Evolution of CMTS/CCAP Architectures

Karthik Sundaresan
CableLabs

Abstract

The CMTS has evolved over time, becoming a CCAP supporting both high-speed data and video services. As bandwidth capacity needs grow rapidly, there is increasing pressure on the infrastructure, and new Distributed CCAP Architectures are emerging to address the need for smaller scope, more flexible options.

Distributed CCAP implementations can bring significant benefits in certain HFC network deployments. It will enable higher PHY layer performance for DOCSIS 3.1, and reduce space and power needs at the headend. The paper describes various CCAP architectures and how they handle data and video.

BASELINE ARCHITECTURE

Cable operators implement and deploy IP High Speed data, Linear Broadcast Video, Video on Demand, Voice, and various other integrated services to their end customers. The primary choice of access technology for this today is DOCSIS over a hybrid-fiber/coax (HFC) cable network.

Data Service

The DOCSIS system allows transparent bi-directional transfer of Internet Protocol (IP) traffic between the cable system head-end and customer locations over the HFC network. The Cable Modem Termination System (CMTS) is the central platform in enabling High speed Internet connectivity over the Cable HFC network. The CMTS platform has evolved over time along with the DOCSIS Specifications. The CMTS provides the MAC and PHY layer connection to the Cable Modem (CM) at customer premise.

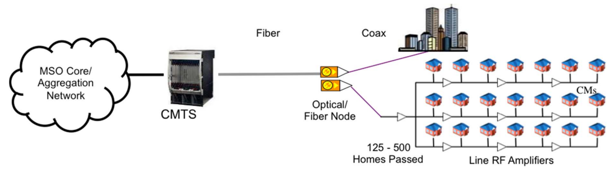

DOCSIS Access Network

The CMTS consists of various logical functional components, at a high level these are as follows:

- DOCSIS PHY Layer
  - Upstream Receiver
  - Downstream Transmitter
- DOCSIS MAC Layer
  - Upstream MAC and Scheduler
  - Downstream MAC Processing
  - DOCSIS QoS
  - Security
- RF Output block
- L2 forwarding block
- L3 forwarding block
- IP Processing for DHCP
- SNMP agent / CLI, etc.

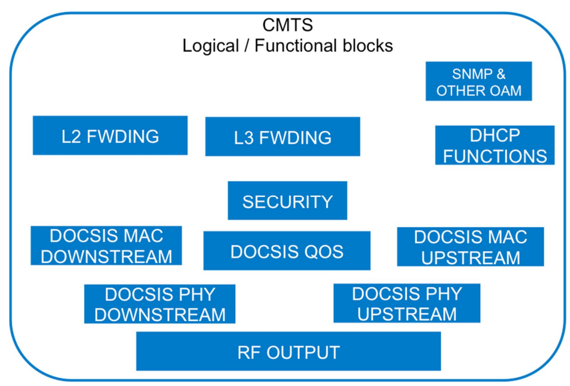

CMTS Functional Blocks

Video Distribution

The Video EdgeQAM is a key piece of equipment in any headend, hub site to enable video services. The role of the EQAM in the video-on-demand and switched-digital-video architecture is to receive an IP unicast or multicast stream containing MPEG transport stream packets, and then produce that transport stream on one or more RF outputs for transmission over the HFC cable plant. The EQAM contains the multiple Gigabit Ethernet (GigE) data inputs. The content sent to an EQAM over the GigE interface is framed in an MPEG-2 transport stream (TS). An EQAM will support both Single Program Transport Stream (SPTS) and Multiple Program Transport Streams (MPTS). The EQAM can also re-stamp PCR timestamps for de-jitter processing, which helps reduce network impairments.

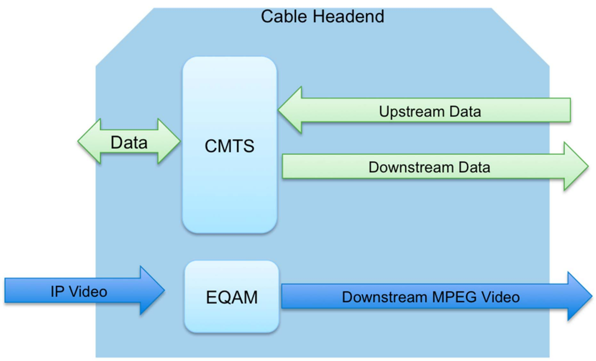

Baseline Architecture, Data & Video

MODULAR HEADEND ARCHITECTURE

Over time there have been many steps in the evolution of the CMTS platform. One of the first steps was the creation of the Modular Headend Architecture (MHA), which essentially separated out the DOCSIS downstream PHY layer out of the CMTS and moved it to a separate EQAM device.

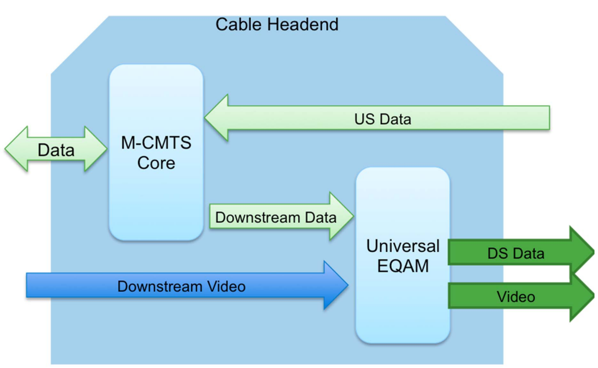

Modular Headend Architecture

A new interface called the DEPI (Downstream External PHY interface) was defined to support sending the data from the CMTS Core to the EQAM. The idea was to reuse the EQAM to modulate the bits on to the wire for both downstream DOCSIS data as well as MPEG video. The video EQAM now becomes a universal EQAM handling both video and DOCSIS data as inputs. The upstream receiver remains at the CMTS core. Since the DOCSIS MAC and PHY were separated in the MHA architecture, a new DOCSIS Timing Interface was introduced to keep the two devices closely synchronized. The Modular CMTS was essentially a two separate platforms as compared to the fully integrated CMTS.

INTEGRATED CCAP

The next big step in the evolution of CMTS platforms was the Converged Cable Access Platform (CCAP). The CCAP was intended to provide a new equipment architecture option for manufacturers to achieve increased Edge QAM and CMTS densities that MSOs require. The CCAP leverages existing technologies, including DOCSIS and also newer technologies such as Ethernet optics and EPON (Ethernet Passive Optical Network). The CCAP unifies the CMTS, Switching, Routing, and QAM functions at the headend, so that all data, video, voice functions can be handled over IP before conversion to RF or Optical signals. The CCAP eliminates the need for the combiner functionality in the headend.

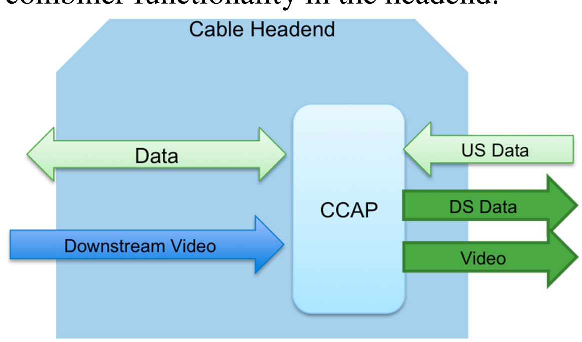

Integrated CCAP

DOCSIS 3.1

DOCSIS 3.1 is the next generation in the evolution of DOCSIS and it brings some fundamental changes to the technology. DOCSIS 3.1 introduces Orthogonal Frequency Division Multiplexing (OFDM) as the new PHY layer technology, and allows for wide channels from 24 MHz to 192 MHz wide, moving away from the legacy 6 MHz sizes. DOCSIS 3.1 introduces Low Density Parity Check (LDPC) based Forward Error Correction (FEC).

DOCSIS 3.1 also introduces options for additional HFC spectrum, by allowing for expansion of the US Split (Mid-split or High-split). It also allows expansion of the downstream spectrum at the higher end. DOCSIS 3.1 LDPC FEC allows data transmission approaching the theoretical limits, and this enables 50% more efficient modulations such as 4096 QAM, harnessing more capacity on existing HFC networks. In the long run, DOCSIS 3.1 could support data rates reaching 10 Gbps downstream and 1 Gbps upstream.

CABLE OPERATOR CHALLENGES

Cable operators today face numerous challenges. The customer and market demands for higher bandwidth and data rates are increasing at a fast pace, and the demand for data and video services is forcing operators to continuously upgrade the plant capacity. As bandwidth demand continues to grow, facility space, power, and other related factors start becoming a concern. Below are some issues that factor into the discussion.

Analog optical networks between the hub and the fiber node:

- The length of optical link can be a limiting factor in managing the SNR performance.
- For DOCSIS 3.1 technology, capacity in the plant can be optimized without the analog optical noise floor.

Digital optics:

- Digital optics promise reduced OPEX as compared to analog optics. For digital optics, MSOs can use low-cost small form-factor pluggable (SFP) lasers instead of high-priced distributed feedback (DFB) lasers used in AM fiber. In this manner, they avoid the operations costs of maintaining rigorous performance on AM links as plant conditions change.
- The increased reach of the technology can be used to reduce facilities expense.
- Digital optics also provide a lower CAPEX for capacity growth. With a shift to digital optics, the throughput on the fiber becomes much greater (wavelengths can be packed much closer spectrally in the Dense Wavelength Division Multiplexing (DWDM) than with WDM spacing for analog optical signals. Typical spacing between wavelengths in AM mode is 100 GHz, which allows up to 40 wavelengths on a single fiber, whereas spacing for digital wavelengths is 25 GHz or less, which leaves room for 160 or more wavelengths. DOCSIS 3.1 modulations and/or capacity can be increased with no linear optics noise, and better SNR.

Facilities costs (space, power, HVAC):

- Increasing cost of distribution hub facilities cost is a concern for operators.
- As the power consumption in the hub escalates, distributed CCAP architectures allow reduction of facilities and power costs associated with hub and headend facilities.

Node + 0 architectures:

- There is a need to move to smaller service groups and to N+0 architectures to increase the bandwidth available to the user. The number of service groups needed might increase by a large factor with Fiber Deep HFC architecture and the question is if the current CMTS platforms can keep up.

Low cost deployments:

- In some global markets, the need for low-cost smaller-scope deployment options are important to enable the success of cable broadband.
- Chinese MSOs in particular face unique challenges. Their plant architecture consists of a digital optical packet network (using point-to-point, EPON, GPON, etc.) to the MDU (Multi Dwelling Unit) and coax within the MDU. A centralized CMTS is not as economical in this architecture.

Smaller CMTS step functions:

- The industry needs smaller increments in cost and capacity when increasing capacity for just a few more users. The cost of adding a adding an entire new CMTS/line card for a minor increase in needs is sometimes cost prohibitive.

Flexibility and other benefits

- Distributed architectures move portions of the CMTS functionality into nodes; this gives operators flexibility and another option in their tool belt for deploying DOCSIS.
- There are also various benefits of a high rate digital backhaul which can be used for other services, such as wireless backhaul, business fiber extensions, and so forth.

Timing:

- To get to the higher bandwidth channels allowed DOCSIS 3.1 in the upstream, an operator may be considering moving the US-DS split in the cable plant. This could be a disruptive change to the customers and would be a big effort for the operator. The same challenges exist for a Distributed Architecture deployment, which also will need to touch the fiber nodes.
- If the operators decides to make both changes, they could benefit by timing both the upstream split and deploying a distributed solution together.

As a result, MSOs today are investigating distributed CCAP solutions, in addition to leveraging digital optics in the access network. The question is whether the benefits of moving away from AM optics outweigh the costs and risks of placing processing and RF hardware in an environmentally hardened chassis on the outside plant or in an MDU basement.

DISTRIBUTED CCAP ARCHITECTURES OVERVIEW

There are multiple distributed CCAP architectures are emerging in the marketplace. The basic idea is to distribute some or all of the functionality of the CMTS/CCAP down to a remote location, like the Fiber Node.

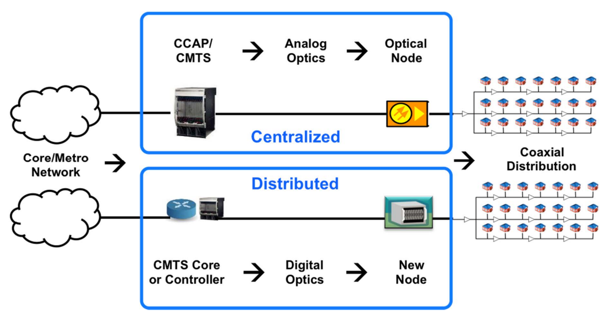

Digital Optics and Shifting of Functionality

There are 3 distributed architectures that have come forth so far. These are the Remote PHY, Remote MAC-PHY and the Split-MAC variations.

The concept behind Remote MAC-PHY is to move the entire CMTS/CCAP into the Remote node. The idea behind the Remote PHY is to split the CMTS between the MAC and the PHY Layers and move the PHY layer to the Remote Node.

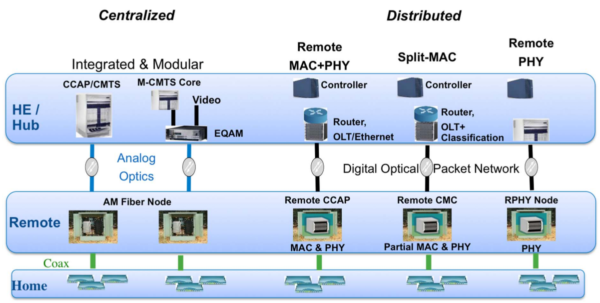

Centralized and Distributed Architecture Summary

The Split-MAC is in between the above two options, where some of the MAC functionality is defined and left at the headend and the remaining MAC and PHY functionality are moved to the Fiber Node. This option is defined in the C-DOCSIS System specification. The rest of this paper is focuses on the Remote MAC-PHY and Remote PHY options, as these have gained more traction in the American, European and China markets, whereas the Split-MAC has gained limited traction in the China market.

The foundation for Distributed CCAP Architectures is a digital optical plant, which essentially makes the connection between the head-end and the fiber node a Layer 2 Ethernet connection. In a digital HFC plant, the fiber portion utilizes a baseband network transmission technology such as Ethernet, EPON (Ethernet over Passive Optical Networks), GPON (Gigabit Passive Optical Network), or any layer 2 technology that would support a fiber-based PHY layer.

REMOTE PHY

The Remote PHY technology is also know as MHAv2 (Modular Headend Architecture version 2) as in many ways it builds on the original MHA architecture. MHAv2 uses a Layer 3 pseudowire between a CCAP Core and a set of Remote PHY devices. One of the common locations for a Remote PHY device is the optical node device that is located at the junction of the fiber and coax plants.

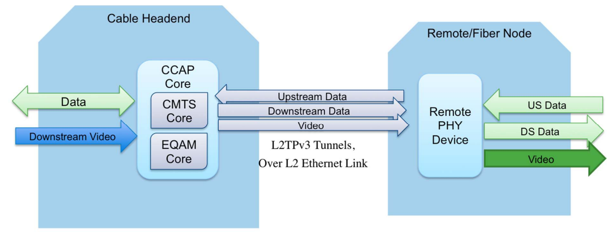

Remote PHY Architecture

In a Remote PHY System, the integrated CCAP is separated into two distinct components. The first component is the CCAP Core and the second component is the R-PHY Device (RPD). The CCAP Core contains both a CMTS Core for DOCSIS and an EQAM Core for Video.

The CMTS Core contains the DOCSIS MAC and the upper layer DOCSIS protocols. This includes all signaling functions, downstream and upstream bandwidth scheduling, and DOCSIS framing. The DOCSIS functionality of the CMTS Core is defined by the existing DOCSIS Specifications. The EQAM Core contains all the video processing functions that an EQAM provides today.

The Remote PHY Device contains mainly PHY related circuitry, such as downstream QAM modulators, upstream QAM demodulators, together with pseudowire logic to connect to the CCAP Core. The RPD platform is a physical layer converter whose functions are:

- To convert downstream DOCSIS, MPEG video and OOB signals received from a CCAP Core over a digital medium such as Ethernet or PON to analog for transmission over RF or linear optics.
- To convert upstream DOCSIS, and OOB signals received from an analog medium such as RF or linear optics to digital for transmission over Ethernet or PON to a CCAP Core.

REMOTE MAC PHY

The Remote MAC PHY technology moves both the DOCSIS MAC and PHY layers down to the Fiber Node. The connection between the Headend and the node is essentially a Layer 2 Ethernet connection. There are two options for this, which are different based on how video is handled. In both cases the data forwarding CMTS functionality is at the remote node. A compact CMTS is deployed at the fiber node and the CMTS NSI connects through the digital optical network back to the cable headend. For video, there are two options, as described below: Remote CCAP or Remote CMTS+Divided EQAM.

Remote CCAP

The RemoteCCAP term applies to an architecture where both the data and the video functions are moved to the Remote node. The CMTS functionality and the EQAM functionality are completely moved to the Fiber Node, and hence the term Remote CCAP. The Video and data transit the L2 Ethernet link like any other IP traffic. The Video needs to be encrypted to protect from unauthorized access at the Remote node.

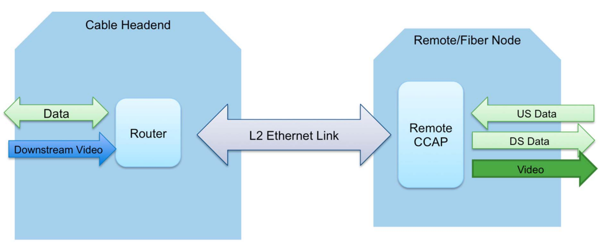

Remote MAC-PHY Architecture: Remote CCAP

Remote CMTS + Divided EQAM

The Remote CMTS term applies to an architecture where only the data/CMTS functionality is moved into the remote node. The video/EQAM functionality is divided between the headend and the remote node, as in the Remote PHY architecture. The video MPEG packet processing is handled in the headend by an EQAM core device and the EQAM PHY inside the Remote CMTS handles the modulation of the video onto the wire.

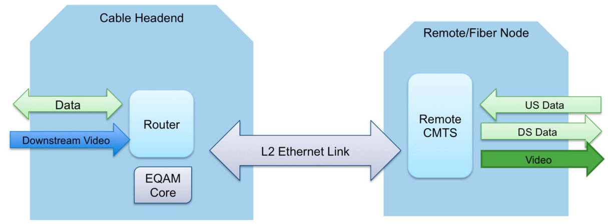

Remote MAC-PHY Architecture: Remote CMTS + Divided EQAM

COMPARING DISTRIBUTED ARCHITECTURES

Due to the digital optical plant needed, any of the distributed architectures have PHY layer performance gains, which are very helpful in getting to the higher order modulations as defined in DOCSIS 3.1 (1024 QAM/4096 QAM). The digital fiber link can give 3–7dB gains, which can bump up the possible modulation order, increasing capacity.

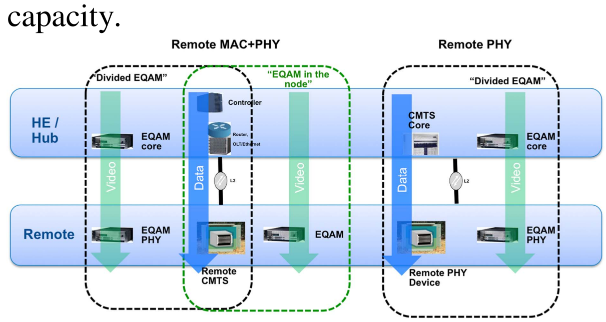

Remote MAC-PHY and Remote PHY Architecture Comparison

The Remote PHY architecture keeps the Remote Node device relatively simple, with only the PHY layer modulation in the Remote Node. There is also added communication between the CMTS MAC and PHY layers to setup the connections as needed, but this is needed only at startup.

In Remote MAC-PHY architectures the CMTS is condensed and sized for a node, although now with the MAC processing adds more intelligence to the device. But given that the CMTS design is known and has been used for a long time, the operators may prefer this familiar provisioning and management environment.

SYNERGIES WITH VIRTUALIZATION

As computation power available in the microprocessor platforms today increases, the cost of Common-Off-The-Shelf (COTS) servers is going down dramatically. This is leading to a huge boom in virtualization, where operators are virtualizing various functions within their networks. Considering this new functionality, which can be implemented in the customer premise, instead of implementing new ASIC designs and rolling out a new CPE device, operators now have the ability to create those services on the cloud and direct user traffic through each of those virtual functions.

Similar ideas on virtualization will affect the CCAP devices and their implementations. The main PHY layer functionality, which modulates the bits onto the wire need to be on the cable plant. But essentially all other functionality could become candidates for virtualization.

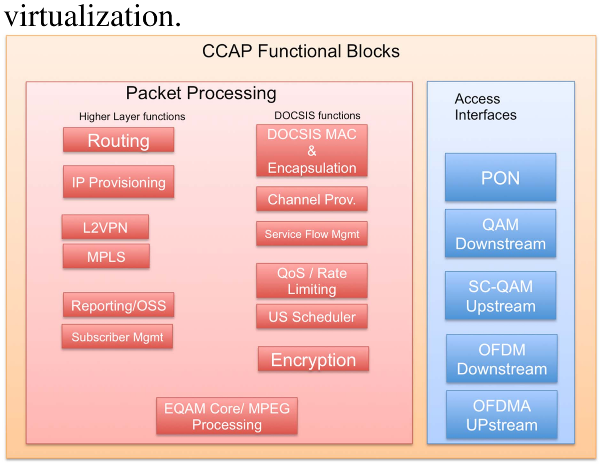

Virtualization Candidates within a CCAP

In a distributed CCAP architecture there are various pieces, which naturally gravitate towards virtualization. It is possible to run many of the CCAP packet-processing functions inside a virtual machine in the headend. This can start of with some of the easy candidates (labeled *Higher layer functions* in the diagram above) such as routing functions, IP/DHCP provisioning functions, OSS and reporting functions, and subscriber management functions.

For a Remote MAC-PHY architecture, the above set of features may be all that can be achieved, as the remaining functions would already be implemented in the hardware at the Remote CMTS device. In the Remote PHY architecture, since the MAC level functionality is being implemented in the CMTS -Core at the headend, many of the higher-level functions as well some of the core DOCSIS functions (as labeled in the diagram) can be virtualized and be run from a remote data center. The main concerns at this point are the latency delays of such architecture. Whether those problems are surmountable remain to be seen. If the trend towards virtualization persists, and the latency problems are ironed out, one can imagine operators in the future running lightweight DOCSIS CCAPs in the cable plant. The CCAPs will act as a media converter, and all the intelligent functionality of the CCAP will be implemented and controlled from the cloud.

EVOLUTION OF HFC NETWORK

This paper has described the various CMTS/CCAP architectures seen by the cable industry so far. The Integrated CMTS/CCAP is a staple of the operators and that platform here is to stay for quite some time as it meets many of the current needs. Driven by specific future needs and competitive pressures in different parts of the plant, the cable operators will assess their need to go to Distributed CCAP architectures. The choice between technologies is complex and involves numerous tradeoffs.

In the future, operators will be transitioning most of their services including video to run over IP/DOCSIS. The different CCAP architectures, either the integrated or the distributed platforms, will be used in in different parts of the network. The platform, which helps the operators gain the most bandwidth and operational efficiencies for the cost and help in the IP transition of services, will become the favored approach.

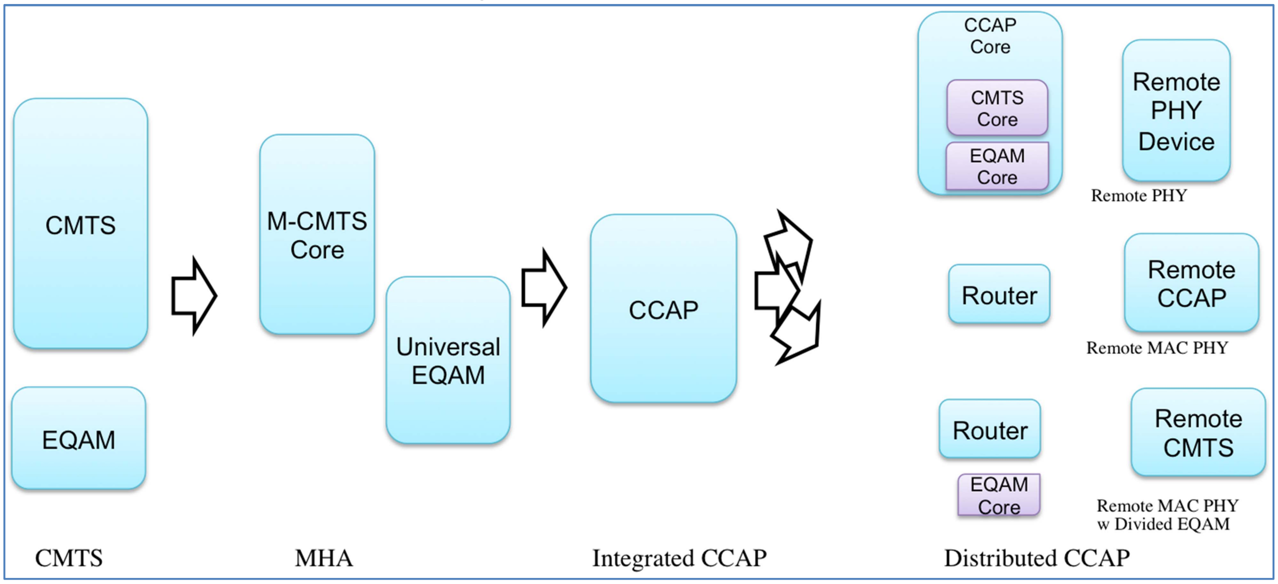

Evolution of the CCAP Architecture