Enabling Distributed Base Station Architectures with CPRI

White Paper

by Gerry Leavey

Issue No. 1: February, 2006

Legal Information

Copyright

PMC-2051877 (issue1)

Disclaimer

None of the information contained in this document constitutes an express or implied warranty by PMC-Sierra, Inc. as to the sufficiency, fitness or suitability for a particular purpose of any such information or the fitness, or suitability for a particular purpose, merchantability, performance, compatibility with other parts or systems, of any of the products of PMC-Sierra, Inc., or any portion thereof, referred to in this document. PMC-Sierra, Inc. expressly disclaims all representations and warranties of any kind regarding the contents or use of the information, including, but not limited to, express and implied warranties of accuracy, completeness, merchantability, fitness for a particular use, or non-infringement.

In no event will PMC-Sierra, Inc. be liable for any direct, indirect, special, incidental or consequential damages, including, but not limited to, lost profits, lost business or lost data resulting from any use of or reliance upon the information, whether or not PMC-Sierra, Inc. has been advised of the possibility of such damage.

Trademarks

For a complete list of PMC-Sierra's trademarks and registered trademarks, visit:
http://www.pmc-sierra.com/legal/

Abstract

Wireless operators are increasingly challenged to support a diverse array of emerging broadband, data-oriented mobile services. Simultaneously, they must find ways to reduce capital and operational costs of the networks that will deliver both new and existing services. Two major equipment trends are beginning to emerge as enablers for deploying more cost-effective, next-generation wireless access networks:

- Distributed base stations using Remote Radio Heads (RRHs); and
- Vendor-driven open standards for next-generation base stations.

As a leading contributor to both of these trends, the Common Public Radio Interface (CPRI) industry initiative has developed an open, publicly available specification that standardizes the digital interface to RRHs and enables the realization of this exciting new base station deployment topology. This white paper will:

- Describe the distributed base station concept, RRHs and the many benefits that this new network architecture can bring for wireless operators
- Introduce CPRI and illustrate how it fuels the growing paradigm shift to distributed base stations
- Explore the features of the CPRI protocol in detail

About PMC

PMC-Sierra™ is a leading provider of high speed broadband communications and storage semiconductors and MIPS-Powered™ processors for Enterprise, Access, Metro Optical Transport, Storage Area Networking and Wireless network equipment. The company offers worldwide technical and sales support, including a network of offices throughout North America, Europe and Asia. The company is publicly traded on the NASDAQ Stock Market under the PMCS symbol and is included in the S&P 500 Index.

About the Author

Gerry Leavey is a Technical Advisor in the Communication Products Division at PMC-Sierra. Previously, he has held design and management positions with British Telecom, Lucent Technologies, and Toucan Technology. Gerry graduated with a bachelor of Electronic Engineering degree from the National University of Ireland, Galway in 1988.

Revision History

Issue No.	Issue Date	Details of Change
1	February 2006	Document created

Table of Contents

List of Figures

List of Tables

1 Introduction

Wireless operators are increasingly challenged to offer a diverse array of emerging broadband, data-oriented mobile services. To support this new era of data-driven growth, service providers must deploy base stations with more sophisticated enabling technologies, more flexible support of multiple radio standards, and more capacity to meet ever-increasing traffic demands. However, the deployment of traditional base stations is challenging how operators can maintain and expand their service business models in the face of ever increasing capital and operational expenditure costs.

In response to these issues, two major architectural trends are beginning to emerge as enablers for reduced cost next-generation wireless access networks. The first of these trends is the concept of distributed base stations and Remote Radio Heads (RRH). The use of centrally located base station “hotels” that connect via lossless fiber links to remote radio units dramatically reduces costs associated with site acquisition, site leasing, and energy consumption for carriers. In addition, allowing several remote antenna sites to be controlled by one central base station offers more rapid and more scaleable network deployment.

The second of these trends is the emergence of vendor-driven open standards for next generation base stations. The benefits of open standardization are already well accepted in other industries: lower individual R&D costs as the total costs are shared over a wider community; shorter time-to-market; increased competition among component suppliers leading to greater innovation and downward pressure on pricing; and acceleration of the adoption and evolution of new best-in-class technologies. Base station vendors are now just beginning to embrace open standards so that they may leverage the same powerful market forces into their industry.

Harnessing these two powerful forces will enable base station vendors to revolutionize the way they develop their systems. In response to the aforementioned trends, the Common Public Radio Initiative (CPRI) industry association has developed an open, publicly available specification that standardizes the digital interface to RRHs.

This white paper first describes the distributed base station concept and the myriad benefits it can bring for wireless operators. Secondly, it provides an introduction to CPRI, the key interface technology that enables this exciting new base station deployment topology to be realized. This is followed by a detailed technical overview of the CPRI specification.

2 Distributed Base Stations

Base station OEMs are continuously striving to find new and innovative product architectures that not only provide greater flexibility to their operator customers but also help them to reduce the capital and operational expenditure incurred with each new equipment or service deployment. The concept of distributed base stations is one such recent innovation that is garnering significant interest.

2.1 Traditional Base Station Model

In most current deployments, base stations are located close to the antenna, for example in an environmental enclosure or “hut” at the base of an antenna tower or in the basement of a tall building. In this traditional base station configuration, the Radio Frequency (RF) transceiver and Power Amplifier (PA) are co-located with the other functional elements of the base station (such as baseband processing and backhaul transport) in a single enclosure. From this enclosure, each PA drives their associated antenna over electrical cables, where the antenna is situated at the top of the mast or building. Figure 1 below illustrates such a typical base station deployment using a co-located radio head.

Figure 1: **Monolithic Base Station with Co-Located Radio Transceivers**

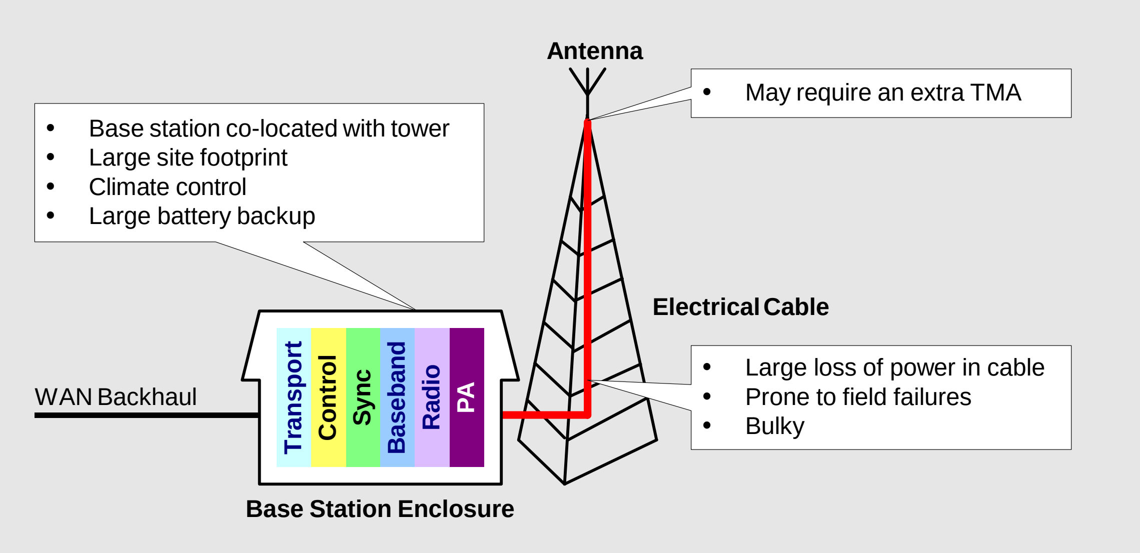

Deploying base stations in this traditional manner has a number of drawbacks for operators:

- Because the base station drives the antenna(s) over lossy electrical cable, it must be physically located very close to the antenna(s). This can create problems with acquiring suitable sites for new cell build-outs, for example because a larger site footprint is required to install and secure the base station hut than is geographically available or because the base station enclosure must be located on a structurally reinforced rooftop.

- If a suitable site can be acquired, the usually unattractive size and appearance of a base station hut or enclosure can be a cause of environmental complaints from local authorities or the general public. This factor can introduce significant delays in the process of applying for and receiving planning permission.
- Many sites are not purchased outright by the wireless carrier but are leased from the property owner, where the volume of real estate space consumed by the base station enclosure is a large factor in determining the cost of the lease.
- Installation may be challenging because of the size and location of the base station enclosure. For example, locating an enclosure on a rooftop may require heavy lifting gear to place it as a pre-assembled entity or may require the entire base station to be assembled from scratch from its piece parts and commissioned on-site, thereby greatly increasing the time and cost of installation.
- Once installed, another factor used in determining the lease cost for a tower site is the number of cables and their weight. Because each radio transceiver co-located in the base station enclosure will drive its associated antenna(s) directly via individual, bulky electrical cables, cabling requirements can have a large impact on overall lease costs.
- The use of lossy electrical cables to drive the antenna(s) has a dramatic bearing on the energy consumption of the base station. The loss in these cables and their associated connectors can range from a typical value of 3db to as much as 10dB in extreme cases, meaning that between 50% and 90% of the radio transceiver’s output power is dissipated in cable transmission before it even reaches the antenna.
- The same cable losses may require the deployment of additional tower mounted amplifiers (TMA) to boost the uplink signals received at each antenna before they can successfully drive feeder cables back to the radio receivers in the base station enclosure.
- All this extra power required to drive the electrical feeder cables means that higher output power amplifiers must be deployed. These high power amplifiers are more expensive and have traditionally had poor operating efficiencies of around 10%, further compounding the problem of high energy consumption by base station.
- The wasted power generated by these low-efficiency amplifiers is dissipated as thermal energy and requires the base station enclosure to have sophisticated metal enclosures with climate control facilities such as air conditioning.
- When deploying traditional base stations as part of a major network upgrade (e.g. when deploying 3G as an overlay to an existing 2G network), it often requires the installation of an additional new base station at an existing cell site. However, it may be very difficult (if not impossible) to co-locate two independent base stations at the same site because of space constraints. Even if it is possible to physically co-locate them, the extra space and energy consumed by the second base station leads to a large increase in the operational cost of that site.
- Base stations require a backup power source should their primary electricity feed from the mains be interrupted. Because of the high energy consumption of a typical base station, this backup power source is implemented using either a large array of lead-acid batteries or an oil-fired generator. Both of these options further increase the expensive footprint of the cell site and pose additional environmental and operational challenges.

Clearly, the deployment of large numbers of traditional monolithic base stations is an expensive investment in time and resources for wireless service providers as they deploy their networks and easily forms the largest portion of their day-to-day operational expenditure as they deliver their services. As a result, it is obvious that these operators would be attracted to a new deployment paradigm that can significantly reduce these capital and operational costs. Given these very compelling incentives, the base station vendor community has developed the new concept of distributed base stations to address the problem.

2.2 Distributed Base Station Model

In contrast to the local base station model, the distributed base station architecture splits the RF transceivers from the rest of the base station and relocates them next to their associated antennas such that the antennas are driven directly with minimal transmission power loss. The digital baseband data is transported between the baseband processing located in the central base station enclosure and the remotely located RF transceivers over a flexible, loss-free optical fiber, as shown in Figure 2.

Figure 2: Central Base Station Hotel with Remote Radio Heads

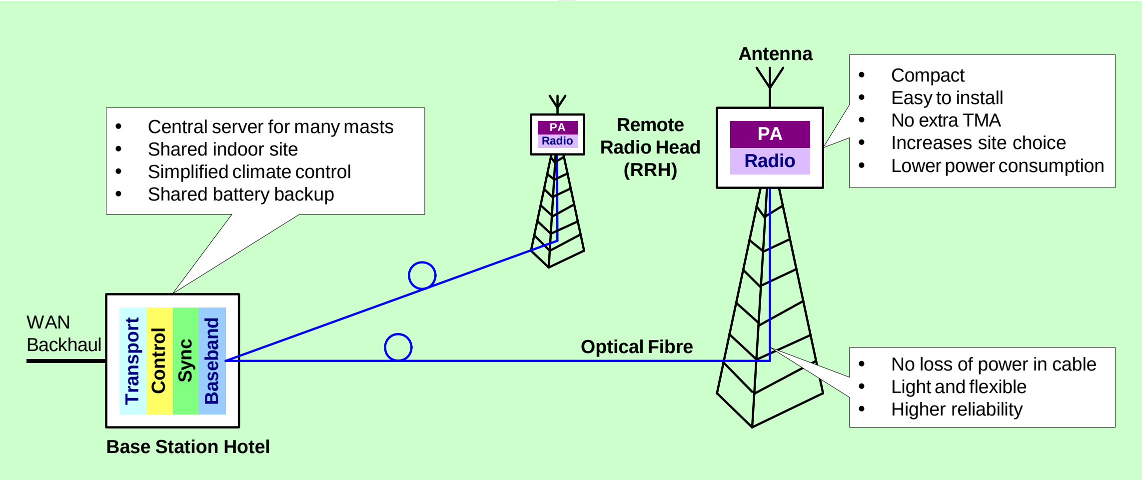

Additionally, the use of fiber interconnect provides the opportunity to remotely locate the RF transceiver a significant distance from the main base station cabinet such that a single, central base station "hotel" can centrally serve a large number of these remote RF transceivers (referred to as Remote Radio Heads (RRH)). This base station hotel concept applies equally well to both outdoor and in-building deployments.

The distributed base station architecture offers a compelling alternative to the traditional base station deployment model. It addresses all the drawbacks identified in section 2.1, as follows:

- With the distributed base station model, there is no longer a requirement to have a large enclosure located close to the antenna(s). For example, the base station enclosure can now be located in an existing building close to but not necessarily adjacent to the cell tower. This dramatically reduces the requirements for cell site footprint, making many new locations viable as alternatives and greatly reducing the cost of site acquisition.

- The small physical size of an RRH means that, with their associated antennas, their presence has a much-reduced environmental impact. This can expedite the often protracted process of attaining planning permissions with local authorities and consultations with community groups.

- As the physical space requirement at each antenna location is greatly decreased, so are the proportional costs of leasing the site.

- RRHs are compact equipment units that can be easily handled and installed by individual field engineers without a need for large teams or for special heavy lifting equipment. As a result, the overall time and cost of site installation can be reduced significantly and the pace of network deployment can be accelerated.

- By removing the need for bulky electrical cables and substituting this interconnect with light, flexible optical fibers, another component of the site lease costs can be reduced.

- Undoubtedly, the single biggest issue with the monolithic base station deployment model is the very large loss experienced in the electrical feeder cables and the increased energy consumption that this incurs. In the distributed base station scenario, by feeding the RRH with digital baseband data over optical fiber, transmission to the antenna location is now virtually loss-free with the exception of some small attenuation in the short electrical cable connections between the RRH and the antennas.

- By co-locating the RRH receiver with the receive antenna, it is no longer necessary to have an additional TMA act as a booster amplifier for the uplink signal path.

- Removing the loss in the signal path due to transmission over the feeder cables means that the output power of the amplifiers can be reduced. As their efficiency is a percentage of their total output power, reducing the maximum output requirements for these PAs results in further energy savings for the system.

- With this significantly reduced power, there is a large reduction in the amount of thermal energy dissipated by the system. This means that the RRH can be designed to passively manage the thermal issues without the need for any expensive climate control facilities at the remote site. Additionally, the central base station hotel that supports the RRHs no longer needs to be deployed in the field and can be implemented using standard telecommunications equipment practices to operate in the more benign environmental conditions of an indoor facility.

- Once a base station hotel has been installed, it can become a central shared unit that serves more than one remote cell site. This breaks down the previous economic barrier of requiring one base station per cell site and introduces new economies of scale, for example where the transport backhaul costs from a number of cell sites to their controller can now be aggregated and shared. Additionally, it provides greater scalability and deployment flexibility for introducing new cell sites, where a new location can be very quickly brought online by simply deploying a RRH and performing all commissioning of the cell from the same shared central location.

- The distributed architecture also facilitates the rapid deployment of new base station equipment to existing cell sites. The radio components usually make up the largest portion of the physical volume of a traditional base station, as well as determining the capacity of climate control equipment required at the site. If distributed base stations are used to deploy new network capability at existing sites, all the high power radio components are now located remotely while the remainder of the base station consumes a much smaller volume and can often be retro-fitted into an existing enclosure. This enables quick installation and, more importantly, allows the deployment of the new technology without consuming additional space and energy.

- The dramatic reduction in remote site power consumption also reduces the amount of backup power capacity required, offering the opportunity to use new, compact, lower capacity backup solutions and further shrinking the footprint of the installation.
- Another benefit of the distributed architecture is that it provides a very cost effective approach for deploying picocells in an indoor environment to improve network coverage and capacity e.g. an office block or a shopping mall. Small, low power radio heads are unobtrusive and can be deployed very easily inside a building. Once installed, they are connected back to a central base station hotel using optical fibers drawn through existing in-building wiring conduits.

As illustrated above, the RRH deployment scenario promises many significant cost savings for operators by reducing site acquisition and management expenses, centralizing system complexity for easier maintenance and service enhancements, and improving network reliability through reduced complexity of equipment installations in the field. Because of the growing interest in this new deployment scenario, CPRI has addressed RRH networking in its specification.

3 CPRI

3.1 Introduction

Founded in April 2003, the Common Public Radio Interface (CPRI) is an industry co-operation between five large base station OEMs: Ericsson AB, Huawei Technologies Co. Ltd., NEC Corporation, Nortel Networks and Siemens AG. Although membership is strictly limited to these five OEMs, the CPRI specification is publicly available and can be downloaded directly from the CPRI web site. CPRI's goal is the open standardization of the internal digital interface between the baseband processing and the radio transceiver in the base station and was the first industry specification to address the specific interface needs of RRH networking.

The first CPRI specification, Version 1.0, was released in September 2003. Since then, there have been four subsequent releases that have made incremental recommendations, clarifications and corrections to the initial draft. The current release, Version 2.0, has been published and stable since October 2004 and supports the following fundamental RRH requirements:

- Transport of digital baseband radio data over either optical or electrical media
- Deployment in point-to-point, tree-and-branch, chain, and ring network topologies
- In-band transport of a high-stability network frequency and time references
- In-band transport of RRH control messaging
- In-band transport of vendor-specific application data
- Accurate calibration of the link delay

The original scope of CPRI was focused upon supporting the UMTS air interface standard. Therefore, many of the specific characteristics of CPRI such as link bit rate and frame format were chosen to ensure that it is a very efficient transport solution for UMTS user plane data. However, CPRI's basic transport mechanisms are specified in a flexible way such that other air standards may be supported.

3.2 CPRI System Architecture

The CPRI specification divides the traditional radio base station functionality into two distinct subsystems: the Radio Equipment Controller (REC) and the Radio Equipment (RE). The REC subsystem is responsible for processing radio information in the digital baseband domain while the RE subsystem performs the conversion between digital baseband signals and analog radio frequency signals that interface with the antenna.

The CPRI interface defines a digital point-to-point interface between the REC and the RE, as illustrated in Figure 3. By using a digital interface, systems designers have the flexibility to either co-locate the REC and RE within a single enclosure or to remotely locate the RE from the REC in a distributed topology. Additionally, using long-range optical transmission technologies, REs may be located in very remote locations over 10 km from a central REC.

Figure 3: CPRI Architecture

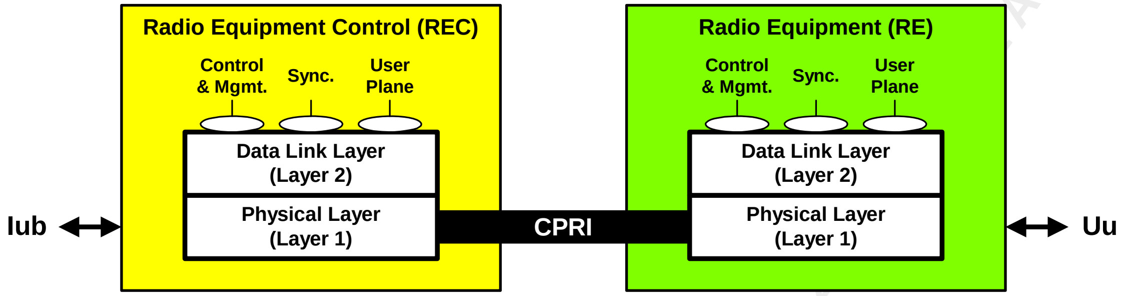

3.3 CPRI Information Flows

Remotely locating REs means control and synchronization information must be communicated to and from the REC over the same digital interface as the digital baseband traffic. To specify the operation of the complete interface, CPRI defines both the physical and data link processing layers of the protocol and the characteristics of the information that is transported in the different user plane payload, control and synchronization sub-channels supported by the protocol. The CPRI protocol stack is illustrated in Figure 4 below:

Figure 4: CPRI Protocol Stack

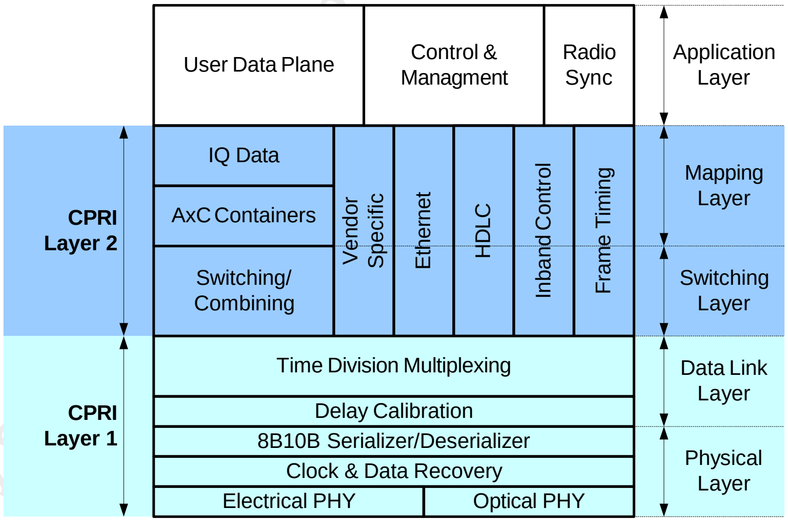

3.3.1 User Data Plane Information

User Data Plane data is transported as baseband In-Phase/Quadrature-phase (IQ) samples. The CPRI transport payload is subdivided using simple time division multiplexing (TDM) to simultaneously support multiple independent streams of IQ samples, where each bidirectional stream transported is the digital baseband data associated with the transmission and reception of one wireless carrier at one antenna. These IQ sample streams are referred to in CPRI as antenna carriers (AxC).

3.3.2 Control and Management Information

CPRI control data is divided into two different types of information; a control flow to maintain the link itself and a control flow to remotely manage the operation of the RE radio functions. To address the former requirement, CPRI has a dedicated sub-channel running a simple bit-oriented control protocol, the L1 Inband Protocol, to handle basic link configuration and alarms that are independent of the RE application. The RE application control requirements are supported through two dedicated sub-channels, a lower bandwidth High level Data Link Control (HDLC) and a higher-bandwidth Ethernet alternative. These two standard data-link options support the communication of vendor-specific Operations, Administration, Management and Provisioning (OAM&P) messaging between centralized control in the REC and the RE.

3.3.3 Synchronization Information

The synchronization information carried by the CPRI link consists of two components: a highly accurate network frequency reference and a very precise air interface frame timing reference. Through these two different components, RRH radio transmission and reception can be precisely synchronized and coordinated in both frequency and time with the rest of the network.

3.4 CPRI Network Topologies

The CPRI architecture splits and distributes the base station functionality between the REC and the RE and, in its most basic form, a CPRI link provides a single point-to-point connection between a REC and one RE. However, each REC will typically be required to support more than one RE. Therefore, the simplest deployment topology for CPRI-connected REs involves networking them via multiple point-to-point links from a centralized REC in a star topology, as illustrated in Figure 5 below.

Figure 5: CPRI Star Topology

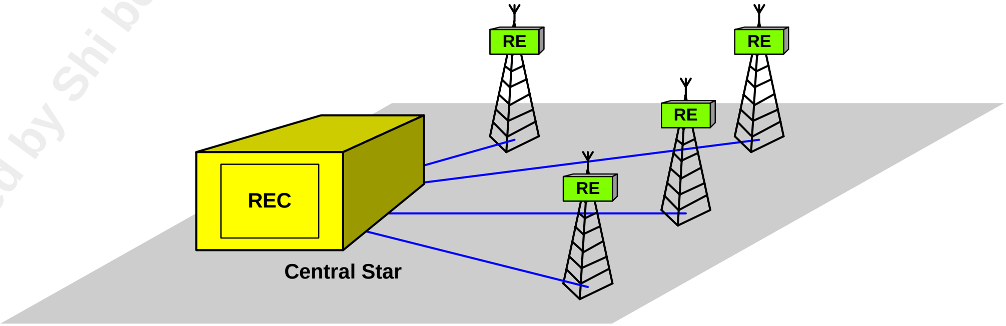

However, there are practical network deployment situations where a central star topology may be too difficult or too costly to implement. To address this issue, the CPRI group added additional networking functionality to Version 2.0 of the specification. These networking features support the concept of RE-to-RE CPRI links, as illustrated in Figure 6.

Figure 6: CPRI Networked System Architecture

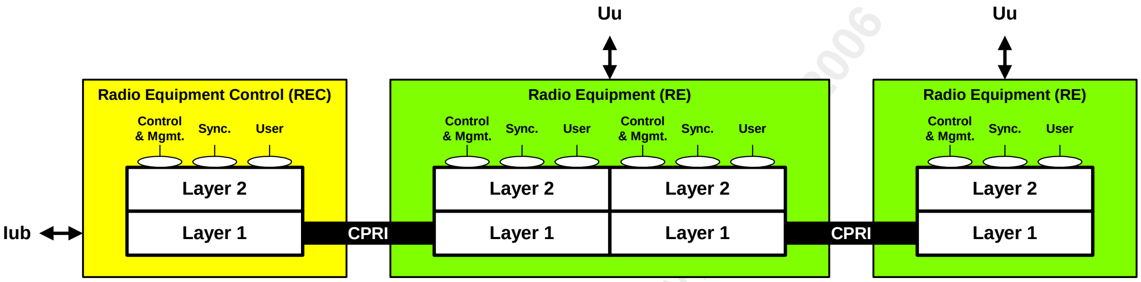

By allowing RE-to-RE connectivity, RRHs may be deployed in many different networked configurations that can address the problems of deploying base station infrastructure in real network environments. The three fundamental topologies that are enabled are chains, rings and tree-and-branch, as illustrated in Figure 7.

Figure 7: CPRI Network Topologies

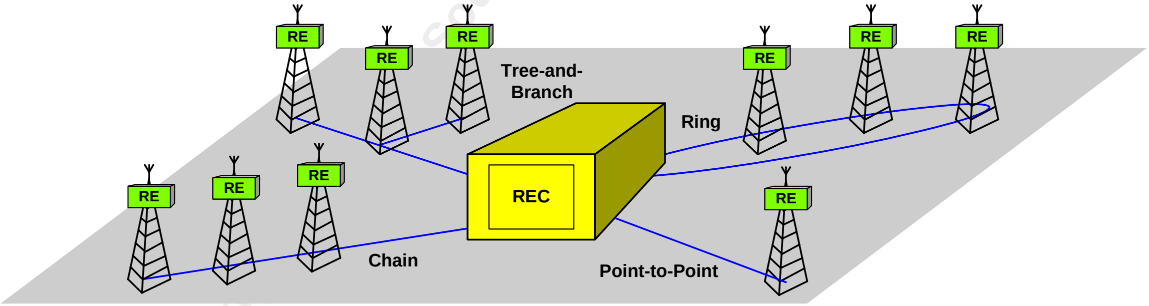

Each networked topology has benefits and drawbacks, as summarized below:

- The simplest networked topology is a chain, where a number of REs are cascaded and share a single connection to the REC. The advantages of a chain are that it maximizes the use of REC ports, where a single REC connection can be shared by a number of REs, and that the chain minimizes the amount of fiber deployed such that REs are only connected to their nearest neighbor and do not need to have an independent fiber connecting each RE back to a centralized REC. However, simple chain networks are not very resilient because a link failure at a single RE will result in link failures for all REs that are cascaded beyond it.
- A tree-and-branch network shares the advantages of chain networks by also maximizing the use of each REC port and minimizing fiber requirements. In this network, a CPRI link from the REC is terminated at a single remote location before being split out to a

number of REs over individual point-to-point links. The tree-and-branch network addresses the resilience problem of the chain because no RE can be a single point of failure for the network. However, the hub point for the individual branches is itself a single point of failure.

- The ring network’s main advantage over a chain is that it addresses the issue of network availability by “closing” the chain and providing an alternative path to maintain connectivity between the REC and all REs in the presence of a link failure at any one segment in the ring. However, a ring requires two dedicated ports at the REC per network and an additional independent fiber network to provide the redundant protection path.

3.5 CPRI Physical Layer

In defining the physical layer of the CPRI specification, the members were mindful of the need to standardize on common solutions while retaining flexibility for vendors to support different interconnect media and topologies for different system configurations. For example, a primary focus for CPRI remains the connectivity between a central base station hotel and remotely located radio heads in a distributed system using optical transmission. However, it could equally be used to interconnect baseband and radio transceiver cards that are co-located within a single base station enclosure over an electrical backplane or cabling. Therefore, CPRI has defined common physical layer requirements that are independent of the media supported, and has provided technical recommendations for other physical layer features that are media-specific.

3.5.1 Media-Independent Physical Layer Requirements

In terms of media-independent features, CPRI specifies requirements in the following areas:

- The transmission line coding that must be used
- The link rates that are allowed
- The quality of the reference clock that can be recovered from the line

Regardless of the selected interconnect medium, the data is 8B/10B coded before transmission. This line code is appropriate for supporting both electrical and optical interconnect and, as will be described in Section 3.6.1 in more detail, it can be directly reused by the higher-level framing protocol to provide periodic time alignment information with no additional overhead.

Once encoded, three line rates are specified to provide three scaleable levels of transport capacity over the link: 614.4 Mbit/s (the base rate), 1228.8 Mbit/s (2x the base rate) and 2457.6 Mbit/s (4x the base rate). These link rates have been chosen as integer multiples of the 3.84 MHz UMTS chip rate to simplify the recovery of a high quality network reference clock.

Finally, the recovery of a very accurate network timing reference from the link is an extremely important part of the CPRI specification. The fundamental system requirement that CPRI must address is that the UMTS standard stipulates a frequency accuracy of better than ±0.05 ppm at the radio interface [3]. In the traditional local base station model, the network element operates as a fully synchronized system where all cards are timed from a shared high quality source that is traceable to a primary network reference. This high quality source distributes its timing signals directly and very simply over the local interconnect in the base station, thereby guaranteeing the availability of a highly stable, low jitter reference to each radio transceiver card. However, in a distributed base station, the only path for transporting this timing to the RRH is over the CPRI link. Therefore, a CPRI implementation at the REC must be able to generate a low jitter transmit output that is frequency locked to the central timing reference. Conversely, at the remote RE, the CPRI implementation must recover timing from the 8B/10B-encoded line rate and regenerate the original reference frequency to the required UMTS accuracy, typically using a clean-up Phase Locked Loop (PLL) to attenuate the jitter. To help bound this implementation challenge, CPRI specifies the maximum jitter transfer bandwidth of the PLL (Requirement R-17), the maximum contribution that transporting the timing over the end-to-end link can make to the total frequency deviation of the recovered reference (Requirement R-18), and the maximum settling time of the timing recovery circuit (Requirement R-30).

3.5.2 Media-Specific Physical Layer Recommendations

Beyond the requirements that all CPRI interfaces must share at a physical level, there are features that may vary from solution to solution depending upon whether electrical or optical transmission media is used. The CPRI specification provides recommendations for these features based on direct reuse of existing, well adopted, industry-standard physical interfaces.

In terms of electrical interface technology, the original CPRI specification [2] provided full scope for vendors to implement CPRI links using any available physical standards as long as they supported the media-independent requirements detailed previously in Section 3.5.1 and a Bit Error Rate (BER) of better than 10^{-12} for both User Data Plane and Control & Management data sub-channels. However, to reduce the number of implementation options and to improve vendor interoperability, CPRI revised the specification from Version 1.2 onwards to recommend two preferred alternatives: XAUI, as a single interface that can support all line rates; and 1000Base-CX, a higher voltage alternative that can support the two lowest line rates. CPRI directly reuses these electrical standards with few modifications other than scaling the baud rate to support the CPRI-specific transmission speeds.

For optical interfaces, re-use of readily available Fibre Channel or Gigabit Ethernet transceivers based upon the Small Form-factor Pluggable (SFP) modular form factor is recommended. Additionally, it is recommended that the optical transceiver solutions should support the LV and/or HV electrical interfaces. Using the SFP form factor, this approach allows a single CPRI interface implementation to flexibly support either active optical or passive electrical SFP modules and cabling.

The connection between two modules is not limited to a single CPRI link. Many links can interconnect two modules in order to increase the transport bandwidth but only one link is defined as an active link while the others are defined as passive links. Only the active link carries Control and Management information as well as synchronization information.

3.6 CPRI Data Link Layer

3.6.1 Framing and Sub-channel Multiplexing

Above the physical layer, the CPRI data link layer provides a channelized transport mechanism for the different RRH information flows in a hierarchical frame structure, illustrated in Figure 8:

Figure 8: CPRI Frame Structure

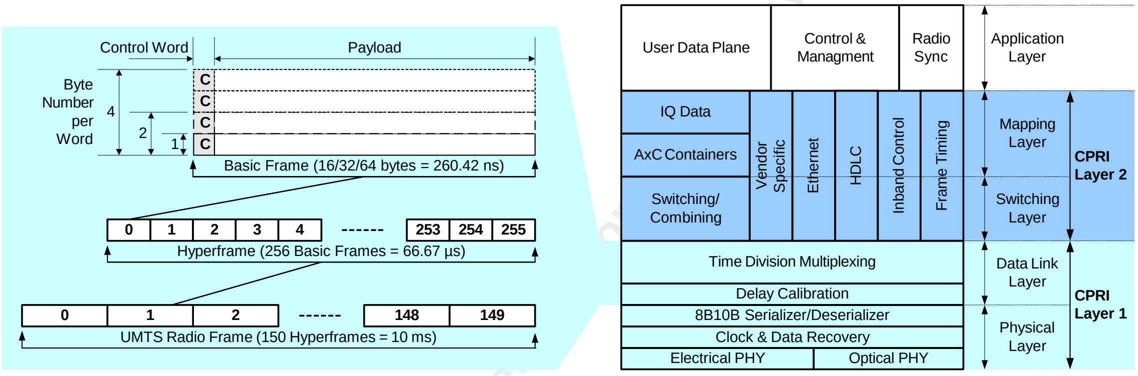

Frame Structure **Protocol Stack**

The fundamental unit of the CPRI frame structure is a 16-word sub-frame called a Basic Frame (BF). Within each 16-word BF, 15 words are reserved for transporting user data payload and 1 word is reserved as a control word. A fundamental cornerstone of the CPRI frame structure is that the frequency of BF transmission is 3.84 million per second – that is, the BF is periodic and repeats at the UMTS chip rate. In this way, AxCs sampled at the UMTS chip rate can be directly mapped to CPRI basic frames without requiring significant buffering for rate adaptation. This provides a very predictable, consistent and low latency mechanism for transporting UMTS AxCs over the link.

The CPRI link rate can be scaled from the lowest rate of 614.4 Mbit/s to 1228.8 Mbit/s and 2457.6 Mbit/s, providing three increasing levels of transport capacity. However, to maintain the fixed periodic relationship between the BF repetition rate and the UMTS chip rate, scaling the CPRI line rate does not change the periodicity of the BF. Instead, this linear scaling of capacity in power-of-two multiples is achieved by changing the word size per BF from 1 byte at the 614.4 Mbit/s rate to 2 bytes at the 1228.8 Mbit/s rate or to 4 bytes at the highest 2457.6 Mbit/s rate.

To build up a larger hierarchical frame, 256 BFs are grouped together to form a multi-frame called a Hyperframe (HF). Each HF is delimited by the use of an 8B/10B K28.5 comma character as the first control word in the first BF of a HF. As mentioned previously, the reuse of an 8B/10B comma character as the CPRI frame alignment word provides a method for time division multiplexing with very low framing overhead bandwidth. The HF structures the 256 control words transmitted within it such that the bandwidth they provide over the link is divided into a set of time multiplexed sub-channels, as illustrated in Figure 9.

Figure 9: CPRI Hyperframe Structure

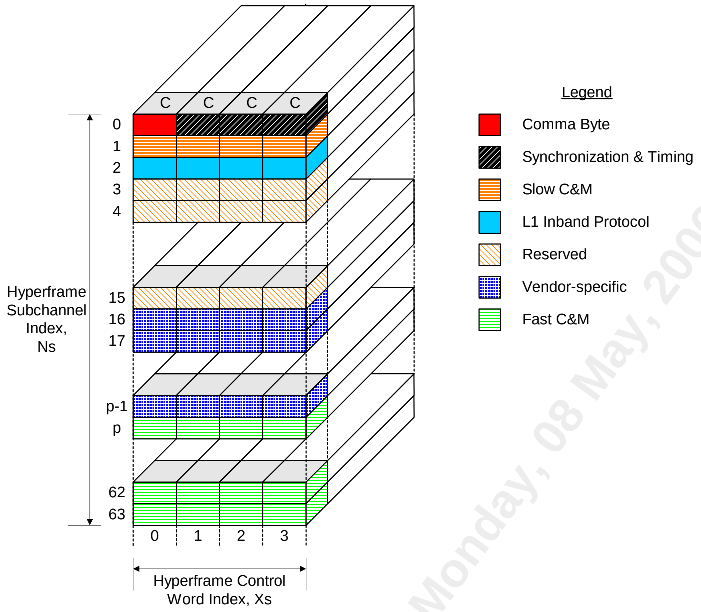

The following summarizes the functionality of each of these sub-channels:

- To align the UMTS baseband data being transported with its associated air interface frame period, air interface alignment data is mapped into a synchronization sub-channel. Each HF carries both a Node B frame number and a number to identify which HF it is of the 150 HFs that span a UMTS 10 ms radio frame. This time-aligned transport from multiple source nodes to multiple sink nodes supports the time-ordered arrival of baseband data for channel combining and the synchronization of air interface transmission over a network of distributed antennas.
- The Slow and Fast Control and Management (C&M) sub-channels support the transfer of OAM&P messaging between the REC and its associated REs using HDLC and Ethernet respectively as data link protocols for reliable transmission. Both sub-channels may be configured to support a range of HDLC and Ethernet bit rates as shown in Table 1. The configured bit rate depends upon the CPRI link rate and the application-specific control bandwidth required by a particular vendor's RE implementation.

Table 1: CPRI Slow and Fast C&M Bandwidth Options

CPRI Line Rate	Slow C&M (HDLC) Bandwidth		Fast C&M (Ethernet) Bandwidth	
(Mbit/s)	Min (kbit/s)	Max (kbit/s)	Min (Mbit/s)	Max (Mbit/s)
614.4	240	480	0.48	21.12
1228.8	240	960	0.96	42.24
2457.6	240	1920	1.92	84.48

- The L1 Inband Protocol sub-channel provides a dedicated mechanism for configuration and maintenance of the CPRI link. Via bit-oriented fields, the L1 Protocol supports the end-to-end exchange of capabilities to determine the bit rates of the Slow and Fast C&M channels. In addition, it allows for the rapid signaling and acknowledgement of critical link status conditions using bit-level fields.
- The Vendor-specific sub-channel is provided for user flexibility, where up to 192 control words in each Hyperframe can be used for vendor-specific purposes. This sub-channel can be used to provide additional capacity for augmenting the existing user plane data and/or control sub-channels. However, the Vendor-specific sub-channel shares its capacity with that of the Fast C&M sub-channel, allowing the user to tradeoff transmission bandwidth in one versus the other. At a minimum, 16 control words are dedicated per Hyperframe exclusively for Vendor-specific sub-channel use.
- To support the addition of new official extensions to the CPRI protocol, 52 control words in each Hyperframe are marked as reserved. New functions associated with these control words may be defined in future releases of the specification.

3.6.2 Delay Calibration

Remotely locating the RF transceiver from the main base station introduces a new complexity associated with overall system delay that the CPRI interface must address. For example, the 3GPP UMTS standards include some very stringent delay constraints to support key features like transmit diversity [3] and user equipment positioning [4]. In the case of transmit diversity, two antennas operating as a diversity pair must have their transmission times aligned such that they are both transmitting the same signal to the air within a maximum acceptable period of one another. For user equipment positioning, if the mobile network can measure the transmission delays between the user's terminal and a number of serving base stations, it can resolve these delays into distances and determine the geographical location of the terminal.

For both features to comply with their respective 3GPP UMTS standards, the system delays incurred in the path between the base station and the user terminal, across all the various segments of processing blocks and transmission media, must be static and measurable to a high degree of accuracy. The introduction of CPRI links to remotely located radio transceivers adds a new extra degree of complexity in managing these system delays. Specifically, it introduces three additional delay components in each direction of CPRI transmission:

- The transport delay through the CPRI transmitter
- The propagation delay over the cable link (electrical or optical)
- The transport delay through the CPRI receiver

The cable propagation delay value will change with different transmission media and cable lengths and is determined by how and where the RE is deployed. The other two transport delay components are determined by the CPRI interface implementation and should ideally have known fixed values. However, in a real-world implementation, while a CPRI transceiver implementation can be designed to minimize transport delays to within a certain degree of accuracy, there will always be factors that introduce small degrees of delay variation. For example, variations can occur for a single transceiver between power-up initializations where the initial board-level startup conditions are different each time. These delay variations are not in themselves a significant problem, as long as they do not change dynamically during the normal operation of the transceiver and are measurable. Therefore, to bound the amount of allowed implementation variation in the transceiver, CPRI has specified the following key delay requirements:

- The one-way delay in the downlink direction of a CPRI link (i.e. from the REC to the RE), excluding the cable propagation delay component, shall be known to an accuracy of better than ±1/32 of a UMTS chip period (Requirement R-19).
- The round-trip delay of a CPRI link (i.e. from the REC to the RE and back), excluding the cable propagation delay component, shall be known to an accuracy of better than ±1/16 of a UMTS chip period (Requirement R-20).

In addition, a CPRI interface solution must ensure that the delay variation between different links does not exceed these tightly toleranced values and that the delay per link is accurately measurable within these tolerances to support link delay calibration and equalization. With a sound implementation foundation based upon fully meeting these requirements associated with the “known” link delay components, accurate measurement of the “unknown” cable propagation delay component can be performed. To this end, CPRI specifies the following requirement:

- The round-trip cable propagation delay of a single CPRI link (i.e. from the REC to the RE and back) shall be measurable to an accuracy of better than ±1/16 of a UMTS chip period (Requirement R-21).

To meet this stringent requirement, a CPRI interface solution must include special dedicated hardware to perform very accurate link delay measurements.

3.7 CPRI Mapping and Switching Layers

The CPRI data link layer provides the previously described sub-channels for transporting the various higher-level information flows over the end-to-end connection. Above this layer, the CPRI mapping layer specifies how radio data, control and synchronization applications insert and extract their respective information flows from these sub-channels.

3.7.1 User Plane Data Mapping and Switching

The major consumer of CPRI link bandwidth is the user plane data associated with the radio application of the RE. User plane AxC data is mapped as IQ baseband samples into the payload of each BF. To provide flexibility for vendors, sample widths of different sizes are supported – from 4 to 10 bits per sample in the uplink direction and from 8 to 20 bits in the downlink. In addition, either two times or four times oversampling may be used in the uplink direction.

The mapping of each AxC is performed in two stages. Firstly, for each IQ sample, the I sample and the Q sample are bit-interleaved. Secondly, the bit-interleaved IQ samples associated with each supported Antenna-Carrier (AxC) are multiplexed into fixed timeslots in the BF payload. This two-stage mapping process is illustrated in Figure 10 below for an example mapping of 10-bit IQ samples with no oversampling into a BF operating at the 614.4 Mbit/s line rate:

Figure 10: CPRI User Data Plane Mapping Options

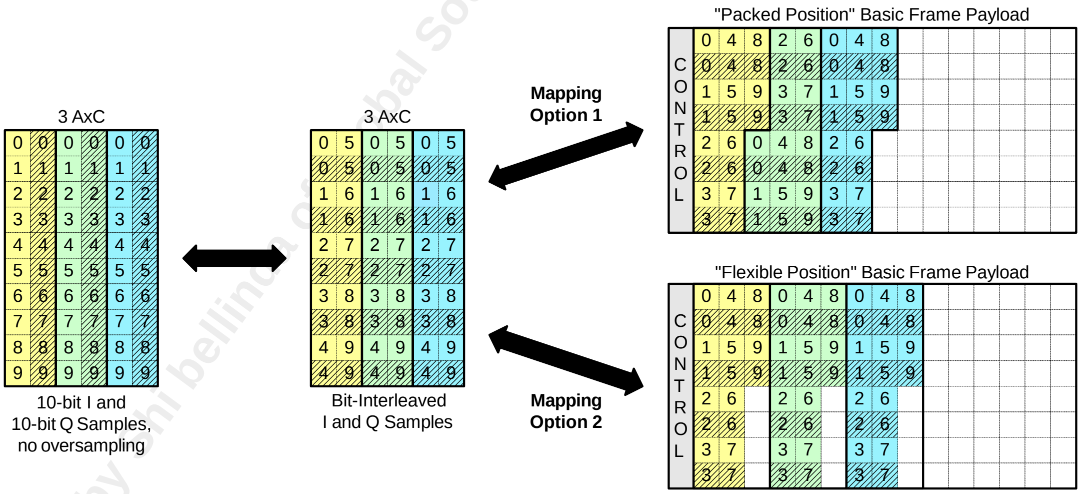

The location of a specific AxC stream is uniquely defined by its fixed timeslot location in the BF payload. In mapping the AxC blocks to each BF, CPRI supports two options:

- A “packed position” mapping that allows AxCs to be mapped contiguously with no white space (reserved bits) between them. This option maximizes the payload bandwidth available for transporting AxCs at the expense of making the timeslot mapping more complex.

- A “flexible position” mapping where AxCs are mapped to non-contiguous timeslots with white space between them. This option does not make best use of the available payload bandwidth but it can provide other system flexibilities to the radio application e.g. byte-aligned timeslot mappings for application processing simplicity or the use of timeslot optimal locations that minimize buffering requirements.

Using these transmission rules creates specific periodic timeslots in the CPRI frame structure for the time-ordered delivery of IQ data associated with a specific AxC. Although not explicitly defined in the CPRI specification, the use of this regular transport structure supports the switching of AxC containers. For networked topologies where REs are connected in chains or rings, each RE must act as a switching element where AxC data associated with that RE is added and dropped from the shared traffic flow and AxC data associated with other REs is transparently forwarded through the node. The time-division multiplexed approach to transporting AxCs allows this add/drop/continue switching to be implemented efficiently with minimal latency.

All the above AxC mapping functionality is defined within CPRI in terms of UMTS radio data as the specification focused on supporting WCDMA as its initial target application. However, these basic mapping principles are generic and do not preclude CPRI from being re-used as a transport vehicle to support other air interface standards such as TD-SCDMA and WiMAX.

3.7.2 C&M Mapping and Switching

Control data for transmission over the Slow C&M sub-channel is first encapsulated in standard HDLC frames. These HDLC frames are mapped transparently to the channel in a bit-oriented manner, without using any byte alignment. In this way, the high-level OAM&P application that is sourcing and sinking control data can use standard HDLC flow control, error checking and retransmission techniques to manage the end-to-end exchange of data and to perform rate matching of the information flow to the bandwidth provided by the Slow C&M sub-channel.

Where networking of multiple REs is required, routing of the HDLC-framed control traffic must be performed at each intermediate RE. Here, data addressed to the local RE is added and dropped and data destined for subsequent nodes is forwarded. CPRI does not provide specific direction as to how this networking should be performed but two possible alternatives are:

- HDLC-based Layer 2 frame switching, where the addressing mechanisms provided by the HDLC frame structure are used to multiplex traffic to and from multiple REs.
- IP-based Layer 3 packet switching, where the HDLC framing is used only to provide point-to-point data links between nodes. The networking of the nodes is then performed through addressing and routing at the higher IP layer.

Both options have benefits that make them attractive solutions. HDLC-based frame switching has the advantages of requiring less networking software and being more bandwidth efficient but it will use proprietary mechanisms that vary from implementation to implementation. On the other hand, although requiring more software complexity, an IP-based solution will use standard protocol stacks that offer many advanced networking features. In addition, an IP-based approach can also be shared with the higher bandwidth Ethernet-based Fast C&M sub-channel.

3.7.3 Fast C&M Mapping and Switching

The Fast C&M sub-channel offers a higher bandwidth alternative to the Slow C&M sub-channel. Control data for transmission over the Fast C&M sub-channel is encapsulated in standard Ethernet Media Access Controller (MAC) frames. These MAC frames are 4B/5B encoded before transmission, as per the 100BASE-X Physical Coding Sublayer (PCS) of IEEE Standard 802.3-2002, such that the CPRI sub-channel mimics the functionality of a dedicated Fast Ethernet physical layer device to the higher-level Ethernet MAC processing. The 4B/5B-encoded data stream is mapped in a bit-oriented manner to the sub-channel but no flow control mechanism is specified. Therefore, each solution must implement a mechanism for adapting the standard rate of the Ethernet data source to the non-standard bandwidth of the Fast C&M sub-channel.

As per the Slow C&M sub-channel, routing of the Ethernet-framed control traffic must be performed at each intermediate RE in a networked topology. However, the use of Ethernet means that networked REs can take advantage of readily available hardware- and software-based Layer 2 switching technology to perform forwarding based upon MAC addresses.

3.7.4 L1 Inband Protocol Mapping

The L1 Inband Protocol sub-channel is a dedicated point-to-point communications path that uses the low-level Hyperframe bit fields illustrated in Figure 11 below:

Figure 11: L1 Inband Mapping

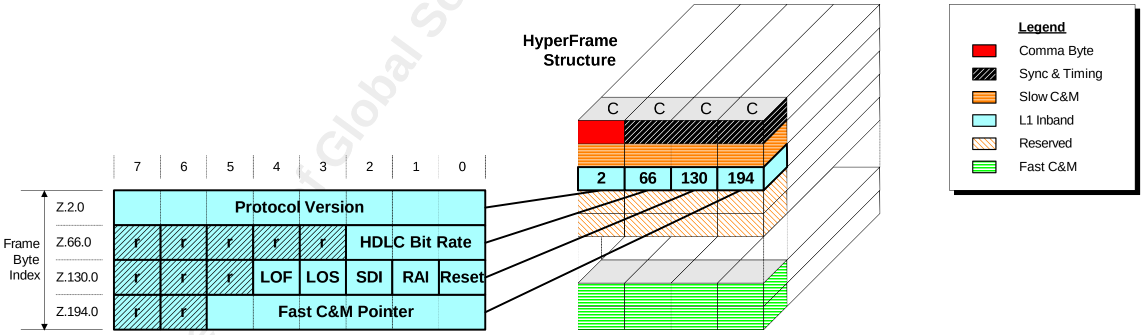

It provides the following functions associated with control and maintenance of the CPRI link:

- To communicate the version of the CPRI protocol supported at either end of the link
- To configure the operational bit rates of the Slow and/or Fast C&M channels
- To rapidly signal and acknowledge local and remote CPRI link error conditions, for example a Loss of Signal (LOS) failure or a Remote Alarm Indication (RAI)
- To force a reset of the remote RE should all other communication paths fail

4 Summary

Carriers and OEMs are under tremendous pressure to reduce the overall costs associated with base station development, deployment, and operation. The distributed base station concept is a revolutionary new approach to building and deploying the next generation of wireless access networks. The system flexibilities it offers to service providers in terms of addressing challenging problems like cell site acquisition, power consumption and network scalability give it a unique and compelling value proposition.

By developing and promoting a digital baseband interface with all the application-specific features required for base station synchronization and control, the CPRI organization has created an open standard that is already facilitating the development of equipment to meet the new deployment paradigm of distributed base stations. The adoption of this common interface over the next few years promises to enable a new standards-oriented approach to building base stations that can both sustain and invigorate the industry.

PMC-Sierra is excited to support the CPRI initiative and is already taking an active role by offering the standard-compliant BRIC™ chipset for the distributed base station architecture. The PM7830 BRIC-6 and PM7832 BRIC-2 are full-featured 6-port and 2-port termination devices that fully support the CPRI specification for wireless base station interconnect. PMC-Sierra's interface solutions make it possible for OEMs and ODMs to offer all defined RRH network topologies with a standard digital interconnect for distributed base station systems. The BRIC product family:

- Accelerates the adoption of distributed wireless base station architectures, which dramatically reduces capital and operational expenditures for carriers;
- Supports multiple air interfaces, including WCDMA, CDMA2000, TD-SCDMA, and WiMAX over CPRI with a single interface solution; and
- Reduces customer development and bill-of-material (BOM) costs as well as qualification cycles.

For more detailed information on the BRIC-2 and BRIC-6 devices, please see the online PMC-Sierra Product Directory at http://www.pmc-sierra.com/products/.

5 References

1. CPRI Specification V2.0, “Common Public Radio Interface (CPRI); Interface Specification”, October 2004.

2. CPRI Specification V1.0, “Common Public Radio Interface (CPRI); Interface Specification”, September 2003.

3. 3GPP TS 25.104, “Base Station (BS) radio transmission and reception (FDD)”, Release 5, V5.8.0, December 2003.

4. 3GPP TS 25.133, “Requirements for support of radio resource management (FDD)”, Release 5, V5.11.0, June 2004.

Glossary

AxC	Antenna Carrier
BTS	Base Transceiver Station
C&M	Control and Management
CPRI	Common Public Radio Interface
HDLC	High-level Data Link Control
HV	High Voltage
IEEE	Institute of Electrical and Electronic Engineers
IQ	In Phase & Quadrature
LV	Low Voltage
MAC	Media Access Control
OAM&P	Operations, Administration, Maintenance and Provisioning
PHY	Physical Layer
RE	Radio Equipment
REC	Radio Equipment Control
RF	Radio Frequency
RRH	Remote Radio Head

Contacting PMC-Sierra

PMC-Sierra
100-2700 Production Way
Burnaby, BC
Canada V5A 4X1

Tel: +1 (604) 415-6000
Fax: +1 (604) 415-6200

Document Information: document@pmc-sierra.com
Corporate Information: info@pmc-sierra.com
Technical Support: apps@pmc-sierra.com
Web Site: http://www.pmc-sierra.com