.2 dB	0.30%	26.13 dB	7.25%	Fraction of OFDM Symbol	Asymptotic 50 dB at 12 bits
Frequency domain 12-bit BFP quantization	0.46	72 dB	0.025%	72 dB	0.025%	One OFDM symbol	BFP overhead is 1 scaling factors per 24 data samples
Frequency domain 8-bit BFP quantization	0.31	48 dB	0.39%	48 dB	0.39%	One OFDM symbol	

Algorithm Implementation

FPGA technology provides the ideal platform on which CPRI processing, including compression, can be implemented. It provides the flexibility required to optimize data mapping onto the CPRI frame, and resource types compatible with a range of compression algorithms. A possible implementation for CPRI processing including data compression is shown in Figure 6. The figure shows generic processing that could be part of an RRU digital radio FPGA, a BBU interface FPGA, or a stand-alone CPRI switch.




Figure 6: CPRI Processor Implementation Top-Level Block Diagram

CPRI processing comprises CPRI termination for one or more connections, I/Q and Ethernet switching, and I/Q data compression. In the example of Figure 6, compression/decompression is performed directly on the samples sent to (and received from) the CPRI interface block. Consequently, the I/Q data at the switch boundary is uncompressed, and switching is performed on uncompressed data. This approach has the advantage that the basic switch functionality is relatively simple and is similar to that described in the Basic Radio Sample Connectivity section. There are, however, drawbacks associated with this arrangement:

- The number of switch ports and overall switch throughput increases due to the fact that more CPRI sector/carrier I/Q data streams (referred to as AxCs in CPRI terminology) can be supported on each CPRI link.

- In instances where the module performs only data switching functionality, the received samples must initially be decompressed, then switched, then recompressed, and finally transmitted to the destination. Here, the processes of decompression and recompression perform no useful function other than simplifying switch design.

The alternative approach is to perform switching on compressed data. In this case, if the CPRI processor is used to simply route data from one module to another, no decompression and recompression is performed, and the data is switched directly from source to destination. In this case, the processes of compression and decompression are performed only at the source and destination modules. This reduces switching data rates and simplifies overall system design, although the design of the switch itself could be more complex. In practice, such an approach is feasible provided that some limitations are imposed on data to be switched: for example, it might be necessary to limit word lengths to certain discrete values.

The specific processing performed by the compression block is dependent upon which algorithm is employed, and typically comprises optional data buffering (though *required* only for certain algorithms), compression/decompression, data formatting, and mapping to the CPRI frame. In many compression schemes, the processes of compression and decompression are similar in nature. For others, however, they are quite different; in some such cases, decompression is not even required.

The implementation characteristics of the compression algorithms under consideration are given in [Table 6](#). The table provides a description of the processing performed, along with FPGA resource estimates for a single instance of a compression/decompression block assuming a CPRI data rate of 9.81 Gb/s (i.e., CPRI Rate-7). The resource estimates, which are derived from preliminary architecture proposals, apply to the complete compression/decompression functionality including associated data buffering and formatting for the CPRI interface. In addition, the table also gives the number of CPRI 20 MHz AxCs and associated sample rates that can be supported *after* compression using standard CPRI procedures for generation of stuffing bits and reserved words. The compression/decompression blocks are dimensioned to support the quoted number of AxCs. For reference, a single CPRI Rate-7 link without compression can support seven 20 MHz AxCs, assuming that I/Q data is encoded on 16 bits; this quantity can be increased to eight if 15-bit I/Q encoding is used.

Table 6: Implementation Description and Resource Estimates for Compression Techniques

Compression Scheme Description	CPRI AxCs	Max I/Q Sample Rate (MSPS)	Implementation Description	Implementation Complexity		
				LUT FF	Block RAM (18K)	DSP
Time domain 12-bit resampling. Resampling rate 3/4.	12 ⁽¹⁾	368.64	Compression: Perform time-domain fractional rate resampling (decimation filtering), round and saturate to 12-bits. Decompression: Perform time-domain fractional rate resampling (interpolation filtering)	3,700	0	84
Frequency domain 12-bit BFP quantization	16	491.52	Compression: Perform at radio unit on RX data samples. Store data block, remove CP, compute FFT on data block, compute scaling factor for each RB of each symbol, apply scaling factors to all RB elements, round and saturate to 16/12 bits, and format data block (add scale factor word). Plus PRACH DDC. Decompression: Perform at radio unit on TX data samples. Derive scaling factor for data block, multiply received data samples within block by scaling factor, store data block, compute IFFT, add CP to IFFT output samples, and round and saturate to the required number of bits for the DUC.	26,300	440	86
Frequency domain 8-bit BFP quantization	24	737.28		39,000	618	134

Notes:

1. In practice, up to 13 AxCs can be supported using the algorithm configuration under consideration in this white paper. However, 12 is used in the examples, as it is a more convenient number from the cellular-system viewpoint.

The time-domain resampling technique’s logic and block RAM utilization is relatively low. However, DSP resource utilization is high, because a large number of filter taps are required to meet the sharp cutoff requirements (123 taps are assumed in this example). Resampling filter complexity can be reduced by trading CR and/or EVM performance against implementation complexity. For example, if 7/8 resampling is implemented, then the resource utilization can effectively be reduced by a factor of two, with an associated reduction in CR from 0.56 to 0.66.

The frequency-domain data transfer resource utilization appears to be very high. It is important, however, to bear several points in mind with respect to the viability of this technique:

- When viewed from the system level, the overall processing requirements are unchanged from the conventional system; in effect, functionality has simply been migrated from the BBU to the radio unit with no new functionality introduced. This is quite different from other techniques, where compression/decompression functionality must be implemented at both ends of the link, resulting in an overall increase in resource requirements.
- The frequency-domain processing itself must only be performed at the radio so the full compression/decompression algorithm is relevant only to the radio application; no such processing is performed in the BBU or stand-alone switch. Radio units are generally served by a relatively low number of CPRI links, which means that the number of compression/decompression blocks is low. Moreover, the radio switching application is generally limited to CPRI daisy-chaining and redundancy management. Consequently, if switching is performed on compressed data, most often only a single instance of the compression/decompression block is required. This contrasts with other applications which, in general, might operate on a large number of CPRI links, thus requiring a significant number of compression/decompression blocks.

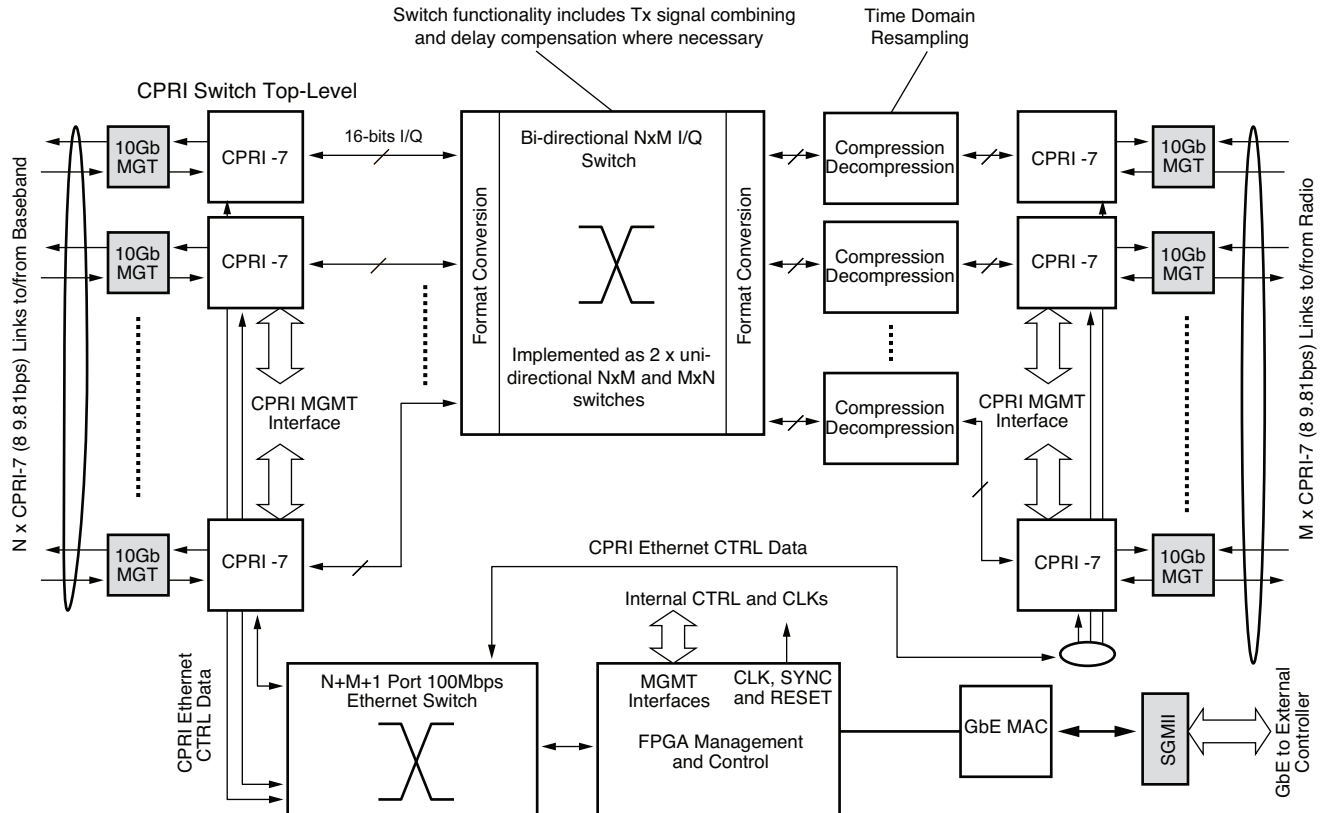
- The number of AxCs serviced is potentially very high, because the CR that can be achieved using frequency-domain data transfer is significant. For example, assuming frequency-domain data transfer with 12-bit BFP, the number of 20 MHz AxCs that can be supported is 16. This is sufficient to support four 20 MHz LTE carriers on four antennas, which represents a relatively complex radio application. Hence, the resource requirement for compression and decompression could be modest compared with the digital radio itself.

System Level Application

In this section, two system level application examples are considered:

- A BBU switching module CPRI processor plus I/Q data compression block, based on time-domain resampling plus quantization resolution reduction
- A radio CPRI processor plus I/Q data compression block, based on frequency-domain data transport with BFP signal representation

The architecture for the baseband CPRI processor is shown in Figure 7. The functionality is identical to that described for the basic I/Q processor in the Basic Radio Sample Connectivity section, with the exception that a compression/decompression block is added to each of the *M* CPRI links to the radio modules. In this implementation, compression is performed only on the radio links, and the switch operates on uncompressed data. The compression/decompression operations are performed using the time-domain resampling algorithms outlined in the Algorithm Description section.



WP450_07_040814

Figure 7: Baseband CPRI Processor with Data Compression Functionality

A top-level resource utilization breakdown is shown in Table 7 for an $N=12$, $M=12$ CPRI Processor plus data compression block. The compression block is dimensioned to support a maximum of twelve 20 MHz AxCs per CPRI link. Consequently, each of the twelve connections to the radio and baseband modules can support up to three-sector operation with 80 MHz of used system bandwidth.

Table 7: Resource Utilization: 12x12 CPRI Processor plus Radio Sample I/Q Compression

Logical Block	FF	LUT	18K Block RAM	DSP	Primary Clock (MHz)	Block Functionality Assumed
CPRI BBU and RRU interfaces (IP v7.0)	80,664	62,616	96	0	245.76	9.83 Gb/s, master functionality, Ethernet support, delay measurement support, AXI Lite CTRL interface
CPRI switch and associated functionality	99,952	52,434	420	0	245.76	BBU-to-RRU and RRU-to-BBU switching plus signal addition on BBU-to-RRU interface. 32-bit data interface. Total throughput per port supported 15.728 Gb/s (2x basic data rate). FPGA operating frequency: 245.76 MHz
Delay Compensation for TX and RX path for 8 radio connections	3,360	3,360	384	0	245.76	Delay compensation for up to 8 km for both TX and RX paths
CPRI Ethernet functionality	26,550	14,236	87	0	245.76	25-port light-weight Ethernet switch, GbE MAC, and 1000Base-X SGMII
Radio I/Q sample compression/decompression	44,400	44,400	0	1,008	368.64	Compression implemented as 3/4 123-tap resampling filter plus quantization to 12 bits I/Q. Decompression implemented as a 4/3 164-tap resampling filter
FPGA management and CTRL	7,700	7,700	40	0	156	Implementation in MicroBlaze processor with PLB, DMA, and AXI-Lite
AXI interconnect	2,220	11,100	0	0	156	AXI-Lite interconnect to all managed components
Total	264,846	195,846	1,027	1,008		

The principal impact of the addition of the compression/decompression block on the basic 12x12 CPRI processor (device mapping for which is given in Table 4) is that the LUT utilization increases by ~22% and the DSP utilization increases from 0 to 1,008 blocks. Nevertheless, the design still maps well to Kintex-7 XC7K355T and UltraScale Kintex XCKU060 devices with the resource utilization shown in Table 8.

Table 8: Device Utilization: 12x12 CPRI Processor plus Radio I/Q Sample Compression

Device Type	FF	LUT	18K Block RAM	DSP	Transceivers
XC7K355T	59%	88%	72%	70%	100%
XCKU060	39%	59%	48%	37%	75%

Essentially, the I/Q compression functionality can be considered as complementary functionality to CPRI processing, because it makes significant use of the otherwise unused DSP blocks; the impact on logic utilization is relatively well constrained.

The architecture for the radio CPRI processor plus I/Q compression functionality is shown in Figure 8. The compression/decompression block is connected directly to the digital radio processor. Consequently, the CPRI switch must operate on compressed CPRI data. This means that switch design and the packing of data into the CPRI frame is likely to be more complex than in a

conventional application. However, the numerology chosen for the compression block (as outlined in the [Algorithm Description](#) section) ensures that the complexity increase is marginal.

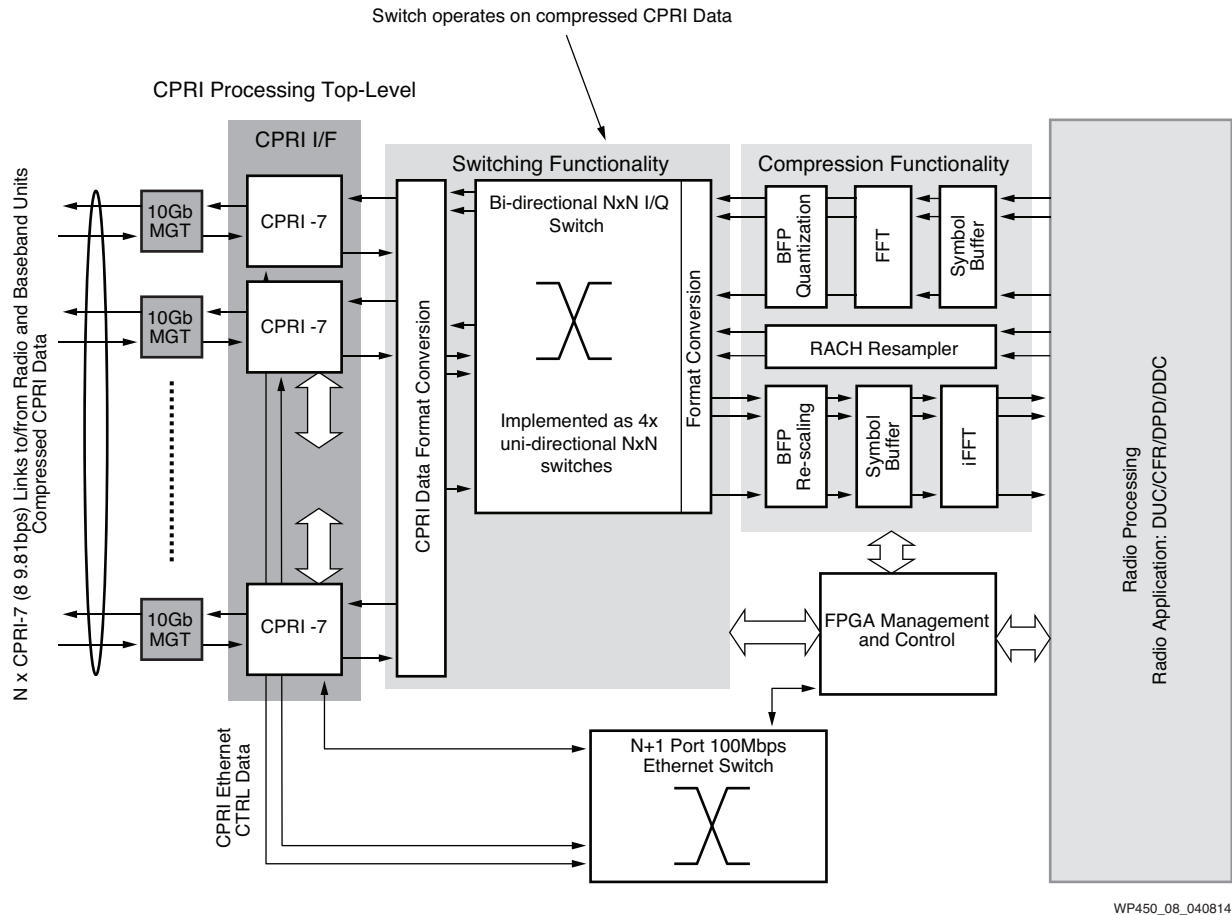


Figure 8: Radio CPRI processor with data compression functionality

The compression/decompression functionality comprises five basic processes:

- **UL OFDM demodulation:** radio sample buffering, cyclic prefix removal, and FFT processing.
- **Block floating point encoding of each of the UL RBs:** computation of the scaling factor and application of this to the data block.
- **PRACH resampling:** channel extraction and down-sampling by 24, 12, or 6 for LTE bandwidths of 20 MHz, 10 MHz, and 5 MHz respectively.
- **DL block floating point decoding of each of the DL RBs:** determination of the rescaling factor and application of this to the data block.
- **DL OFDM modulation:** QAM data buffering, iFFT processing, and cyclic prefix addition.

Other functionality (e.g., CPRI interface, CPRI Ethernet switching, and FPGA control) is similar to that performed in a conventional radio CPRI processing block.

The resource utilization of a 2x2 CPRI switch processor plus I/Q compression block is listed in [Table 9](#). The application assumes that there are three CPRI Rate-7 connections to the radio. One connection is used to transfer I/Q data to the local digital radio, one is used either for redundancy

or to increase CPRI bandwidth, and one is used for daisy-chaining to another radio unit. In all cases, it is assumed that a maximum of sixteen 20 MHz AxCs are used by the local radio. This represents a relatively large configuration that can support up to three-sector operation over 100 MHz of used system bandwidth.

The compression/decompression functionality is dimensioned in accordance with the maximum number of AxCs supported, and comprises a single instance of the sixteen-AxC frequency-domain processor plus BFP encoder described in the [Algorithm Implementation](#) section.

Table 9: Resource Utilization for a 2x2 CPRI Processor plus I/Q Sample Compression

Logical Block	FF	LUT	18K Block RAM	DSP	Primary Clock (MHz)	Block Functionality Assumed
CPRI v7.0 Baseband Interface	10,083	7,827	12	0	245.76	9.83 Gb/s, master functionality, Ethernet support, delay measurement support, AXI Lite CTRL interface.
CPRI Switch and associated functionality	1,904	1,258	12	0	245.76	BBU-to-RRU and RRU-to-BBU switching plus signal addition on BBU-to-RRU interface. 32-bit data interface. Total throughput per port supported 15.728 Gb/s (2x basic data rate). FPGA operating frequency: 245.76 MHz.
CPRI Ethernet Functionality	6,800	4,800	28	0	245.76	4-port light-weight Ethernet switch, GbE MAC, and 1000Base-X SGMII.
Radio I/Q sample compression/decompression	26,300	26,300	440	86	245.76	Compression implemented as DL iFFT, UL FFT, PRACH resampler and BFP processor.
AXI Interconnect	745	2,200	0	0	156	AXI-Lite interconnect to all managed components.
Total	45,832	42,385	492	86		

In this implementation, no resource is included for TX/RX delay adjustment, because it is assumed that coarse adjustment is performed in the baseband or switching module, and that fine adjustment is performed in the digital radio itself. Moreover, no resource is included for the controller, because it is assumed that this is included as part of the digital radio design.

For this application, no specific device is identified for implementation of the design. This is because CPRI and compression functionality are assumed to be integrated within the digital radio FPGA solution. In this case, device selection is normally made as a function of the characteristics of the digital radio solution, which tends to dominate the overall implementation.

Hardware Acceleration

A possible enhancement to the connectivity architecture that might be appropriate for certain system configurations is *implementation of a hardware acceleration functionality* within the connectivity processor. The primary motivations for such an approach are:

- **To enable efficient implementation of highly complex functions** that require the high levels of parallel processing only offered by a hardware-based solution.
- **To increase overall system capacity, throughput, and/or performance** by offloading functions performed by heavily loaded system components.

- **To increase efficiency of the system level solution.** Many functions are more efficiently implemented in a central location, rather than in a distributed fashion where the processing is divided over several interconnected components. The connectivity processor is ideally situated to provide this type of functionality, because it typically has access to data from numerous BBUs and RRUs.
- **To enhance evolutionary possibilities of the system.** Allocation of FPGA resources to certain functions allows the possibility of functional upgrade to support new functionality and/or improve performance. This could be problematic if such functionality is performed using other technologies.

The types of acceleration that can be envisaged include baseband signal processing, higher-layer and transport acceleration, and preprocessing for radio signals. Baseband signal processing acceleration can be envisaged for all radio air-interface types, and is partially applicable to advanced system configurations and functionality requiring significant computational complexity. The following baseband processing functions might be candidates for hardware acceleration:

- **LTE baseband processing:** OFDM modulation/demodulation (iFFT/FFT), MIMO decoding for high-order MIMO applications, complete DL PHY processing, PRACH correlation, PUCCH, and SRS processing
- **WCDMA processing:** DL scrambling; spreading and power weighting; DL MIMO processing; UL correlation and Rake receiver implementation; and PRACH correlation
- **GSM:** TX modulation and RX equalization
- **Multi-mode encoding/decoding acceleration for GPP-based CRAN applications:** Convolutional encoding/decoding, Turbo encoding/decoding, Reed-Muller encoding/decoding, etc.

Higher-layer and transport processing is conventionally performed in GPP or network processor technology. However, such technology is not necessarily well suited to all functions performed at the higher layer, and hardware acceleration is often required. The functions suited to hardware implementation include:

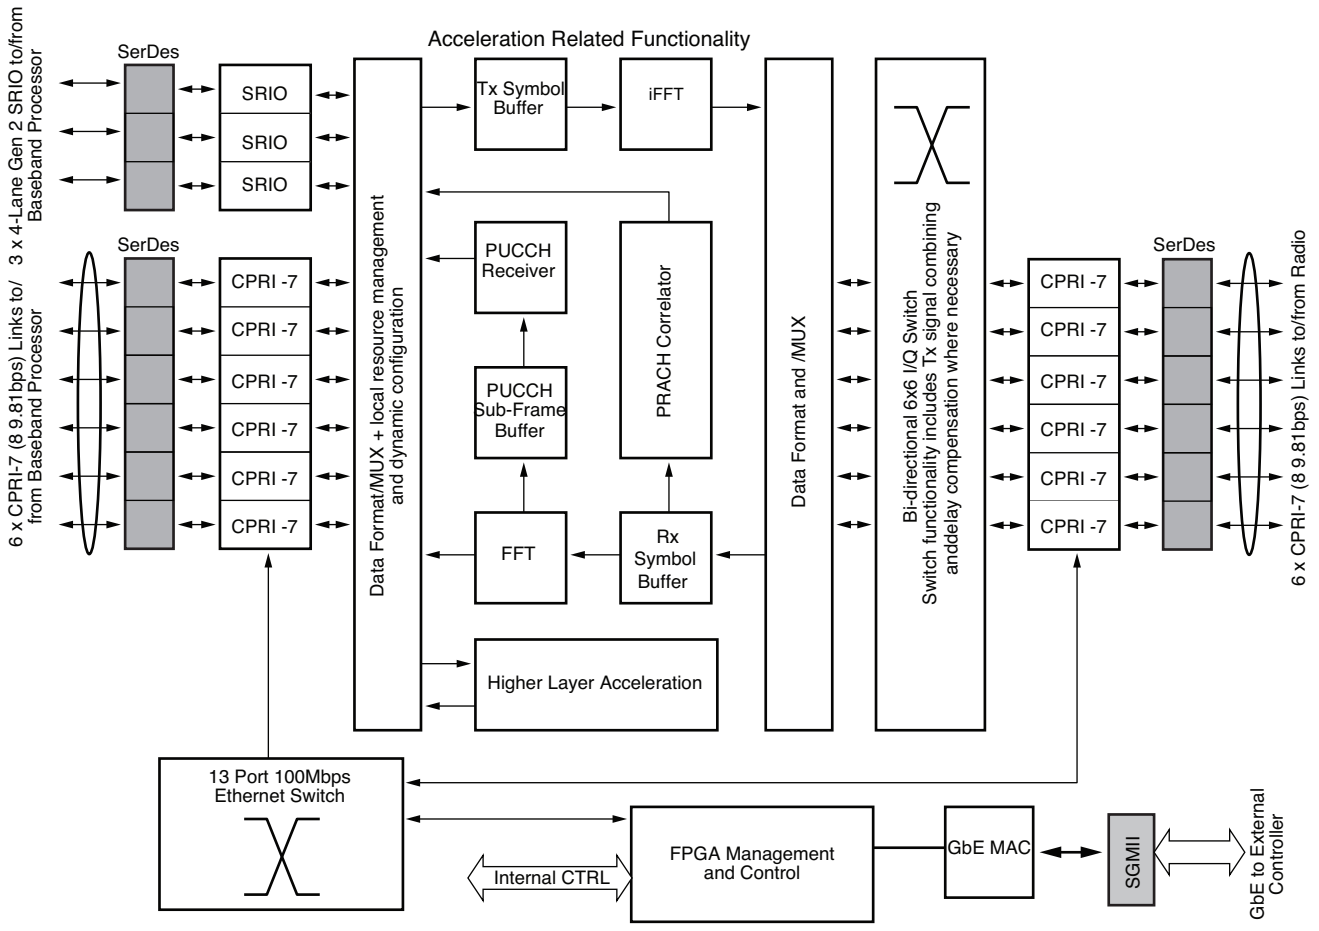
- **Security applications:** SNOW3G, AES, ZUC, and network security (IPSec and MACSec)
- **RoHC acceleration**
- **Synchronization:** GNSS support, network-driven timing support: IEEE Std 1888-2008/PTP SyncE, local time/frequency reference generation [\[Ref 25\]](#)

Preprocessing and postprocessing of radio signals is becoming increasingly important as systems evolve to support a more diverse range of applications. The types of functionality envisaged include:

- **Power management:** TX/RX power measurement and control
- **AAA processing:** TX/RX path calibration, adaptive beamforming and signal combining, and beamforming coefficient derivation

In this section, baseband and higher-layer acceleration is considered. (AAA processing is described in the [AAA Processing](#) section.) This type of acceleration is of primary relevance to the BBU switch application; consequently, this example is used to illustrate the implementation possibilities. A

possible switch/accelerator architecture is shown in Figure 9. The system-level application assumed is an LTE baseband module connected to up to six radio units and performing L1 plus higher-layer processing functionality for a three-sector single carrier 4x4 20 MHz LTE system.



WP450_09_040814

Figure 9: Possible Architecture for H/W Acceleration Implementation within a Baseband I/Q Switch FPGA

The functionality performed by the FPGA can be summarized as follows:

I/Q Sample Interconnect and Switching

- **Six CPRI interfaces to radio modules and six CPRI interfaces to baseband components:** Rates 2 to 7 supported, with core configuration as described in the [Basic Radio Sample Connectivity](#) section. In practice, to support the configuration under consideration, CPRI Rate 3 would be sufficient (2.4576 Gb/s); however, for consistency, the interfaces and switching functionality are dimensioned to support any CPRI Rate up to 7.
- **I/Q switching between the six baseband and six radio units,** providing bidirectional connection between the radio and baseband components
- **Ethernet switching support for CPRI radio and baseband interfaces,** plus one local master port for external connection or local termination (embedded or external controller if required)

Hardware Acceleration Functionality

- **Baseband acceleration:** DL iFFT and UL FFT, PUCCH receiver, and PRACH correlator
- **Higher layer acceleration:** IPSec and RoHC
- All acceleration functionality is dimensioned to support three-sector 20 MHz 4x4 LTE operation.
- The baseband and higher-layer functionality under consideration is used as an example to illustrate the types of functions that can be accelerated, and is not exclusive; many other types of function can also benefit from hardware implementation.
- These particular functions are chosen because they have efficient FPGA implementations but require significant processing resource in a software-based implementation.

Internal Control and Interconnect Functionality

- **Three 4-lane Gen 2 SRIO connections** to the baseband processing elements (one per sector) carrying user data and dynamic configuration parameters for the accelerator blocks.
- **Internal FPGA controller** used for configuration and management of the acceleration, switching, and interface blocks. This could be implemented using a hardware state machine or a MicroBlaze embedded processor (resource estimates assume MicroBlaze processor is used).
- **Optional 1GbE connection** to external controller (GbE MAC plus SGMII PHY).
- **Internal interconnect** based on AXI4 crossbar providing connection from the embedded controller to the control interface of each accelerator, interface, and switching block.

The resource utilization for the I/Q switch/accelerator application is shown in [Table 10](#). The resource estimates are derived from the following sources:

- **CPRI functionality:** Xilinx CPRI IP v7.0 and Xilinx internal reference designs for the I/Q and Lightweight CPRI Ethernet switches
- **Hardware acceleration functionality:** Baseband accelerator resource estimates are derived from standard Xilinx LTE baseband IP available as part of the IP catalogue. Resource estimates for higher layer acceleration are derived from Xilinx partner IP (Elliptic Technologies Inc.)
- **FPGA management and internal interconnect:** Xilinx MicroBlaze processor and AXI4 IP infrastructure

Table 10: Resource Utilization: 6x6 CPRI Processor plus Baseband and Higher Layer Acceleration Functionality

Logical Block	FF	LUT	18K Block RAM	DSP	Primary Clock (MHz)	Block Functionality Assumed
CPRI BBU and RRU interfaces (IP v7.0)	40,332	31,308	48	0	245.76	9.83 Gb/s, master functionality, Ethernet support, delay measurement support, AXI Lite CTRL interface.
CPRI switch and associated functionality	22,160	12,258	84	0	245.76	BBU-to-RRU and RRU-to-BBU switching plus signal addition on BBU-to-RRU interface. 32-bit data interface. Total throughput per port supported 15.728 Gb/s (2x basic data rate). FPGA operating frequency: 245.76 MHz.
Delay compensation for TX and RX path for six radio connections	1,680	1,680	192	0	245.76	Delay compensation for up to 8 km for both TX and RX paths.
CPRI Ethernet functionality	15,000	9,500	67	0	245.76	17-port light-weight Ethernet switch, GbE MAC, and 1000Base-X SGMII.
SRIO interfaces	21,900	19,350	24	0	250	SRIO Gen 2 (spec rev 2.2) x4 at 5.0 Gb/s per lane.
iFFT plus buffers	10,200	10,200	75	27	307.2	3 x Radix-4 4-channel FFT. Dimensioned for 20 MHz 4x4 three-sector operation
FFT plus buffers	11,400	11,400	75	27	307.2	3 x Radix-4 4-channel FFT. Dimensioned for 20 MHz 4x4 three-sector operation.
PRACH correlator	38,034	29,103	102	174	245.76	Three-sector, 32-root sequences per sector, 2 ms latency and four antennas per sector.
PUCCH receiver plus sub-frame buffer	6,400	6,400	44	38	307.2	Three-sector processing with 100 format 1 and 50 format 2 users per sector.
IPSec accelerator	12,800	12,800	12	0	245.76	IPSec throughput ~1 Gb/s for UL and DL.
RoHC accelerator	20,000	20,000	12	0	245.76	RoHC throughput ~60 Mb/s (compressed header data) for UL and DL.
FPGA management and CTRL	7,700	7,700	40	0	156	Implementation in MicroBlaze processor with PLB, DMA, and AXI-Lite.
CTRL interconnect	1,200	3,700	0	0	156	AXI-Lite interconnect to all managed components.
Total	208,806	175,399	775	266		

Such an application maps well to the Xilinx Kintex-7 XC7K355T and the UltraScale Kintex XCKU060 devices, both of which have sufficient numbers of GTX transceivers to service the application. Device mapping is summarized in [Table 11](#):

Table 11: Device Utilization: 6x6 CPRI Processor plus Baseband and Higher Layer Acceleration Functionality

Device Type	FF	LUT	18K Block RAM	DSP	Transceivers
XC7K355T	47%	79%	54%	18%	100%
XCKU060	31%	53%	36%	10%	75%

AAA Processing

In general, an AAA base station comprises one or more radio units each supporting a large number of antennas (e.g., eight, sixteen, or more) connected to a pool of BBUs dimensioned to support the appropriate number of carriers and sectors. Due to the large number of data streams, radio sample connectivity and switching are key problems to be addressed within this architecture. If AAA processing is performed centrally within the base station over several radio modules, then this functionality is often integrated into the connectivity processor within the BBU. In this case, the connectivity solution must be enhanced, not only to switch data streams between radios and BBUs, but also to combine signals such that the required beamforming can be implemented.

One possible architecture for a centralized switch/beamformer FPGA based on CPRI interconnect is shown in Figure 10.

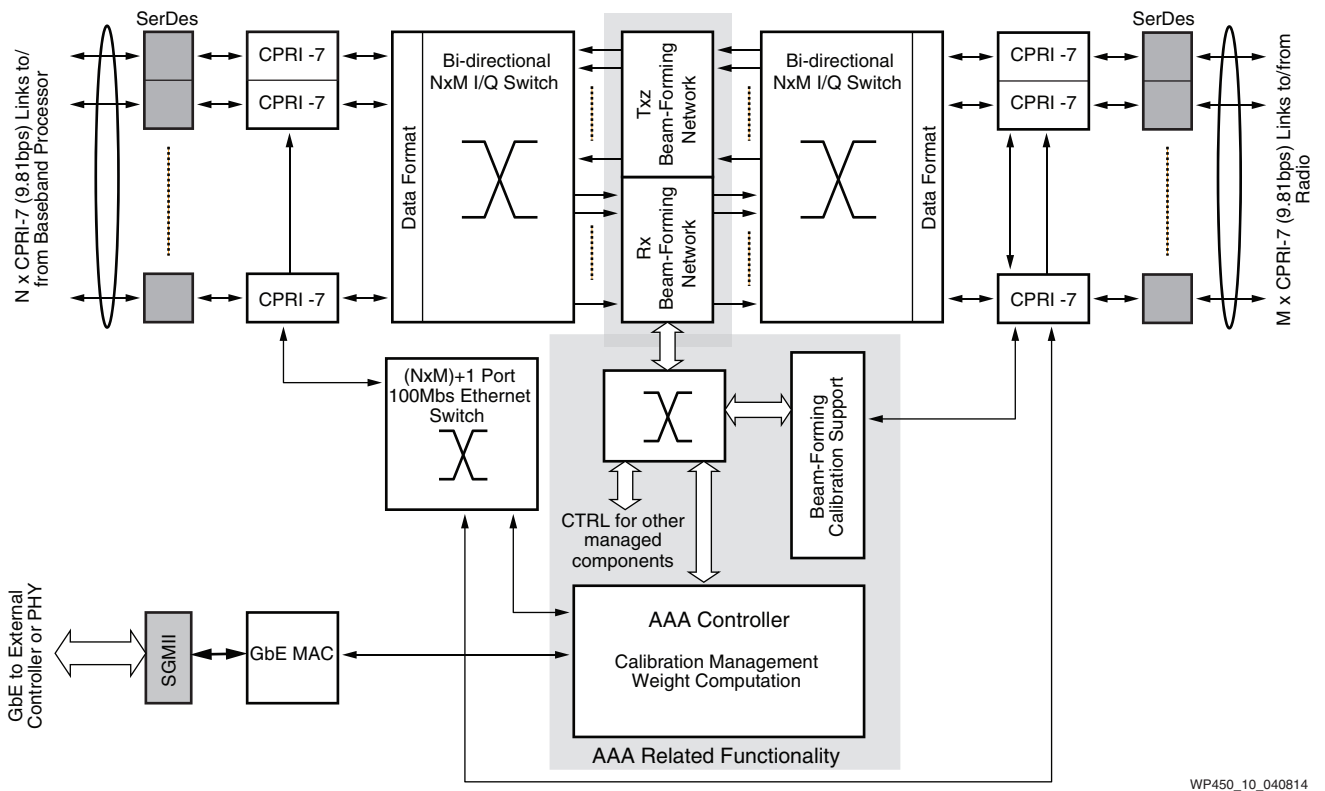


Figure 10: Possible Centralized Switch/Beamformer FPGA Top Level Architecture

The CPRI interface and switch functionality (both I/Q samples and Ethernet) are similar to that described for the radio sample processor in the Basic Radio Sample Connectivity section. However, in this application the demultiplexed AxC data related to both the BBU and RRU are made available to the beamforming networks so that the data streams can be combined as dictated by the AAA controller. As a consequence, the switching function must be modified such that data written into the switch is initially demultiplexed, organized into groups of data samples to be combined, and then sent to the appropriate port of the beamforming network. The beamforming results for each data stream must be written back to the switch and then assigned to the appropriate CPRI transmit AxC. The solution shown in Figure 10 uses two I/Q switches, one at the input and one at the output of the beamforming network, to offer the complete flexibility to switch any input stream to any output stream. If some restrictions are placed on the system configuration, it might be possible to

combine the functionality of the two switches into a single unit, reducing resource utilization. The fundamental functionality of the switch itself is not changed significantly from that described in the [Basic Radio Sample Connectivity](#) section. Switching latency is, however, marginally increased due to the beamforming process.

The additional functionality associated with AAA processing:

- **TX and RX beamforming networks:** These perform complex vector-matrix multiplication of the input data streams with the beamforming weights. The order of the weight matrix is dependent upon the number of baseband data streams P and the number of antennas Q , and is $P \times Q$ or $Q \times P$, depending upon the precise definition of the baseband and radio-sample vectors.

One P or Q vector result for each of the required AAA combinations must be computed at the input sample rate.⁽¹⁾ The beamforming networks are implemented in FPGA fabric and make significant use of the DSP48 blocks.

- **Beamforming hardware calibration support:** This typically comprises functions such as calibration sequence generation (PN sequence generation), TX and RX channel estimation (correlation of the transmitted and received calibration sequences), and power measurement. Again, significant use of the available DSP48 blocks is made.
- **AAA controller:** The controller is responsible for computing the adaptive weights of the beamforming networks and the management of the AAA hardware blocks. The beamforming weights are computed from the channel estimates derived by the calibration function, and are typically updated at a relatively low rate (perhaps several times per second). Such processing is often implemented in a GPP, because computation load is relatively low and flexibility is of principal concern. The GPP functionality could be embedded in the FPGA logic using either a MicroBlaze CPU IP core if conventional FPGA technology is selected, or in the ARM Cortex-A9 processor if Zynq technology is selected.

The complexity of the application depends mainly on the order of the AAA application and the number of channels supported. Resource estimates are provided in [Table 12](#) for an 8x8 CPRI configuration (eight RRUs and eight BBUs), assuming that each connection is configured to support eight 20 MHz data channels. This means that a total of sixty-four 20 MHz channels are supported over the eight data streams.

1. Typical sample rates are given in [Table 1](#).

Table 12: Resource Utilization for an 8x8 CPRI Processor plus AAA application

Logical Block	FF	LUT	18K Block RAM	DSP	Primary Clock (MHz)	Block Functionality Assumed
CPRI v7.0 BBU & RRU Interfaces (v7.0)	53,776	41,744	64	0	245.76	9.83 Gb/s, master functionality, Ethernet support, delay measurement support, AXI Lite CTRL interface.
CPRI Switch and associated functionality	73,216	39,772	288	0	245.76	BBU-to-RRU and RRU-to-BBU switching plus signal addition on BBU-to-RRU interface. 32-bit data interface. Total throughput per port supported 15.728 Gb/s (2x basic data rate). FPGA operating frequency: 245.76 MHz.
Delay Compensation for TX and RX path for 8 radio connections	2,240	2,240	256	0	245.76	Delay compensation for up to 8 km for both TX and RX paths.
CPRI Ethernet Functionality	20,550	14,236	87	0	245.76	17-port light-weight Ethernet switch, GbE MAC, and 1000Base-X SGMII.
TX Beamforming Network	3,200	3,200	32	160	491.52	8x8 beamforming for TX signals.
RX Beamforming Network	3,450	3,450	32	160	491.52	8x8 beamforming for RX signals.
Calibration Hardware Support	3,450	3,450	32	64	245.76	PN sequence generation, correlation, and power meters.
AAA Controller and FPGA Management	8,400	8,400	40	8	156	Implementation in MicroBlaze processor with PLB, DMA, FPU, and AXI-Lite.
CTRL Interconnect	9,200	11,700	0	0	156	AXI-Lite interconnect to all managed components plus AXI-streaming switch.
Total	177,482	128,192	831	392		

It is further assumed that the AAA application is limited to 8x8 functionality (i.e., eight baseband data streams combined over eight antennas). This means that eight TX and RX beamforming networks must be implemented in parallel. The resource utilization breakdown for the entire application is shown in Table 12. The AAA application resource estimates are derived from preliminary Xilinx architecture proposals.

The resource estimates of Table 12 are given assuming the target device uses Kintex-7 technology and that the AAA controller is implemented in logic fabric as a MicroBlaze embedded processor. If a Zynq-7000 device is used, the controller could be implemented in the ARM Cortex-A9 CPU, and the fabric resources given for the controller would be reduced to zero. If UltraScale Kintex technology is chosen, then the number of resources required for the beamforming functionality could be reduced, due to the more efficient implementation of the complex MAC associated with the UltraScale DSP architecture. Essentially, the number of DSPs could be reduced to a total of 192 for TX and RX beamforming, and the number of logic resources reduced by approximately 40%. In all cases, it is assumed that the beamforming network is implemented with a clock frequency of 491.52 MHz, which is viewed as readily achievable for this type of functionality.

The specific devices that could be considered for this implementation are the XC7K325T, the Z-7045, and the XCKU035, for which the resource utilization is as follows:

Table 13: Device Utilization: 8x8 CPRI Processor plus AAA Application

Device Type	FF	LUT	18K Block RAM	DSP	Transceivers
XC7K325T	40%	56%	58%	27%	100%
XCZ7045	39%	53%	73%	43%	100%
XCKU035	42%	62%	77%	16%	100%

CRAN Connectivity

Two different methodologies for the implementation of CRAN base stations are described in the [CRAN Base Stations](#) section; these are referred to as the *evolutionary* and the *revolutionary* approaches. In the evolutionary approach, the base-station architecture and basic technology is similar to that of the distributed macrocell. However, the base station is dimensioned to support a significant increase in radio, baseband, and higher-layer processing capacity. In the revolutionary approach, a similar increase in capacity is targeted, but in addition, the widespread use of commodity technology is proposed to reduce system cost, increase flexibility, and provide a more open hardware platform.

Evolutionary CRAN Base-Station Connectivity

In the evolutionary approach, the major difference between the conventional and CRAN base station is simply one of scale. The CRAN centralized processing unit is dimensioned to support a large number of high-capacity BBUs and RRUs interconnected via a flexible switching network. The implication of the capacity increase on the connectivity architecture is that interconnect becomes so complex that it must often be implemented in several parallel dedicated modules rather than in a single switch module. In addition, it might also be necessary to support several levels of interconnect and switching. Hence, a second level of switching could be required within the BBUs.

An example CRAN or high-density base-station architecture is depicted in [Figure 3](#), where a first level of switching is performed in the parallel *I/Q cross-connect modules*, which are responsible for switching I/Q data between the radio and baseband modules. Each BBU is dimensioned to support very high capacity, which is implemented within a pool of modem processors. Hence, a second level of switching is implemented within the BBU to switch I/Q data to/from individual modems.

At both levels, the functionality to be performed by the interconnect unit is similar to that described in [Basic Radio Sample Connectivity](#), and comprises interfacing to the baseband and radio modules, switching I/Q and Ethernet control data, and internal component management. The top-level FPGA architecture for such a switch, along with resource utilization for various configurations ranging from eight to sixty-four baseband/radio ports, are given in [Figure 3](#) and [Table 4](#) respectively. These numbers assume the interface is based on the CPRI protocol.

Within the CRAN application, the I/Q cross-connect modules would normally be dimensioned to handle high numbers of baseband/radio ports; typical numbers could be around sixty-four or perhaps higher. In the local baseband switch application, the number of ports are in general significantly lower. Typically, the number of ports could be in the range of eight to thirty-two. Device mapping for such configurations are given in [Table 4](#).

Revolutionary CRAN Base-Station Connectivity

The premise supporting the revolutionary CRAN base station is the desire to create a common open-hardware platform based on low-cost commodity technology. The use of commodity technology is key to both the centralized processing platform and the interface solution used to interconnect the system modules.

In terms of platform technology, the vision is to use GPPs in conjunction with flexible FPGA-based hardware accelerators to perform the baseband, higher-layer, and backhaul functions of the base station. The hardware accelerators are required to provide flexible interfacing and to perform wireless specific processing which is not well addressed by the GPP. Such acceleration is required because, unlike ASSPs and SoCs often used in macrocellular applications, GPPs do not support embedded wireless accelerators for functions such as FFT, DFT, correlation, matrix computation, FEC encoding/decoding, and higher-layer processing. Consequently, such functions must be performed by the companion accelerator devices. The resulting architecture is similar to that described in the [Hardware Acceleration](#) section, which considers hardware acceleration for cellular modem applications in general. The key attraction of this approach is that the hardware platform (i.e., GPP and interface/accelerator devices) offers a high level of flexibility and, in effect, represents a general-purpose computing platform. The functionality of the system is defined by the application software and can be easily adapted to meet changing system requirements.

Also a key element of the CRAN vision is the use of standard interface technology to replace CPRI/OBSAI (for I/Q radio data sample transfer). The preferred choice is 10 Gb synchronous Ethernet (SyncE). However, SyncE does not natively support many of the functions embedded within CPRI/OBSAI, which are essential to the transfer of I/Q data within a cellular base station. The functions of critical importance are timing recovery, accurate RTD measurement, and latency management. To overcome these issues, the use of IEEE Std 1588 PTP in conjunction with SyncE is proposed. In this arrangement, SyncE provides the high-quality frequency reference essential to the radio units, and IEEE Std 1588 provides the timing reference necessary to compute the RTD with high precision. The principal remaining issue is that of latency management. The system must be designed such that latency is as low as possible, deterministic and time-invariant. This can be problematic within an Ethernet based system due to its packet based architecture. Nevertheless provided that the system is designed with this requirement in mind acceptable performance can be achieved for many cellular applications using a SyncE plus IEEE Std 1588 PTP based solution.

[Figure 11](#) shows one possible top-level CRAN architecture for a “revolutionary” channel-card interface/accelerator FPGA. The architecture assumes that the channel card performs LTE processing and that interconnect is based on SyncE plus IEEE Std 1588 PTP. The application provides six 10GbE RRU connections, four 10GbE connections to other channel cards (for internal requirements), and six PCIe connections to the modem pool (6x2-lane PCIe Gen 2 is assumed). This architecture requires a minimum of twenty-two transceivers. All external SyncE connections support IEEE Std 1588, and the frequency reference for SyncE is provided by a high-quality external source. The pre- and post-processing acceleration functionality assumed comprises DL iFFT and UL FFT, PUCCH receiver, and PRACH correlator, each dimensioned for three-sector 20 MHz 4x4 LTE operation. Additionally, co-processing functionality is provided for complete PUSCH FEC (Turbo, Viterbi, and Reed-Muller decoding) for a throughput of up to 400 Mb/s shared over three sectors, and IPSec/RoHC acceleration dimensioned for three-sector 20 MHz 4x4 LTE operation.

This functionality is used as an example to illustrate the typical types of functions that can be accelerated in FPGA within the CRAN application. It is not exclusive and many other types of

functions can also be envisaged. These particular functions are chosen because FPGA implementation is in general very efficient, and the functions themselves are relatively independent. This means that the functional partition remains simple and interfaces well constrained.

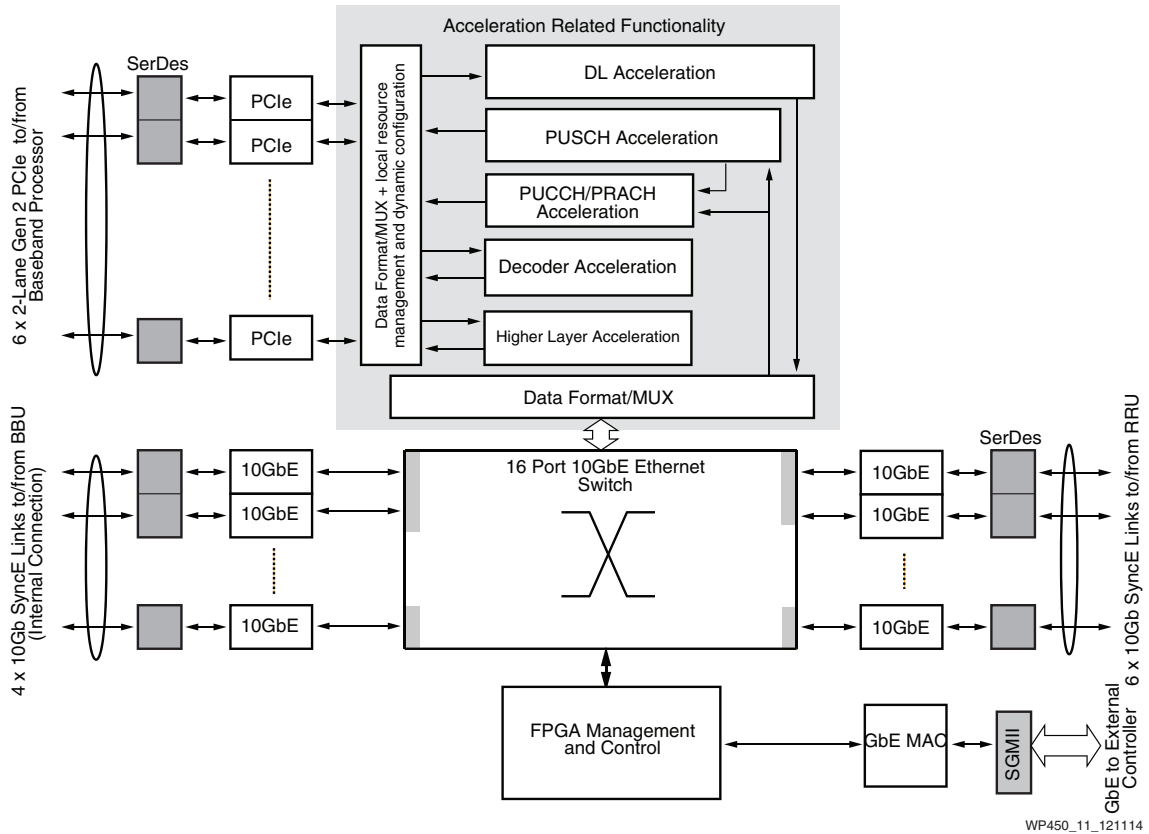


Figure 11: Possible CRAN Channel Card Switch/Accelerator FPGA Top Level Architecture

Resource estimates for this application are provided in [Table 14](#).

Table 14: Resource Utilization 16-Port 10GbE Switch plus H/W Acceleration (CRAN)

Logical Block	FF	LUT	18K Block RAM	DSP	Primary Clock (MHz)	Block Functionality Assumed
10GbE MAC	36,500	37,500	80	0	156.25	10GbE MAC (TX/RX) for ten external ports.
10GbE PCS/PMA	23,000	20,000	30	0	156.25	10GbE PCS/PMA for ten external ports.
Ethernet L2 switch with IEEE Std 1588 support	84,522	96,652	431	0	156.25	Ten external ports and six internal ports. Full L2 switch functionality.
Integrated block for PCIe Endpoint	4,500	3,150	0	0	250	PCIe Gen 2 two lane. Endpoint functionality.
iFFT/FFT acceleration	21,600	21,600	150	54	307.2	3x Radix-4 four-channel iFFT/FFT, complete with data buffers. Dimensioned for 20 MHz 4x4 three-sector operation.
PUCCH/PRACH acceleration	38,034	29,103	102	174	245.76	3x Radix-4 four-channel FFT. Dimensioned for 20 MHz 4x4 three-sector operation.
Decoder acceleration	95,200	95,200	496	374	307.2	Three-sector processing with 100 Format 1 and 50 format 2 users per sector.
Higher-layer acceleration	32,800	32,800	24	0	245.76	IPSec throughput ~1 Gb/s for UL and DL.
FPGA management and CTRL plus GbE MAC	10,250	9,936	42	0	156.25	Implementation in MicroBlaze processor with PLB, DMA, and AXI-Lite.
CTRL interconnect	10,220	19,100	40	0	156.25	AXI-Lite interconnect to all managed components.
Total	356,626	365,041	1,395	602		

This application is relatively complex, because the acceleration functionality is extremely comprehensive and encompasses interfacing, switching, modulation/demodulation, decoding, and higher-layer processing. Moreover, it is dimensioned to provide co-processor functionality for a complex three-sector 4x4 20 MHz LTE system. The application maps to the Virtex-7 FPGA product range, and specifically to the XC7VX690T and XC7VX980T, or alternatively to the UltraScale Kintex XCKU085 device. The resource utilization is summarized in [Table 15](#).

Table 15: Device Utilization: 16-Port 10GbE Switch plus Hardware Accel for CRAN Application

Device Type	FF	LUT	18k Block RAM	DSP	Transceivers
XC7VX690T	41%	84%	47%	17%	28%
XC7VX980T	29%	60%	47%	17%	31%
XCKU085	36%	73%	43%	15%	39%

The most appropriate device depends upon the precise configuration under consideration. The smaller device (XC7VX690T) should be chosen if the application is less complex than the example under consideration. However, the XC7VX980T or XCKU085 are more appropriate for comparable and even higher-complexity applications.

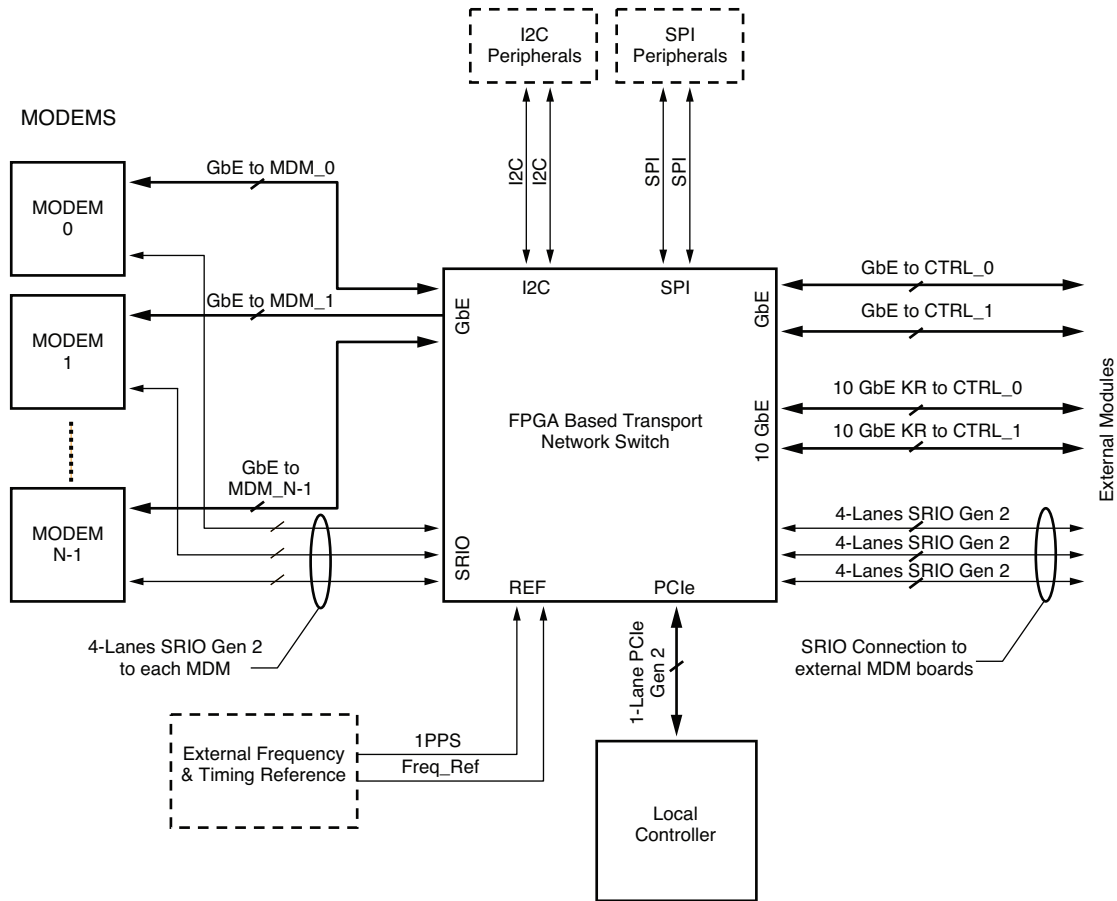
Internal Transport

Internal transport for data other than radio samples is implemented using a variety of protocols chosen to meet the requirements of a wide range of data types. Modern transport architectures are most often based on a combination of Ethernet (both 10GbE and GbE), SRIO, and PCIe. Although switches for such protocols can be found as standard products, FPGA technology is used because it offers the flexibility to adapt operation to suit the fundamental system requirements; it can also combine numerous supplementary functions within a single component. The types of additional functionality can comprise protocol conversion, hardware acceleration functionality, implementation of low-rate signaling protocols used for hardware management (e.g., IIC, SPI, UART, GPIO, etc.), synchronization/timing support, and “glue” logic.

In general, because requirements for internal transport networks are quite diverse, it is often the case that the FPGA that provides this functionality performs several loosely related tasks. The precise functionality performed is base-station architecture dependent, and there is no typical solution. Nevertheless, two broad categories of connectivity architecture can be identified:

- Connectivity associated principally with *intra*-module data transfer
- Connectivity associated with *inter*-module data transfer

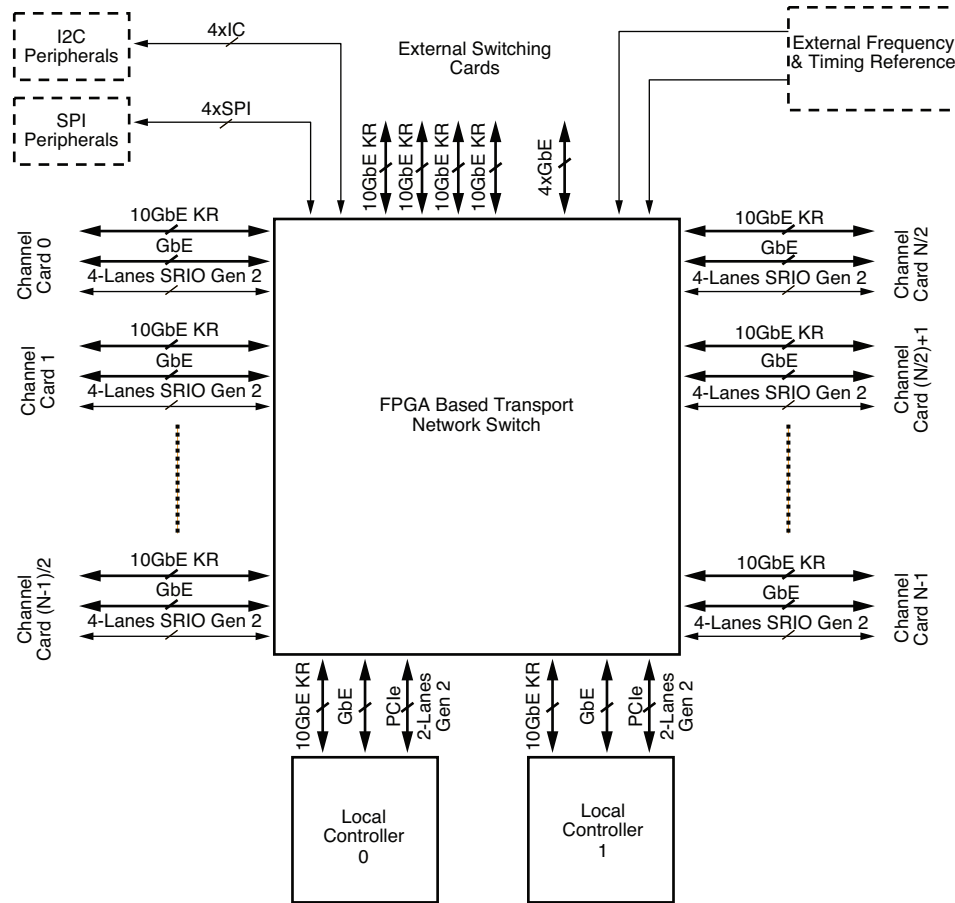
The two architectures are not mutually exclusive, and often both are implemented within a particular base-station design. The principal difference between the architectures is that, in the former, it is the functionality of the module that drives architectural decisions, including choice of transport protocol. As a consequence, connectivity is largely dictated by the specific components on which the module is based. An example of this type of architecture might be a transport network implemented within a channel card. Here the transport network must be adapted to allow data transfer between specific modem and controller devices. A block diagram of a possible module level implementation of such an arrangement is shown in [Figure 12](#).



WP450_12_040814

Figure 12: Channel Card Data Transport Architecture Example

Where connectivity is associated with data transfer between modules, it is often the connectivity requirements themselves that dominate. The transport protocol is chosen to meet the functional requirements of the system as a whole, and the module design is adapted to comply with the system level architecture. An example of this architecture type can be a redundant centralized switching and control module that could be associated with a high-density macrocell or CRAN base station. A block diagram of the implementation of such functionality is shown in Figure 13.



WP450_13_040814

Figure 13: Centralized Switching Module Data Transport Architecture Example

To illustrate the implementation possibilities with Xilinx FPGA technology, application examples of the two types described are considered. The key characteristics of the applications are described in [Table 16](#):

Table 16: Example Internal Transport Configurations Assumed

Application	Internal Channel Card Switch	High-Density Base Station Central Switch
Functionality	Switching user data and CTRL information to local PHYs and external boards	Switching user data and CTRL information to several base station modules
Capacity	Three-sector channel card (three modems): User data rates ~1 Gb/s, CTRL data rates ~250 Mb/s	8x three-sector base station (eight channel cards): User data rates ~24 Gb/s, CTRL data rates ~6 Gb/s
10GbE connectivity	2x 10GbE connections to external switching/ backhaul modules for fast-path user data	14x 10GbE connections to external channel cards, switching/backhaul modules, and internal controller for fast-path user data
GbE connectivity	5x GbE connections to internal modems (x3) and external switching/backhaul modules (x2) for slow-path CTRL data	14x GbE connections to external channel cards, switching/backhaul modules, and internal controller for slow-path CTRL data
SRIO connectivity	6x SRIO four-lane at 5 Gb/s per lane for fast-path connection to local PHYs	8x SRIO four-lane at 5 Gb/s per lane for fast-path connection to external channel cards
PCIe connectivity	1x PCIe Gen 2 single-lane 5 Gb/s end-point connection to board controller	2x PCIe Gen 2 two-lane at 5 Gb/s per lane end-point connection to system controller
Other functionality	2x I ² C, 2x SPI time-base generator	4x I ² C, 4x SPI time-base generator

The top-level block diagram for the channel card internal transport switching architecture is shown in Figure 14:




Figure 14: Block Diagram for Internal Transport Channel Card Switching Architecture

The functionality performed by the FPGA implementation:

- 10GbE support
 - Two 10GbE MACs including PCS/PMR support for KR and AXI-Lite CTRL interface (no support for MAC_SEC and IEEE Std 1588 assumed)
 - Three-port 10GbE L2 switch (dimensioned to support four-port operation)
- GbE support
 - Five GbE MACs supporting frame filter (four entries), 32-bit statistics, and AXI-Lite CTRL interface
 - Six-port GbE L2 switch (dimensioned to support eight-port operation)
- Serial Rapid I/O
 - Six four-lane Gen 2 SRIO end points supporting 5 Gb/s per lane
 - Six-port SRIO switch
- One single-lane Gen 2 PCIe end-point supporting 5 Gb/s
- Other functionality
 - Protocol conversion and CTRL logic to terminate Ethernet connections and route appropriate data to internal and external controllers
 - Embedded soft controller for system control and configuration. Implemented in a Xilinx MicroBlaze embedded processor core
 - 2x I²C plus 2x SPI for local hardware CTRL
 - Local time base derived for external references

The FPGA resource utilization is shown in [Table 17](#):

Table 17: Resource Estimates for the Channel Card Transport Data Switch

Logical Block	FF	LUT	18K Block RAM	DSP	SER-DES	Primary Clock (MHz)	Block Functionality Assumed
10GbE MAC including PCS/PMR (KR)	14,336	14,676	16	0	2	156.25	Two MAC connections to external components. KR support. No MAC_SEC or IEEE Std 1588 support.
Four-port 10GbE switch	16,811	17,076	80	0	0	156.25	Two external ports and one internal port (one port spare).
GbE MAC and SGMII interface	12,750	11,180	10	0	0	157.25	5x 1 Gb/s frame filter (four entries), 32-bit statistics, AXI-Lite CTRL interface.
Six-port GbE switch	30,922	31,252	64	0	0	156.25	Dimensioned for eight-port support. Two ports unused.
SRIO endpoint	43,800	38,700	48	0	24	250	SRIO Gen 2 (spec rev 2.2) 4x at 5.0 Gb/s per lane.
Six-port SRIO switch	30,000	30,000	80	0	0	250	7 External ports 16 Gb/s user data per port.
PCIe Endpoint block	575	400	0	0	1	250	PCIe Gen 2 single lane. Endpoint block functionality.
Protocol conversion and CTRL logic	3,500	3,500	40	0	0	250	Multiplexing data to 10GbE, GbE, and PCIe interfaces and data formatting.
FPGA management and CTRL	7,700	7,700	40	0	0	250	Implementation in MicroBlaze processor with PLB, DMA, and AXI-Lite.
CTRL interconnect	1,200	3,700	0	0	0	250	AXI-Lite interconnect to all managed components.
GPIO, CLK, and random logic	400	400	0	0	0	Various	32 GPIO and 20 SelectIO interfaces for clocks and other CTRL logic.
I ² C interface	600	800	0	0	0	250	Two I ² C interfaces.
SPI interface	10,060	10,060	18	0	0	250	64-bit static configuration. Two interfaces.
Timebase	800	800	2	0	0	250	Compute system time from time reference pulses.
Total	173,454	170,244	398	0	27		

The functionality supported by the centralized switching module FPGA is similar to that described for the channel card transport switch; however, the number of interfaces is significantly higher, and the switching capability (16-port 10GbE switch, 16-port GbE switch, eight-port SRIO switch, and associated support logic) is increased. The overall resource utilization for this application is shown in [Table 18](#) and Xilinx 7 series and UltraScale device mapping for the two applications is shown in [Table 19](#):

Table 18: Resource Estimates for the Centralized Transport Data Switch

Centralized Transport Data Switch	FF	LUT	18K Block RAM	Transceivers	Comments
Total Resources	377,316	377,666	882	50	System clocking is similar to that used in the channel card switch

Table 19: Device Mapping for Example Transport Data Switching Applications

Application	Device	FF	LUT	18K Block RAM	DSP	Transceivers
Channel card transport data switch	XC7K420T	33%	65%	24%	0%	84%
	XCKU060	26%	51%	18%	0%	84%
Centralized transport data switch	XCVX690T	44%	87%	30%	0%	63%
	XCVX980T	31%	62%	29%	0%	69%
	XCKU085	38%	76%	27%	0%	89%

The most appropriate device depends upon the precise configuration under consideration. For example, for the Centralized Transport Data Switch, the XC7VX690T should be chosen if the application is less complex than the example under consideration. However, the XC7VX980T or XCKU085 are more appropriate for comparable and even higher-complexity applications.

Radio Sample and Internal Transport Integration

One natural enhancement to the connectivity architectures presented in previous sections is to combine the radio sample and internal transport switching functionality within a single FPGA implementation. This provides additional flexibility for the base-station transport architecture and allows more optimal use of the switching resource provided by the FPGA. The principal limitation of such an arrangement is the increased number of transceiver devices required to cover all transport protocols. However, Xilinx 7 series and UltraScale architecture technologies offer components that support up to 96/104 high-speed serial transceivers, which are sufficient for most practical implementations currently under consideration.

Switch integration might be of value in the module-based and centralized switching architectures described in previous sections, but it is not normally appropriate for radio applications because the interface is generally provided by CPRI alone.

To illustrate the switch integration possibilities, the examples in the [Internal Transport](#) section are extended to include integrated radio sample implementations. In addition to the transport switching capabilities in the internal channel card switch (detailed in [Table 16](#)), it is assumed that the module must also provide six external CPRI Rate-7 connections to external radios and six internal CPRI Rate-7 connections to the modems (two CPRI connections per modem). The resulting bidirectional 6x6 CPRI switch is therefore capable of switching any radio stream to any baseband stream, and vice versa.

For the high-density central switch, it is assumed that the FPGA provides transport and radio sample switching for up to four baseband modules (reduced from eight in the original example in the [Internal Transport](#) section) with the following interconnect and switching functionality:

- **10x 10GbE KR connections:** One to each of the four baseband modules, four to other switching modules, and two to local processors
- **10-port 10GbE switch**
- **10x GbE connections:** One to each of the four baseband modules, four to other switching modules, and two to local processors
- **10-port GbE switch**

- **4x4-lane SRIO connections:** One to each of the four baseband modules, each connection comprising four high-speed serial connections
- **4-port SRIO switch**
- **2x2-lane PCIe connections** to local control processors
- **16 CPRI Rate-7 connections to multi-carrier radios** (a total of approximately 128 20 MHz LTE carriers, depending upon CPRI configuration)
- **20 CPRI Rate-7 connections to the channel cards** (five CPRI Rate-7 links to each channel card)
- **CPRI switching from any radio stream to any baseband stream** (bidirectional 16x24 CPRI switch)

The overall resource utilization for the two applications and the 7 series and UltraScale device mapping is shown in [Table 20](#):

Table 20: Resource Estimates for the Centralized Transport Data Switch

Application	Resources	Device	FF	LUT	18K Block RAM	DSP	Transceivers
Internal Channel Card Switch	Total Resources	-	232,486	219,418	722	0	40
	Device Utilization	XC7VX415T	45%	85%	41%	0%	83%
	Device Utilization	XCKU085	23%	44%	22%	0%	71%
High-Density Centralized Switch	Total Resources	-	509,708	437,188	2,106	0	64
	Device Utilization	XC7VX980T	42%	71%	70%	0%	89%
	Device Utilization	XCKU085	47%	81%	63%	0%	100%

An alternative integration strategy to increase the CPRI capacity of the design is to remove the SRIO functionality, which could then be performed externally. In the internal channel card switch, this would reduce the transceiver requirement to 16 and the LUT/FF requirement to approximately 150k. In the high-density centralized switch, the transceiver requirement could be reduced to 48 and the logic requirements to ~391k LUTs.

Conclusion

This white paper addresses the application of Xilinx FPGA technology to the implementation of internal communication networks within wireless base stations. This is a critical element of the system design within a range of base-station architectures including conventional macrocell, high-density cell sites, CRAN configurations, and AAA configurations. Within this white paper, two principal aspects are considered for each of these applications:

- I/Q radio sample distribution between the baseband and radio modules
- Base-station internal transport of traffic and control data between functional units

In the case of I/Q radio sample distribution, both conventional networks (which simply provide data transport) and advanced architectures (which provide additional functionality, such as data compression, AAA processing, and hardware acceleration for a range of functions) are examined. Such functionality is considered for several example applications, and FPGA resource utilization for each is presented. For base-station internal transport, both intra- and inter-module communication architectures are considered, and device-functional mappings for specific examples of each are provided. In addition, a combined radio-sample plus internal-transport architecture is presented for a typical macrocell base-station application.

The overall objective of this white paper is to show that Xilinx FPGA silicon technology, in conjunction with connectivity IP available from Xilinx and its partners, is the ideal technology choice for the implementation of wireless base-station connectivity networks. Xilinx connectivity solutions provide high throughput and low latency, with the inherent flexibility necessary to address a wide range of applications. To demonstrate this, numerous use cases are considered, from simple switched networks for I/Q radio data to highly complex multi-functional solutions. Application mapping to current-generation Xilinx all-programmable technology is provided. The analysis demonstrates that 7 series and UltraScale technologies provide the FPGA fabric and flexible I/O resources necessary to address a vast range of connectivity applications. The Kintex UltraScale FPGAs are particularly well suited to base-station connectivity applications, providing up to sixty-four 16 Gb/s gigabit transceivers with over 663k LUTs and 5,520 DSP slices. Moreover, line rates of up to 12.5 Gb/s can be supported in the lowest speed-grade devices, which means that all connectivity protocols commonly used with wireless base stations can be supported by cost-optimized components.

References

1. *Developing and Integrating a High Performance HET-NET* (4G Americas white paper). October 2012.
2. *LTE-Advanced, 3GPP* (3GPP white paper, Jeanette Wannstrom). May 2012.
3. *LTE Release 12: Taking Another Step Toward the Network Society* (Ericsson white paper). January 2013.
4. *LTE Release 12 and Beyond* (Nokia Siemens Network white paper). October 2012.
5. CPRI Specification V6.0, *Common Public Radio Interface (CPRI): Interface Specification*. 30 August 2013.
6. Open Base Station Architecture Initiative: *BTS System Reference Document, Version 2.0*. 27 April 2006.
7. Open Base Station Architecture Initiative: *Reference Point 3 Specification, Version 4.2*. 18 March 2010.
8. ETSI Group Specification, Open Radio equipment Interface (ORI): *Requirements for Open Radio equipment Interface (ORI) (Release 1)*, ETSI GS ORI 001, V1.2.1. August 2012.
9. ETSI Group Specification, Open Radio equipment Interface (ORI): *ORI Specification, Part 1: Low Layers (Release 1)*, ETSI GS ORI 002-1, V1.1.1. October 2011.
10. IEEE Standard for Information Technology – Telecommunications and Information Exchange between Systems – Local and Metropolitan Area Networks, Specific Requirements Part 3: *Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications*, IEEE Std 803.2-2008. 26 December 2008.
11. ITU-T Series G: *Transmission Systems and Media, Digital Systems and Networks, Timing Characteristics of Synchronous Ethernet Equipment Slice Clock (EEC)*, G.8262/Y.1362. August 2007.
12. IEEE Standards, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, IEEE Std 1588-2002. 8 November 2001.
13. [Backgrounder](#), *Moving a Generation Ahead with All Programmable FPGAs, SoCs, and 3D ICs* (Xilinx Backgrounder)
14. [DS180](#), *7 Series FPGAs Overview* (Xilinx data sheet)
15. [DS190](#), *Zynq-7000 All Programmable SoC Overview* (Xilinx data sheet)
16. [WP380](#), *Xilinx Stacked Silicon Interconnect Technology Delivers Breakthrough FPGA Capacity, Bandwidth, and Power Efficiency* (Xilinx white paper)
17. [Backgrounder](#), *Introducing Xilinx UltraScale Architecture: Industry's First ASIC-Class All Programmable Architecture* (Xilinx Backgrounder)
18. [DS890](#), *UltraScale Architecture and Product Overview* (Xilinx data sheet)
19. [WP435](#), *Xilinx UltraScale: The Next Generation Architecture for Your Next Generation Architecture* (Xilinx white paper)
20. [PB012](#), *LogiCORE IP CPRI v7.0* (Xilinx product brief)
21. IEEE Transactions on Wireless Communications: *Compressed Transport of Baseband Signals in Radio Access Networks*, Dragan Samardzija, John Pastalan, Michael MacDonald, Susan Walker, Reinaldo Valenzuela. 4 May 2012.
22. IEEE Globecom 2012, Wireless Communications Symposium: *Baseband Signal Compression in Wireless Base Stations*, Aida Vosoughi, Michael Wu, Joseph R. Cavallaro.
23. IEEE ChinaCom 2012, International ICST Conference on Communications and Networking in China: *CPRI Compression Transport for LTE and LTE-A Signal in C-RAN*, Bin Guo, Wei Cao, An Tao, Dragan Samardzija.
24. ETSI Group Specification, Open Radio equipment Interface (ORI): *Evaluation of Optimum Parameters and Computational Time of the IQ Data Compression Scheme*; ORI(13)VM_004, V. 7 January 2013.
25. WP444, *Synchronization and Timing Support for Heterogeneous Mobile Wireless Networks Using Xilinx All Programmable FPGAs* (Xilinx white paper)

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
03/22/2018	1.2	Updated Internal Control and Interconnect Functionality ; Revolutionary CRAN Base-Station Connectivity ; and Table 20 .
09/29/2015	1.1	Updated Table 4 , Table 15 , Table 19 , and Table 20 .
12/11/2014	1.0	Initial Xilinx release.

Disclaimer

The information disclosed to you hereunder (the "Materials") is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available "AS IS" and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of Xilinx's limited warranty, please refer to Xilinx's Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in such critical applications, please refer to Xilinx's Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>.

Automotive Applications Disclaimer

AUTOMOTIVE PRODUCTS (IDENTIFIED AS "XA" IN THE PART NUMBER) ARE NOT WARRANTED FOR USE IN THE DEPLOYMENT OF AIRBAGS OR FOR USE IN APPLICATIONS THAT AFFECT CONTROL OF A VEHICLE ("SAFETY APPLICATION") UNLESS THERE IS A SAFETY CONCEPT OR REDUNDANCY FEATURE CONSISTENT WITH THE ISO 26262 AUTOMOTIVE SAFETY STANDARD ("SAFETY DESIGN"). CUSTOMER SHALL, PRIOR TO USING OR DISTRIBUTING ANY SYSTEMS THAT INCORPORATE PRODUCTS, THOROUGHLY TEST SUCH SYSTEMS FOR SAFETY PURPOSES. USE OF PRODUCTS IN A SAFETY APPLICATION WITHOUT A SAFETY DESIGN IS FULLY AT THE RISK OF CUSTOMER, SUBJECT ONLY TO APPLICABLE LAWS AND REGULATIONS GOVERNING LIMITATIONS ON PRODUCT LIABILITY.