R-PHY with Remote Upstream Scheduler

A Technical Paper prepared for SCTE•ISBE by

Tong Liu, PhD
Principal Engineer
Cisco Systems Inc
300 Beaver Brook Road, BOXBOROUGH, MA 01719
978-936-1217
tonliu@cisco.com

John T Chapman
CTO Cable Access and Cisco Fellow
Cisco Systems
170 W Tasman Dr, San Joe, CA 92677
408-526-7651
jchapman@cisco.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Introduction

The cable access network is undergoing a radical transformation from the traditional integrated CCAP architecture to a distributed access architecture (DAA), driven by the growing capacity crunches and the cost pressures to deliver gigabit broadband services. With DAA, cable operators are able to push fiber deeper and replace legacy fiber nodes with DAA devices, achieving higher capacity with both better signal quality and reduced service group sizes.

Depending on how the CCAP MAC and PHY functions are separated, there are two basic architectural options to DAA, Remote PHY (R-PHY) and Flexible MAC Architecture (FMA).

In the R-PHY architecture, the PHY element is removed from the CCAP core and added to the fiber node as a Remote PHY Device (RPD). The basic design philosophy is to put the least amount of hardware and software at the endpoints and keep the complexity centralized. It also allows operators to leverage existing CCAP functions as much as possible for a fast and seamless transition to DAA with both integrated PHY and the Remote PHY potentially connected to the same CCAP core.

The FMA, on the other hand, moves both the CCAP MAC and PHY elements to the node, either as an integrated Remote MAC-PHY Device (RMD) or a combination of Remote MAC Core (RMC) and RPD. Essentially, an FMA DAA device is a small scale CMTS without the routing and management functions. Compared to R-PHY, the DAA device requires significantly more hardware and software functions, and inevitably imposes design and deployment challenges when the device is constrained by power and cost. Moreover, given the complexity of the DOCSIS MAC layer, FMA requires a fairly comprehensive standard interface for the upper network layer and management applications to talk to the DAA devices.

From the latency point of view, one architectural difference between FMA and the R-PHY today is the location of the upstream (US) scheduler. With all the MAC layer functions centralized at the CCAP core, R-PHY has been using a centralized US scheduling scheme that requires the request (REQ) and grant (GNT) information to be exchanged across the Converged Interconnect Network (CIN). CIN delay has no impact on the US scheduling latency as long as it is not the dominating factor, which is the case when GNTs are carried in the de facto 2 millisecond (ms) MAPs, and the CIN distance is within the normal 100-mile (160 km) DOCSIS operational range assumed for I-CMTS deployment [1].

As the network keeps transitioning to DAA, there are, however, reported cases where the CIN is stretched beyond the 100-mile mark, for reasons such as hub-side consolidation that relocates a CCAP core to the central headend or a regional data center. Meanwhile, driven by new low latency applications, like cloud gaming and mobile xHaul [2], the DOCSIS REQ-GNT protocol is being tightened to shorter MAPs, such as 1 millisecond MAPs, on DOCSIS 3.1 OFDMA channels [3]. In such circumstances, the CIN delay could be exposed as a significant factor in the REQ-GNT latency equation.

The reason why FMA is immune to CIN delay is because the US scheduler, the MAC element that handles REQ-GNT, is co-located with the PHY where the REQ is received; while in the R-PHY case, the US scheduler is at the core, separated from the RPD across the CIN. This realization leads to the question to be tackled in this paper: Is it possible to put a remote US scheduler at the RPD to help with the latency sensitive REQ-GNT processing?

The remote US scheduler idea was actually considered way back at the beginning of the R-PHY design and development and is mentioned as an option in the current R-PHY specification. However, since the initial R-PHY deployment goal was to replace I-CMTS, it was deferred as a future enhancement. Now the time has come to move forward with the remote US scheduler design to provide the low-latency scheduling (LLS) needed for long-distance R-PHY deployment.

R-PHY with a remote US scheduler adds a new DAA scenario as shown in Figure 1. Latency-wise, it is equivalent to FMA, however, with much less cost and complexity. It offers FMA-lite functionalities with R-PHY's efficiency and simplicity.

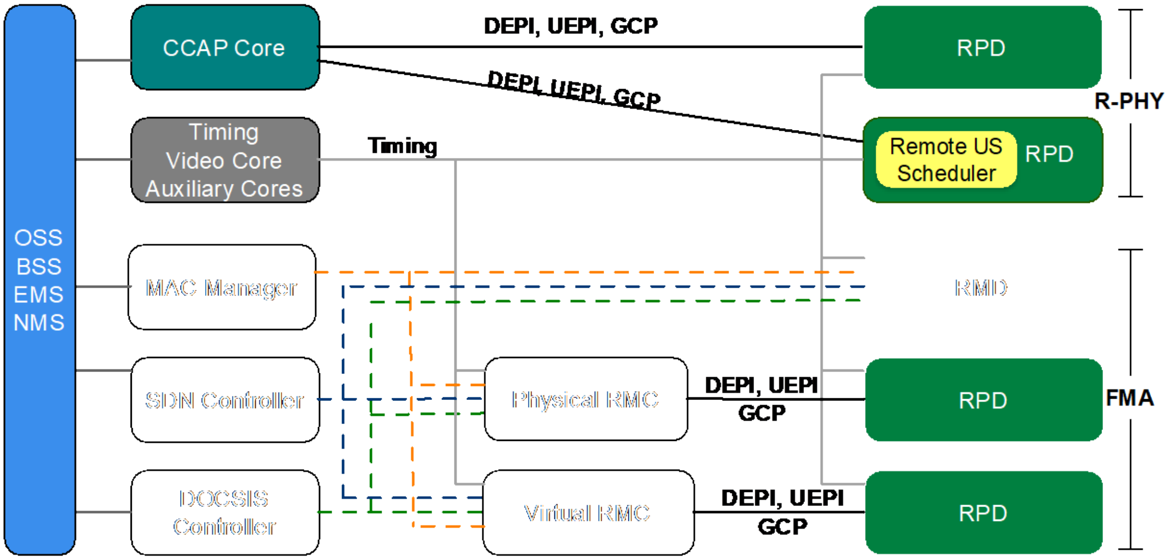

Figure 1 - DAA Options Including R-PHY with Remote US Scheduler

Since the remote US scheduler is internal to the RPD, it can leverage the RPD hardware and software platform and the established forwarding, control and management plane interfaces such as R-UEPI, R-DEPI and GCP. The remote US scheduler APIs will be Yang data model-based and will be able to take advantage of the new control plane infrastructure proposed for R-PHY2.0 [4].

This paper is organized as follows. Section 1 explains the rational for splitting the upstream scheduling between the core and the RPD with a remote RPD US scheduler. Section 2 defines the DOCSIS upstream service model and the remote upstream scheduler service categories. Section 3 examines the split upstream scheduling impact on R-UEPI, and Section 4 examines the impact on R-DEPI. Section 5 discusses the remote US scheduler APIs using the data modeling approach. Section 6 explains how scheduling locality optimization may be applied to improve latency and efficiency for a R-PHY system as a whole. Finally, the paper will be concluded by summarizing the study highlights.

Content

1. R-PHY US Scheduler Location Options

Since the beginning of the R-PHY architecture, there has been a technical debate as to where the US scheduler should be placed. Should it be in the CCAP core where the rest of the software is or should it be in the RPD with the US PHY? To answer this question, there are both business and technical reasons to consider when choosing one location over the other.

From a technical point of view, latency is the main consideration in comparing the two location options. In this perspective, R-PHY with a centralized US scheduler is equivalent to I-CCAP when operating at 2 ms MAPs over a 100-mile plant. R-PHY with a remote US scheduler is expected to provide better latency when operating at shorter MAPs and across a longer CIN distance.

Figure 2 depicts centralized vs. remote US scheduler deployment scenarios and the traverse paths of the REQ-GNT(MAP) messages.

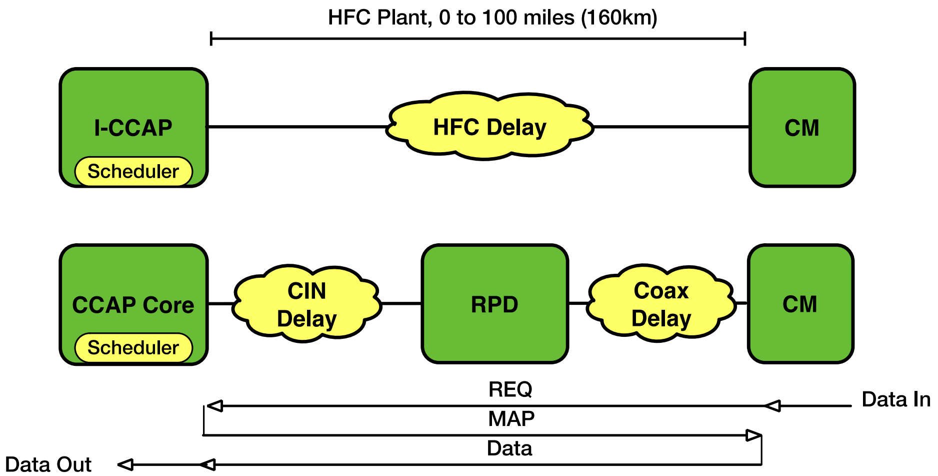

(a) Centralized US Scheduler Scenarios

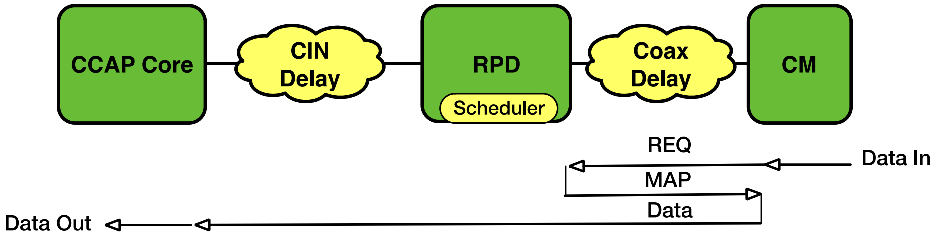

(b) Remote US Scheduler Scenario

Figure 2 - Centralized vs. Remote US Scheduler Deployment Scenarios.

1.1. The Case for a Centralized US Scheduler

The basis of the R-PHY architecture is to move the PHY and replace the analog optical link between the CCAP and the node with a digital link. Just with this initial step, cable operators would be able to get better SNR performance, pull the fiber deeper, rebuild the plant and cut a large N+M service group into much smaller ones. All these can be achieved by simply moving the PHY element out of the CCAP core, while keeping all MAC elements including the DOCSIS US scheduler centralized. This also allows operators to leverage the existing CCAP MAC functions to simultaneously support both integrated PHY and remote PHY for a seamless transition to DAA.

Besides the business reasons, the main technical reason for picking centralized US scheduling is based on whether the CIN delay is the dominating factor affecting the request-grant (REQ-GNT) latency.

DOCSIS uses a REQ-GNT protocol to arbitrate US channel access among CMs. When US data arrives at a CM, the CM sends a REQ to the CCAP US scheduler. The US scheduler arranges an individual transmission opportunity by encoding a GNT for the requesting CM in the bandwidth allocation map (MAP) message and broadcasts it to the listening CMs. For regular sized MAPs, from the CM point of view, at any time, there is always one MAP in use, and one MAP "on deck" that is about to be used. So the shortest possible latency for the CM to receive the GNT from the CCAP is every other MAP, which translates to two MAP intervals in time. Most of the time, the REQ needs to wait at the CMTS for the next MAP to be built, with an average wait time around half of the MAP interval. As a result, the total REQ-GNT protocol minimum delay with a 2 ms MAP interval is about 5 ms.

The REQ and GNT messages also take time to process and transport between the CMTS scheduler and the CMs. In this perspective, the REQ-GNT delay is the sum of the REQ propagation time over the coax plant and CIN (in the case of a centralized scheduler), REQ queueing time at the CMTS, CMTS MAP processing time, MAP propagation time over the CIN (in the case of a centralized scheduler) and coax plant, CM MAP processing, as well as the necessary US and DS PHY serialization / framing time. If the total REQ-GNT processing and transport delay, which includes the CIN delay, is less than or comparable to the REQ-GNT protocol delay, R-PHY centralized US scheduling will have no impact on US latency.

Figure 3 visualizes the composition of the REQ-GNT delay elements, and the impact of the MAP interval and the HFC/CIN distance. The latency values of the corresponding delay elements are listed in Table 1, calculated at different MAP intervals, 2 ms vs. 1 ms, and different HFC/CIN distances, five miles (short) vs. 100 miles (maximum I-CCAP operation range). Note that since the coax delay is negligible compared to the CIN delay in R-PHY, the CIN distance and the HFC distance are used interchangeably in this paper.

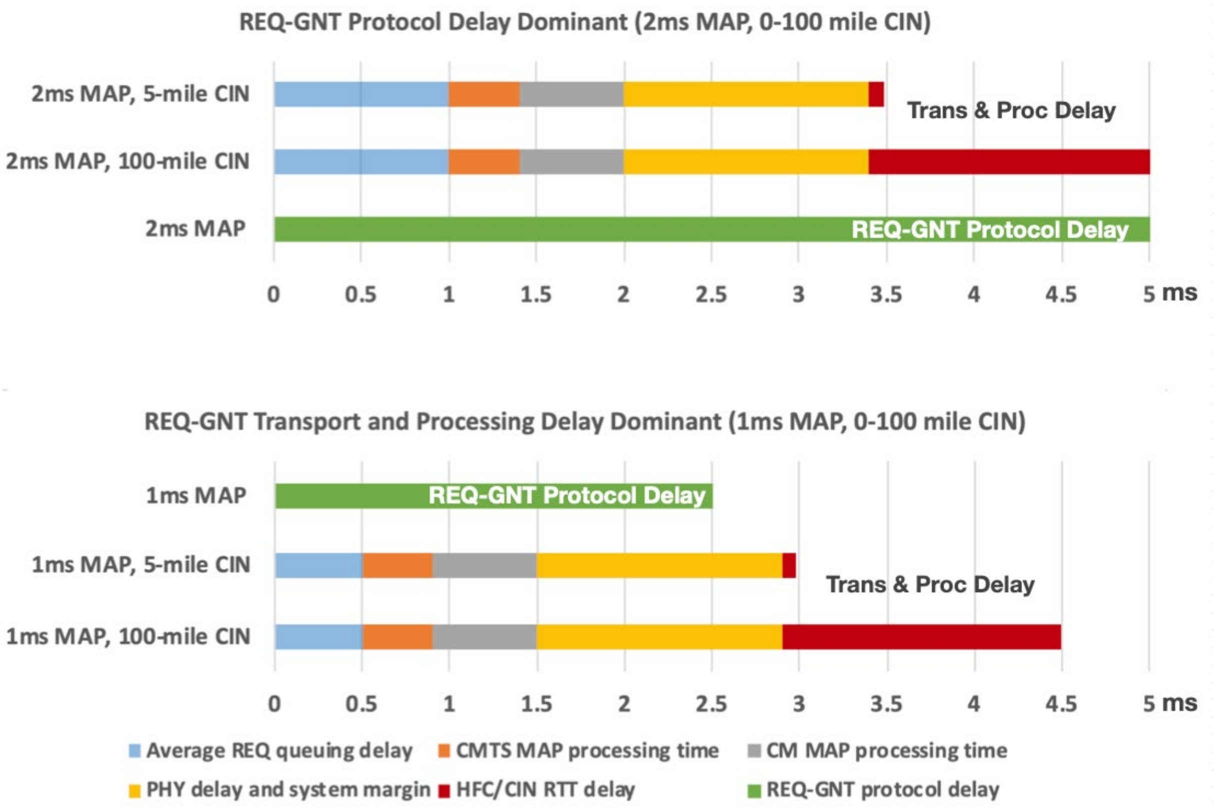

Figure 3 - REQ-GNT Delay Elements and the Impact of the MAP Interval and CIN Distance

As shown, the impact of the MAP interval is mainly reflected in the DOCSIS REQ-GNT protocol delay, which is 2.5 times of the MAP interval as explained earlier. The CIN distance, on the other hand, affects the round-trip CIN delay for transporting the REQ and MAP (GNT) messages. The maxima between the REG-GNT protocol delay and REQ-GNT transport-processing delay determines the minimum REQ-GNT delay. At a 2 ms MAP interval, the REQ-GNT protocol delay is the dominant factor within the I-CCAP maximum operation range of 100 miles. In this case, R-PHY with centralized scheduling is equivalent to I-CCAP, and there is no latency benefit for moving the scheduler to the RPD.

Table 1 - REQ-GNT Delay Elements and the Impact of the MAP Interval and CIN Distance

LATENCY (ms)		TIME VARIANCE			
		2 ms MAP Interval		1 ms MAP Interval	
		5-mile HFC/CIN	100-mile HFC/CIN	5-mile HFC/CIN	100-mile HFC/CIN
REQ-GNT Transport & Processing Delay (Dt)	REQ-GNT (MAP) CIN delay round trip	0.08	1.6	0.08	1.6
	Average REQ queuing delay	1	1	0.5	0.5
	CMTS MAP processing time	0.4	0.4	0.4	0.4
	CM MAP processing time	0.6	0.6	0.6	0.6
	PHY delays and systrm margin	1.4	1.4	1.4	1.4
	Total	3.48	5	2.98	4.5
REQ-GNT Protocol Delay (Dp)	2* MAP Interval + Average REQ queueing delay	5	5	2.5	2.5
MIN REQ-GNT Delay	max (Dt, Dp)	5	5	2.98	4.5

1.2. The Case for a Remote US Scheduler

As cable network evolves, new use cases for remote US scheduling start to emerge, primarily driven by the need for long-distance R-PHY deployment and low-latency support over DOCSIS.

1.2.1. REQ-GNT Protocol Tightening

As part of the low-latency initiatives, the DOCSIS REQ-GNT protocol is being tightened with shorter MAPs (1 ms) and shorter CMTS MAP processing time (400 µs) in order to support low latency over DOCSIS [3]. At the 1 ms MAP interval, the REQ-GNT protocol delay is cut by half, which makes the REQ-GNT processing and transport delay a likely dominating factor of the minimum REQ-GNT delay, as in the example shown in Figure 3. Removing CIN delay in this case does improve the latency performance especially for the longer CIN distance case.

It's also interesting to note from Figure 3 that for long-distance R-PHY deployment, reducing the MAP interval has little impact on the REQ-GNT delay unless the scheduling location is at the RPD. On the other hand, for a short-distance R-PHY deployment, the scheduler location does not matter much if the REQ-GNT processing delay is comparably large with respect to the CIN round trip delay.

Figure 4 illustrates the impact of the CIN delay on the REQ-GNT time when the MAP interval is comparable to the CIN delay. For the REQ sent during the MAP1 interval, the GNT issued by the remote US scheduler at the RPD is available to the CM in MAP3. In comparison, the GNT issued by the core US scheduler for the same REQ is available to the CM at MAP4, as the REQ missed MAP3 building time at the core US scheduler due to the CIN delay.

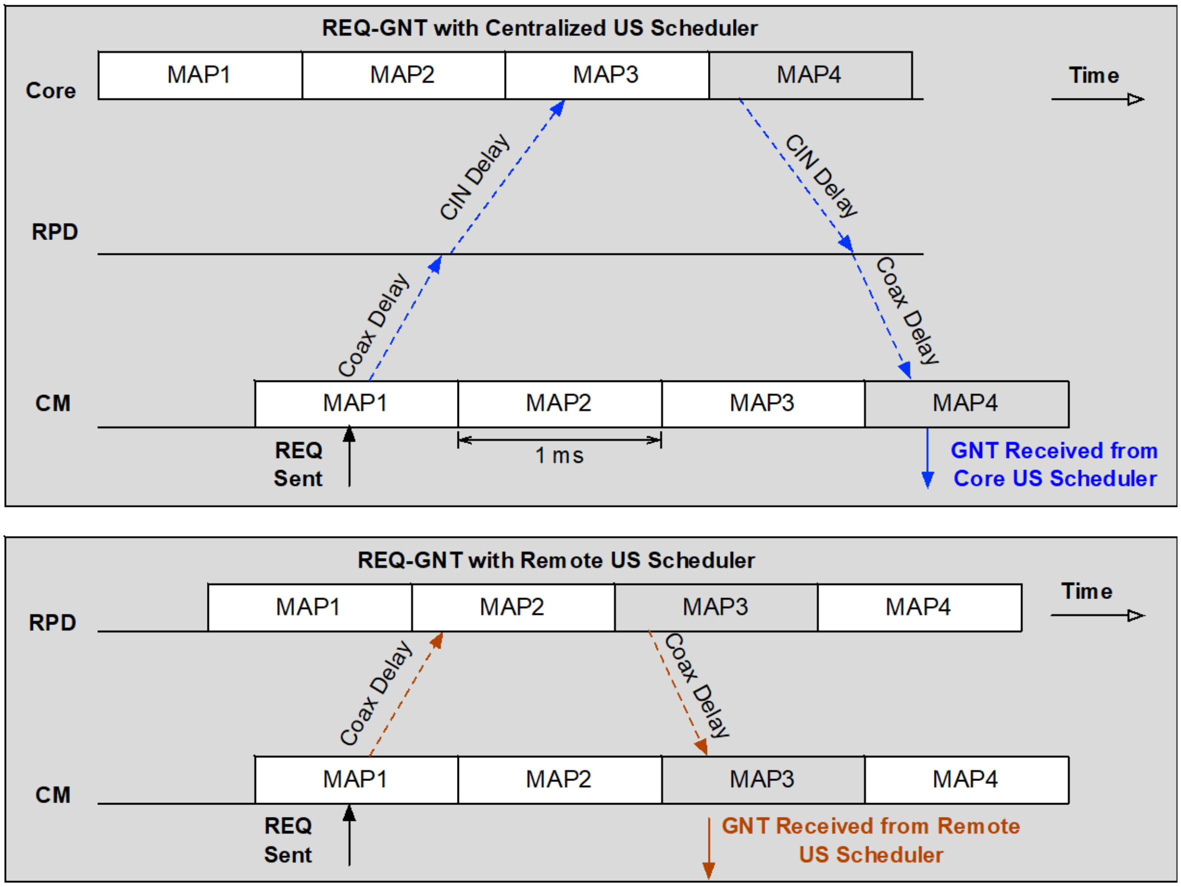

Figure 4 - CIN Delay Impact on the REQ-GNT Latency at Short MAP Intervals.

1.2.2. Long Distance R-PHY Deployment

In a long-distance R-PHY deployment, the CIN distance is beyond the 100-mile range, making CIN delay a dominant factor in the REQ-GNT latency. Long-distance R-PHY deployment may be needed for hub site consolidation where the CCAP core is moved to a central headend or a regional data center. In addition to latency, longer CIN distance may also result in larger jitter due to the queue buildups in the network. Using centralized scheduling for the REQ-GNT handling obviously has challenges to meet the latency and jitter requirements for timing-sensitive applications such as voice and gaming.

On the other hand, adding a remote US scheduler at the RPD can effectively decouple the CIN from the REQ-GNT loop, offering a significant latency improvement for a long-distance R-PHY deployment.

1.2.3. CCAP Realtime Performance Acceleration

DOCSIS US scheduling, characterized by its stringent timeliness, is a real-time task that presents design challenges in the virtualized or cloud CCAP core environment, as the software components with different levels of criticality are all running on the same compute platform. In other words, proper separation and isolation must be in place for the US scheduling task to meet the MAP timing requirements.

With the recent advance of DOCSIS 4.0 technologies, including FDX DOCSIS and extended spectrum, there is a pressing need to scale up the US scheduling capacity to match the upstream bandwidth capacity, which can be up to 50x more than what legacy DOCSIS can offer.

Distributing the timing-sensitive US scheduling tasks to the RPDs can effectively isolate the real-time processing and prevent the disruptions from the non-real time tasks running at the core. In this perspective, the RPD remote US scheduler essentially accelerates the CCAP real-time performance and achieves the scaling required for supporting DOCSIS 4.0 with a distributed scheduling model.

1.2.4. Easy to Add and Flexible to Adapt

From an R-PHY architecture point of view, the remote US scheduler is a component internal to the RPD, and can therefore leverage the RPD hardware and software platform, forwarding, control and management interfaces, and the common configuration and performance monitoring infrastructure.

The R-UEPI and R-DEPI architecture established today can adapt to the remote US scheduler by changing the endpoint location of certain scheduler related pseudowires (PWs) on the MAC side. In other words, no change is needed on the US or DS PHY silicon.

US scheduling with the remote US schedulers is not a decentralized scheduling approach. Instead, it is a distributed scheduling scheme with centralized control performed by a core US scheduler, such that scheduling tasks can be load balanced vertically between the core US scheduler and the remote US scheduler at per-US service flow and / or per-channel basis. This architecture permits latency and efficiency optimizations as well as backward combability with legacy RPDs.

The RPD remote US scheduler is therefore a light-weight, promising solution to address R-PHY latency concerns and accelerate real-time performance for the virtual or cloud CCAP core environment.

2. Remote US Scheduling Services

In a R-PHY system that has remote scheduling, there are actually two schedulers. There is the core scheduler that lives in the CCAP core, and the remote scheduler that lives in the RPD. Together, they form a client-server relationship and collectively provide upstream scheduling services to other MAC elements and applications as shown in **Figure** 5.

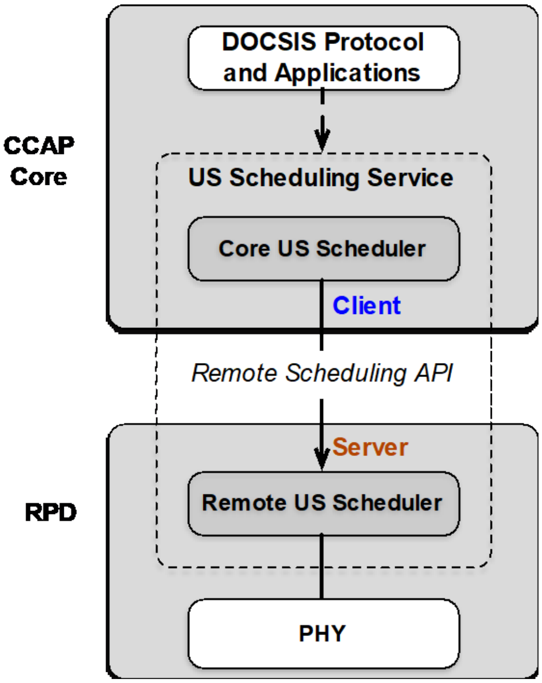

Figure 5 - R-PHY Remote Scheduling Model

The core US scheduler oversees the scheduling needs and manages remote scheduling services for latency sensitive scheduling tasks. It optimizes latency performance for the system as a whole by controlling which scheduling tasks to run remotely and when to trigger them. For the remaining MAC layer and upper layer applications, the core US scheduler provides a top-layer scheduling service, for example enabling an activated service flow, arranging US transmission opportunities for ranging and probing, or coordinating with profile management and proactive management applications with special grant allocations in MAPs.

Since US scheduling is per-US, it is technically possible to have some channels running with a remote scheduler and other channels running with a centralized scheduler. For example, legacy DOCSIS 3.0 ATDMA channels could remain with a fully centralized scheduler and the DOCSIS 3.1 OFDMA channels could run with a distributed scheduler.

The northbound interface between the core US scheduler and the upper MAC and applications is internal to the core, therefore there is no need for standardization. The southbound interface facing the remote US schedulers is a candidate to be standardized via a set of remote US scheduling APIs. This will allow the CCAP core to interoperate with the remote US schedulers from different RPD vendors.

On the RPD side, the remote US scheduler implements the services declared by the remote US scheduling APIs. During the run time, the remote US scheduler decodes the incoming scheduling requests from the core US scheduler, executes scheduling actions and encodes responses.

The remote US scheduler will provide the following services, also shown in **Figure** 6, namely:

- **Remote REQ-GNT Service**: generates GNTs at the RPD where the REQs are received.
- **Remote MAP Builder Service**: encodes the scheduling decisions into per-US channel DOCSIS MAP messages.

- **Remote MAP UCD Replication**: replicates MAPs and UCDs and transmits them in order in DEPI and UEPI formats. For DEPI MAP and UCDs, the MAP and UCD replicas will generated for each downstream channel designated to carry MAPs and UCDs for a given US channel.

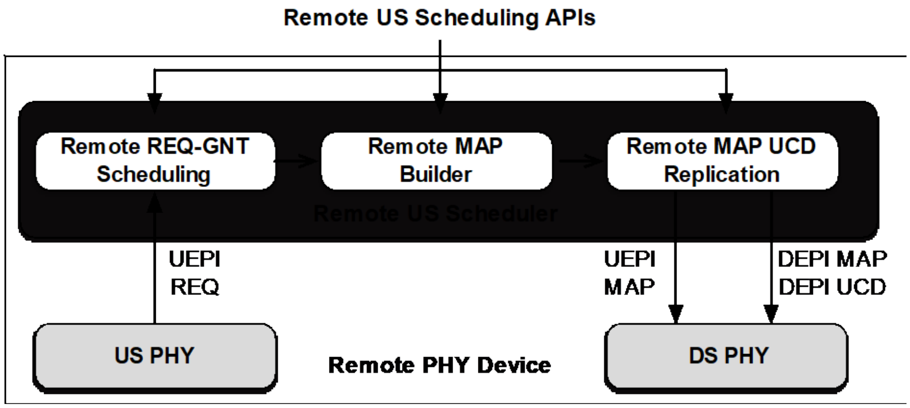

Figure 6 - Remote US Scheduling Service Model

2.1. Remote REQ-GNT Service

The REQ-GNT service implements the reactive scheduling process where grants are given in response to the REQs received from the CMs. When REQ-GNT service is offered remotely at the RPD, it effectively cuts out the CIN round trip delay from the REQ-GNT transport time.

The scope of the REQ-GNT service can be defined on a per-US service flow basis. Once the DOCSIS control plane activates an US service flow, the core US scheduler can request the remote REQ-GNT service by calling the remote REQ-GNT service API. Through the API, the core US scheduler can pass in all the necessary parameters including active QoS parameters, US channel set, and SID to channel assignment associated with the US service flow.

An US service flow without the remote REQ-GNT enabled will remain served by the core US scheduler. With this split-scheduling arrangement, the CCAP core can adjust the remote scheduling workload based on service flows latency requirements and the RPD’s capability and capacity in supporting the remote scheduling.

Depending on how the REQ-GNT scheduling tasks are partitioned, there are two REQ-GNT operational modes, overlapping and non-overlapping, as shown in **Figure** 7. In the overlapping mode, the US channel set for the core REQ-GNT service may overlap with the US channel set used by remote REQ-GNT service. In the non-overlapping mode, there is no channel overlapping between the two scheduling entities.

The non-overlapping REQ-GNT scheme is the basic operation mode which is sufficient to separate out the latency sensitive service flows vs. the latency tolerant service flows between the remote US scheduler and the core US scheduler. The overlapping REQ-GNT scheme is an advanced mode, which offers a bandwidth efficiency benefit by allowing the spectrum resource to be shared between the two scheduling entities. The remote scheduling APIs will be structured to allow the flexibility for supporting both REQ-GNT operation modes.

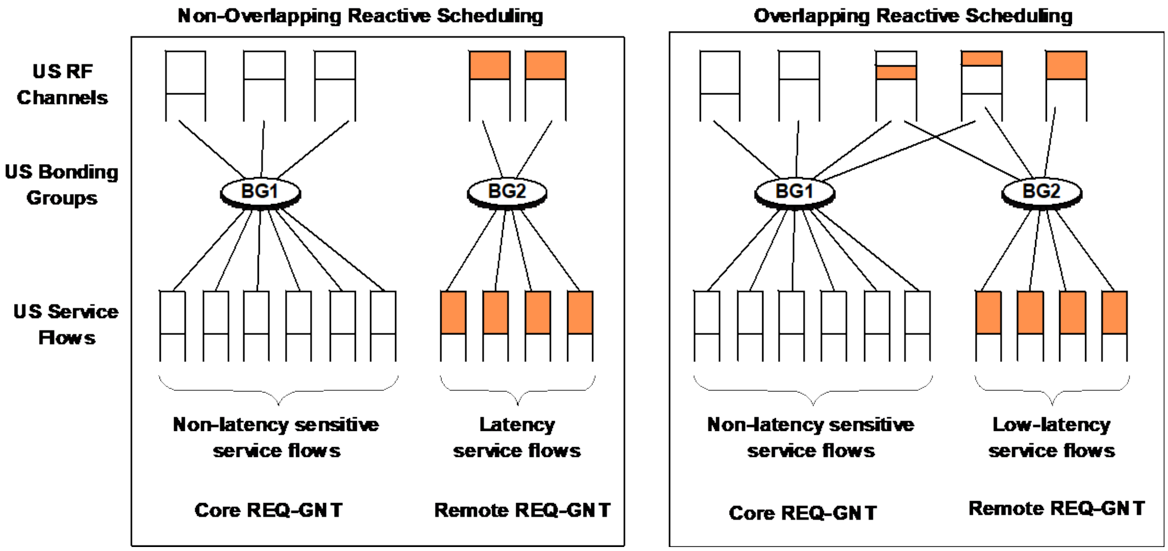

Figure 7 - REQ-GNT Operation Modes

2.2. Remote MAP Builder Service

The remote MAP builder provides the DOCSIS MAP encoding service to translate a scheduling decision into a MAP slot that represents an US transmission opportunity at specific time and frequency. It ensures the nominal MAP message interval and the MAP encoding consistency with the PHY layer configurations carried in the UCD messages.

The remote MAP builder serves the core US scheduler for the portion of the scheduling that is not handled by the remote REQ-GNT service, including the allocations of the ranging request opportunities, probing, OUDP burst transmission opportunities for DOCSIS 3.1 profile test, unsolicited grants for UGS, RTPS and PGS, or proactive data grants based on traffic predictions.

The remote MAP builder service also allows the client to request a MAP slot at a projected time with specifications for periodicity and jitter tolerance. This enables certain core applications like US symbol capture and sounding in FDX, that are latency tolerant however need to be triggered at a specific time in MAP.

Within the remote scheduler, the remote MAP builder is the next node in the service chain after the remote REQ-GNT process. It can therefore be used to serve the REQ-GNT scheduling module internally by converting a grant in bytes to a grant in minislots and encoding it as a MAP IE.

The remote MAP builder service can be enabled on a per-US channel basis. Given the remote REQ-GNT service is enabled on a per-US service flow basis, the remote MAP builder service needs to be enabled on all US channels in the bonding group associated with the US service flow, which in turn requires the remote MAP and UCD replication service described below.

2.3. Remote MAP and UCD Replication Service

After a DOCSIS MAP is built for an US channel, it typically needs to be replicated, as multiple DS channels may be assigned to carry the MAPs for the US channel. Similarly, the DOCSIS UCD message needs to be replicated for each of the MAP carrying DS channels, and sent in sequence with the MAPs to ensure MAC and PHY consistency upon a UCD change.

For any US channel enabled for remote MAP builder service, the MAPs and UCDs must be replicated remotely as the next step. However, remote MAP and UCD replication can be enabled as an independent service by itself even in the centralized scheduling case, which will help offload the CCAP processing if the replication is performed in software, and reduce network traffic load, especially as MAPs and UCDs need to be high priority across the CIN.

To facilitate remote MAP and UCD replication, the DS packet scheduler at the CCAP core must take into consideration the bandwidth consumed by the MAPs and UCDs on the replicated DS channels when shaping the DS traffic flows.

3. R-UEPI in Remote US Scheduling

The Remote Upstream External Physical Interface (R-UEPI) [5] consists of a set of L2TPv3 PWs connecting the US MAC and PHY in between the CCAP core and the RPD. The information exchanged between them includes various DOCSIS US bursts in the RPD to core direction, and DOCSIS MAPs in the core to RPD direction. In the R-UEPI architecture today, the centralized US scheduler at the core is a common end point for four types of UEPI PWs listed below:

MAP PW	From CCAP core to the RPD, containing the DOCSIS MAP messages. There is one MAP PW session for each US channel.
REQ PW	From RPD to the CCAP core, containing DOCSIS REQ information extracted from the REQ burst or the piggybacked data burst.
RNG-REQ PW	From RPD to the CCAP core, containing DOCSIS ranging-request message and the US PHY metrics measured from the ranging burst sent.
Probe PW	From RPD to the CCAP core, containing the US PHY metrics measured from the probes.

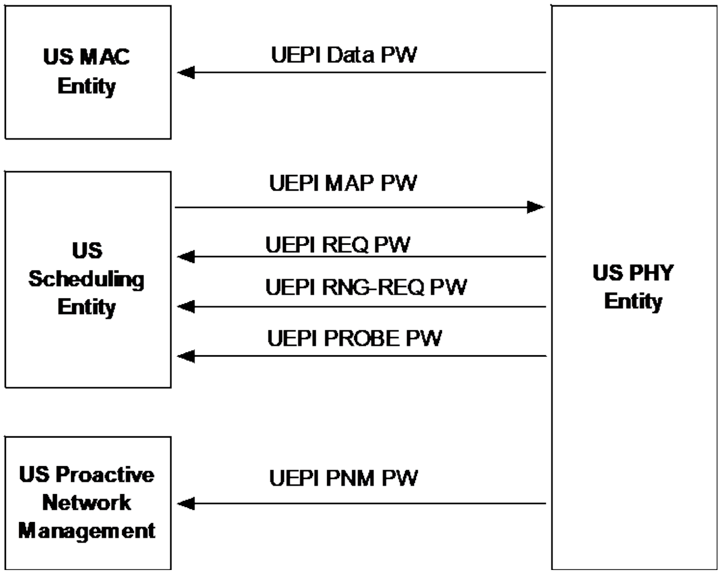

Figure 8 - UEPI Architecture with Centralized US Scheduling

As part of the scheduling function moved to the RPD, the end points of the relevant UEPI PWs will naturally split between the core US scheduler and the remote US scheduler depending on the work load partition between the two scheduling entities.

Figure 9 shows an R-UEPI arrangement option in a remote US scheduling scenario. The UEPI MAP PW and the REQ PW are shortened and terminated at the remote US scheduler co-located with the US R-PHY, the RNG-REQ PW and probe PW remain terminated at the core US scheduler. This arrangement is based on the required MAC processing and response time per DOCSIS protocol. For ranging and probing, the response time limit is 200 ms, well beyond the time scale of the CIN delay. The REQ-GNT time on the other hand needs to be as short as possible, in the 1~2 ms range, for supporting low-latency data services.

If a UEPI PW is terminated locally at the RPD, the UEPI control plane does not need to use L2TPv3 signaling. Instead, the local CPU can do the UEPI configuration through direct register access, same as in the embedded UEPI architecture.

From the US PHY point of view, there is no difference between an embedded PW or an external PW, as the UEPI framing and the forwarding plane setup will remain the same, therefore there is no impact to the US PHY silicon.

In summary, the native UEPI architecture can readily support the R-PHY architecture with remote US scheduling. It has the flexibility required for load balancing the US scheduling tasks to maximize the latency benefit of the remote scheduler and the core computation capacity for non-real time scheduling services.

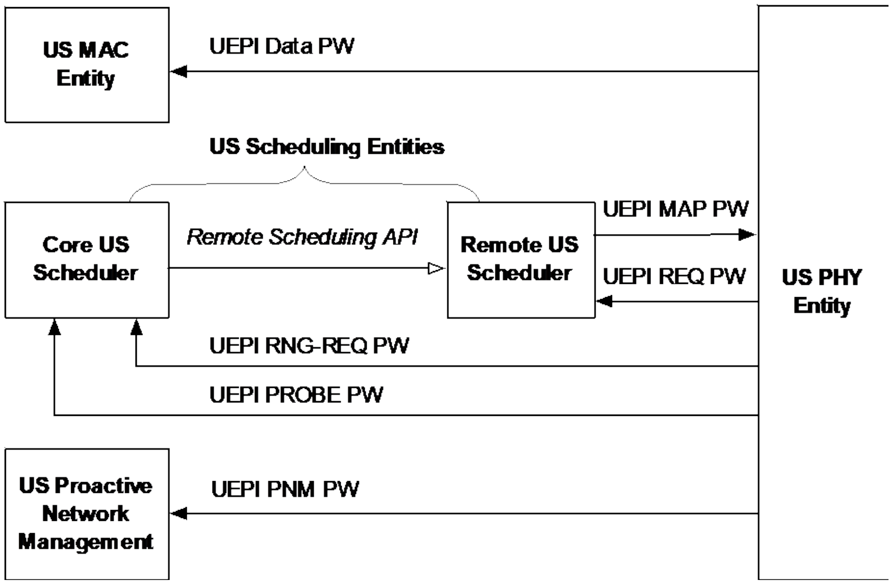

Figure 9 - UEPI Architecture with Remote US Scheduling

4. R-DEPI in Remote US Scheduling

The Remote Downstream External Physical Interface (R-DEPI) [6] is used to carry the DS DOCSIS data and signaling from the CCAP core to the RPD. It uses L2TPv3 PWs and separate packet stream protocol (PSP) flows for data traffic and MAP/UCDs.

Figure 10 shows a typical R-DEPI connection setup with centralized US scheduling. The CCAP core and the RPD form a pair of L2TP control connection endpoints that have a DEPI PW session for each channel. Each session has two PSP flows, PSP flow1 for DOCSIS MAPs and UCDs and PSP flow2 for carrying regular data traffic. PSP flow1 is typically encoded with the DSCP expedited forwarding code point for the MAPs and UCDs to be delivered at high priority across the CIN.

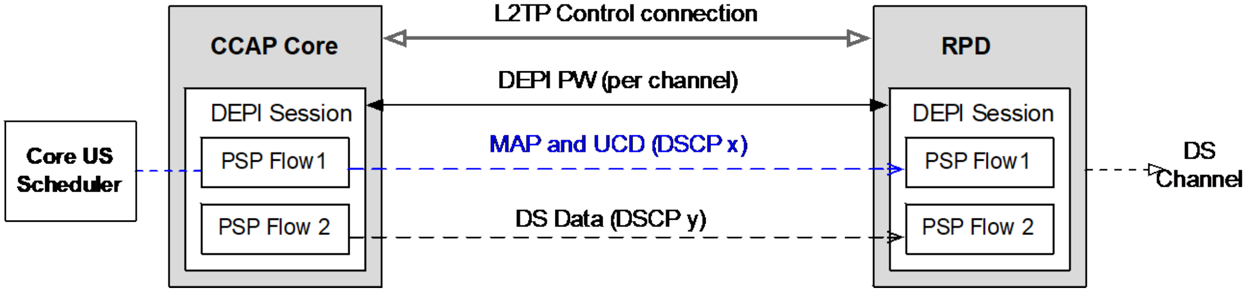

Figure 10 - R-PHY DEPI connection with Centralized US Scheduling

Since the sequence of the DEPI PSP segments is maintained per PSP flow, PSP flow1 can have a completely different forwarding path from the rest of the channel. This flexibility allows the DEPI MAPs and UCDs to be injected from a location totally different from the DS data PSP flow. **Figure** 11 shows the DEPI connections in the remote US scheduling case, where DEPI MAPs and DEPI UCDs are inserted locally at the RPD. From the DS PHY point of view, there is no difference between the centralized US scheduling case and the remote US scheduling case, therefore there is no silicon impact to the DS PHY.

In summary, the remote US scheduling scheme can be readily supported in the R-DEPI architecture.

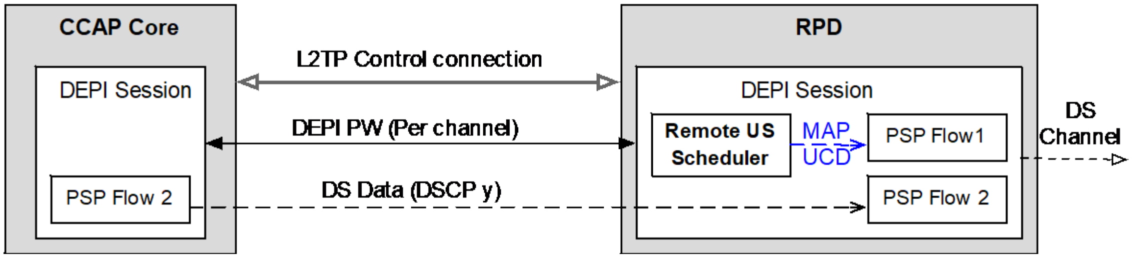

Figure 11 - R-DEPI Connectivity with Remote US Scheduling

5. Remote US Scheduling APIs

The scope of work for defining the remote US scheduler APIs falls into two categories, configuration and operational data modeling which determines the content of the API, and the mechanism to express and convey the APIs between two network entities including transport, encoding and protocol design/selections. **Figure** 12 shows the API stack in today's R-PHY, R-PHY 1.0, and the new options for R-PHY 2.0 [4].

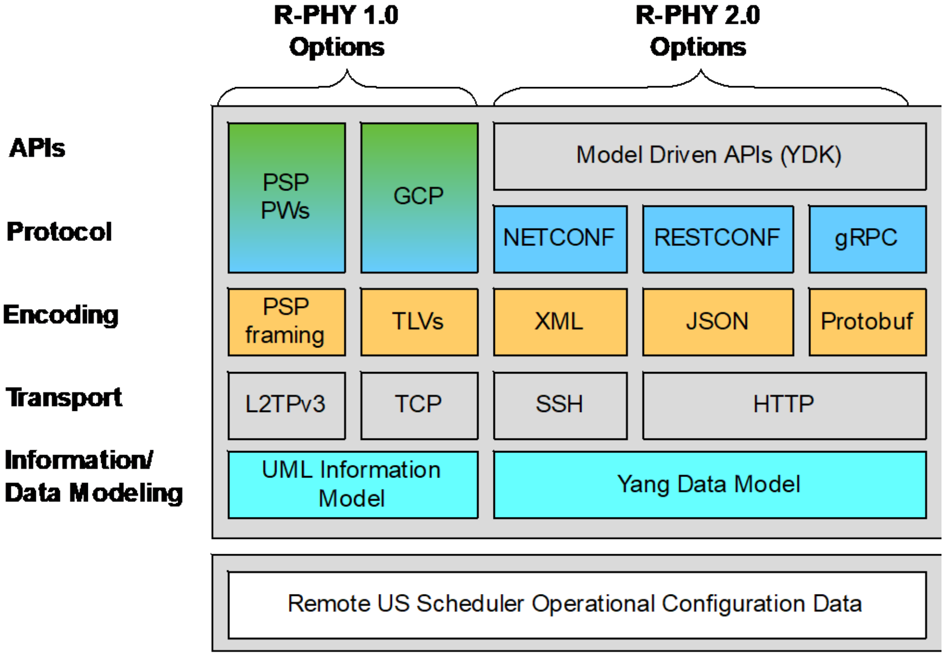

Figure 12 - Remote US Scheduling Service API - Scope and Choices

R-PHY 1.0 uses the UML information model and a messaging scheme that is built on DEPI/UEPI framing and GCP TLVs. APIs are described as rules and requirements declared in the R-PHY specifications published by CableLabs. This process is expected to be overhauled in R-PHY 2.0, where cloud-friendly APIs can be directly enforced based on the Yang data models. Tools like the Yang Development Kit (YDK) can be used to compile the Yang data model to provide APIs in several programming languages. These APIs provide precise definitions of the service contract, automatic data validations, and abstractions of the protocol, encoding and transport details. The data model-based API specifications are expected to significantly simplify the feature development by avoiding the conventional pitfalls caused by different interpretations of the written requirements.

This section focuses on the data modeling portion of the APIs, in other words, what information needs to be exchanged between the CCAP core and the remote US scheduler for configuration management, service enabling and operational state collections. The API messaging mechanism will leverage the R-PHY 1.0 or R-PHY 2.0 messaging methodologies.

5.1. Top-Level RPD Operational Configuration Objects

The remote US scheduler module will be a new managed object visible from the top under the RPD operational configuration root module, as shown **Figure** 13 in conventional UML format to show the relationship with other R-PHY top-level objects as specified in [7].

An RPD that is capable of remote US scheduling contains one *remote-us-scheduler* module, which holds the configuration and operational data objects accessible through the remote US scheduling APIs.

The *remote-us-scheduler* is associated with the existing RPD DS and US objects to access the RF channel configurations through reference needed for MAP building and MAP/UCD replications.

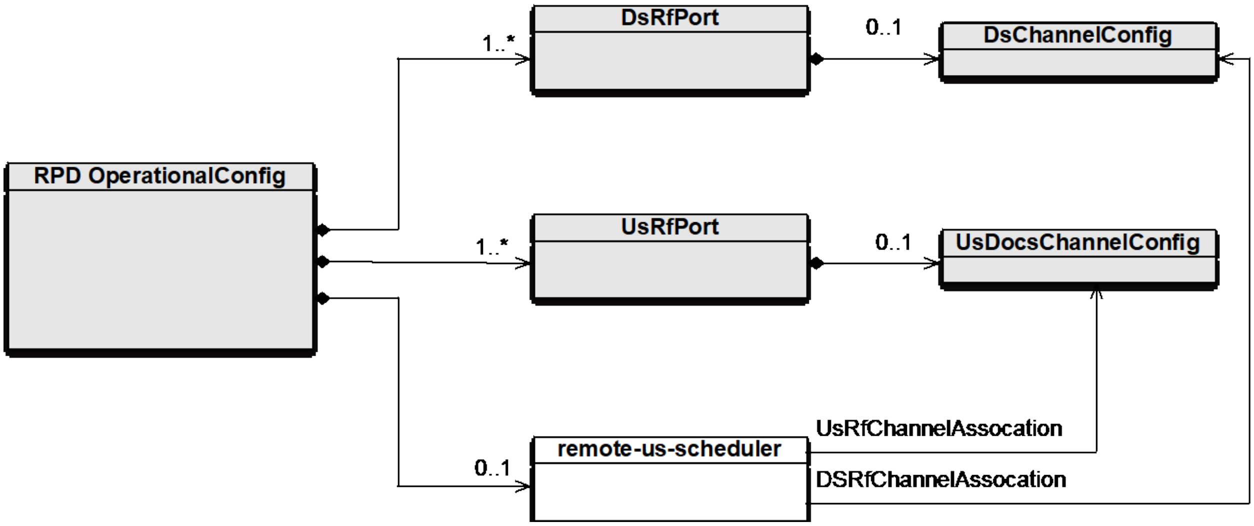

Figure 13 - Remote US Scheduler Module in Relation with Top Level RPD Configuration Objects

5.2. Remote Upstream Scheduler Data Object Tree

The *remote-us-scheduler* module includes all the configuration and operational objects for performing remote US scheduling at the RPD. **Figure** 14 is a simplified Yang data model diagram showing the main branches of the *remote-us-scheduler* data tree. As shown, the *remote-us-scheduler* contains a list of per MAC domain US scheduling instances. Each MAC domain US scheduling instance collects the configuration and operational data elements for all the remote US scheduling service categories.

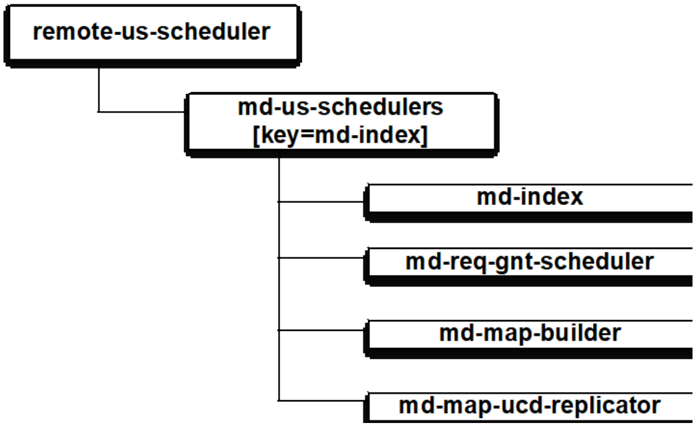

Figure 14 - Remote Upstream Scheduler Data Model

The main elements included in the *md-us-scheduler* are listed below:

md-index	Contains the MAC domain index that identifies a MAC domain from the list
md-req-gnt-scheduler	Contains the information elements for managing the remote US REQ-GNT services for the US service flows in the given MAC domain.

md-map-builder	Contains the information elements for managing the remote MAP-builder services for the US channels in the given MAC domain.
md-map-ucd-replicator	Contains the information elements for managing the remote MAP and UCD replication services for the given MAC domain.

5.2.1. MAC Domain REQ-GNT Scheduler

The MAC domain REQ-GNT scheduler, *md-req-gnt-scheduler*, is the container that organizes the configuration and operational data for a list of US service flows enabled for the remote REQ-GNT scheduling service, as shown in **Figure** 15. In this diagram, the configuration data elements (read-write) are shown as white boxes, and the operational data elements (read-only) are shown as shaded boxes.

Each US service flow element contains the following items:

sf-index	Contains the service flow index that identifies a service flow from the list
sid-clusters	Contains the SID cluster configurations that associate the SIDs carried in the REQ with proper US channel resources
active-qos-prams	Contains the service flow active QoS parameter configurations, for example traffic priority, maximum sustained rate, etc.
enable-disable	Contains the action to enable or disable the remote REQ-GNT service for the given service flow
req-gnt-statistics	Contains the operational data of the REQ-GNT statistics, for example, the total number of bytes requested and the number of bytes granted on the given service flow

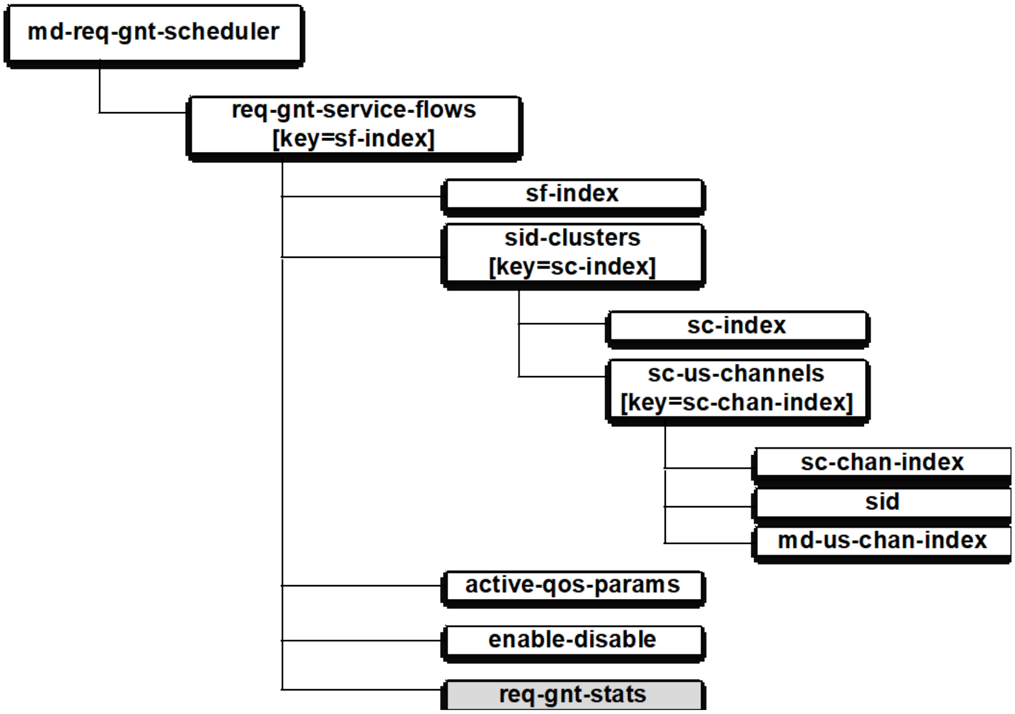

Figure 15 - MAC Domain REQ-GNT Scheduler Data Model

5.2.2. MAC Domain MAP Builder

The MAC domain MAP builder, *md-map-builder*, is the container that organizes the configuration and operational data for a list of US channels intended for the remote MAP-builder service, as shown in **Figure** 16.

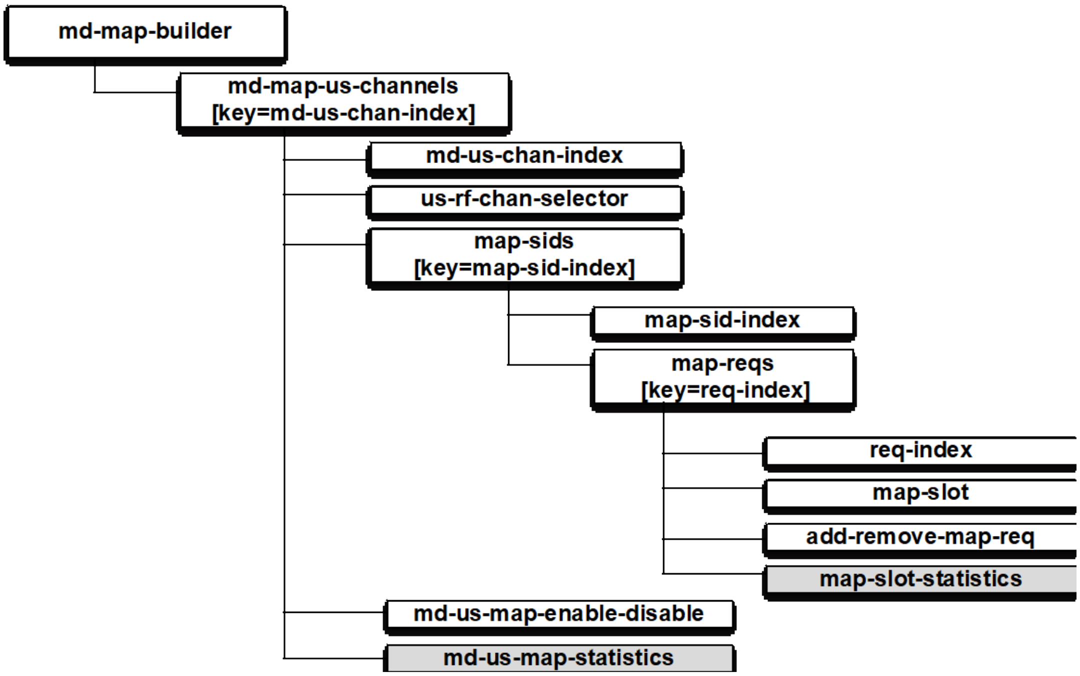

Figure 16 - MAC Domain MAP Builder Data Model

Each MAC domain US channel element contains the following items:

md-us-chan-index	Contains the US DOCSIS channel index that identifies an US channel from the per-MAC domain channel list.
us-rf-chan-selector	Contains the RF channel identifier, as defined in [6], associated with the MAC domain US channel.
map-sids	Contains a list of MAP-SID elements that require transmission opportunities in MAPs. Each MAP-SID element contains a list of requests for transmission opportunities, *map-slots*.
md-us-chan-enable-disable	Contains the action to enable or disable the MAP-builder service on the given US channel.
md-us-map-statistics	Contains the MAP builder statistics, for example, the percentage of minislots allocated for each map-slot type.

The MAP slot request, *map-slot-req*, contains a scheduling decision that needs to be translated into a slot in MAP, or a transmission opportunity from the CM's point of view. The CCAP core uses this API to schedule transmission opportunities that are not handled by the remote US scheduler, such as ranging, probing, profile testing, spectrum capture, unsolicited or proactive data granting, etc. The MAP slot request can be either one-shot for certain ad hoc scheduling needs, or periodic at regular intervals. **Figure** 17 shows the types of MAP slots that the CCAP core can request for MAP-build services.

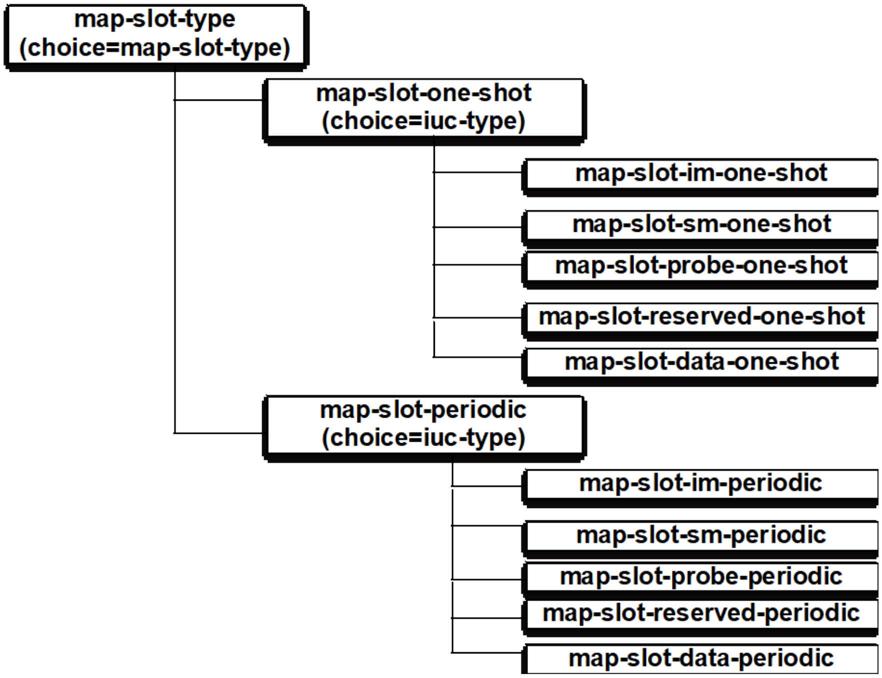

Figure 17 - MAP Slot Base Type and Derived Types

5.2.3. MAC Domain MAP-UCD Replicator

The MAC domain MAP-UCD replicator, *md-map-ucd-replicator*, contains the US channel to downstream channel binding information for the MAP and UCD replication function, as show in **Figure** 18.

Each MAC US channel element contains the following items:

md-us-chan-index	Contains the US DOCSIS channel index that identifies an US channel from the per-MAC domain channel list.
map-ucd-ds-channels	Contains the list of DS channels assigned to carry the MAPs and UCDs for the associated US channel. Each DS channel element contains the association to the DS RF channel identifier, and the MAP UCD transmission statistics.
map-ucd-replication-enable-disable	Contains the action to enable or disable the MAP and UCD replication service for the given US channel.

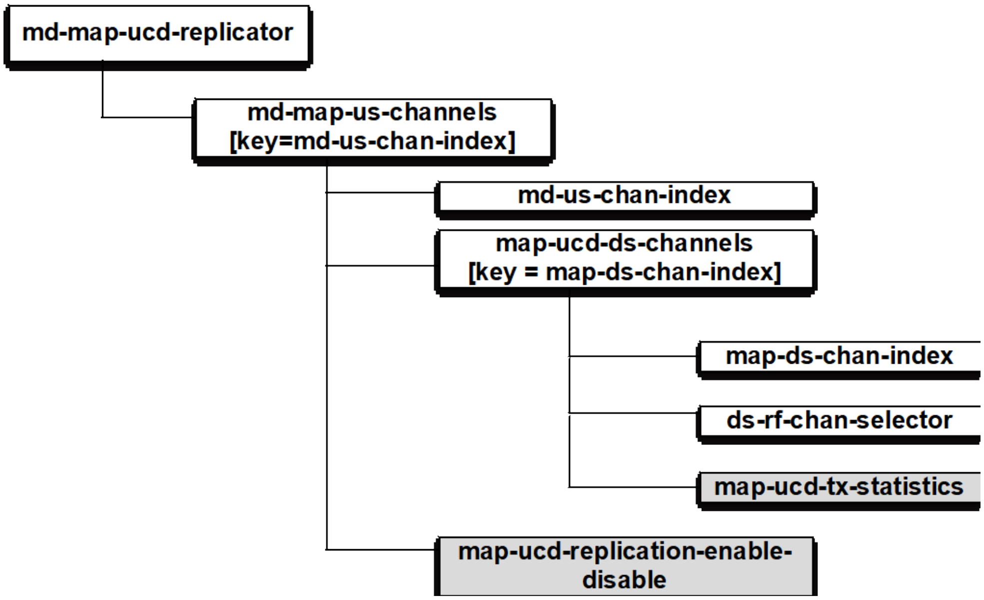

Figure 18 - MAC Domain MAP-UCD Replication Data Model

6. Locality Optimization with Remote US Scheduling

The essence of split US scheduling with the remote US scheduler at the RPD is to enable system-wide locality optimization. By co-locating the reactive granting portion of the scheduling with the PHY where the REQ is received, latency is improved by avoiding the CIN delay in the REQ-GNT process. By keeping the latency-tolerant and computation-intensive part of the scheduling centralized at the core, efficiency is maximized by positioning the core as the common computation platform accessible to all RPDs, as shown in **Figure** 19.

This scheduling model is different from the decentralized scheduling model as used in FMA that has no centralized control from the CCAP core. The core US scheduler in the R-PHY case provides a unique value for global locality optimization that takes into consideration the per-service flow latency requirement, RPD capabilities / constraints, and the CCAP core real-time processing capacity. In this perspective, the ability to do remote scheduling with centralized control gives R-PHY the architecture advantages to achieve low-latency and high efficiency, and remain backwards compatible with legacy R-PHY deployments.

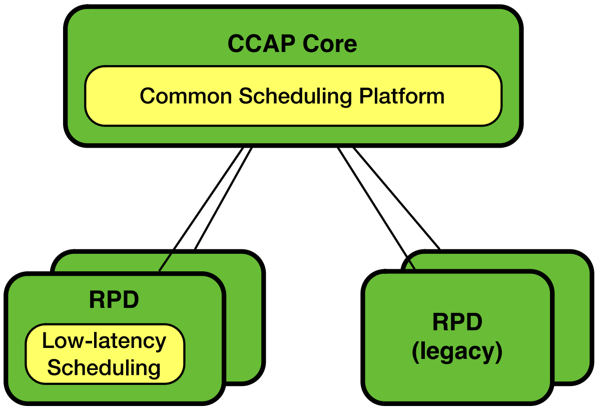

Figure 19 - Upstream Scheduling Locality Optimization

Conclusion

Cable networks have been going through a radical transformation, changing from bandwidth-limited and latency-tolerant networks to a high-capacity, low-latency, multi-service edge access network. Adapting to the change by enabling low-latency US scheduling in R-PHY is a step taken to accelerate this transformation and prepare cable networks for the future.

R-PHY low-latency US scheduling involves moving the latency-sensitive scheduling tasks such as REQ-GNT handling to the RPD, while keeping the latency-tolerant scheduling tasks centralized to retain the centralized MAC advantages. For the US scheduling service, the core and RPD form a client-server relationship, where the RPD remote US scheduler provides services to REQ-GNT low-latency service flows, builds MAPs for both core and the remote schedulers, and replicates MAP UCDs to the proper DS channels. Such services can be precisely defined using data model-based APIs, which can be autogenerated based on published Yang data models.

The RPD remote US scheduler can be built on top of an existing R-RPY platform, which contains the basic MAC and PHY building blocks and the glue logic. The addition of the remote US scheduler has no impact on the US PHY or DS PHY silicon and can be readily supported by the R-UEPI and R-DEPI architecture, as the only change needed is the endpoint location of the UEPI and DEPI PWs on the MAC side.

The addition of the remote US scheduler to the R-PHY US scheduling scheme enables a distributed scheduling model where the core can optimize the scheduling locations and conduct the vertical load balancing between the core and the RPD. This scheduling model is unique to the R-PHY architecture, being able to achieve system-wide optimization in both latency and efficiency, and simultaneously maintaining backwards compatibility with legacy RPDs.

Acknowledgements

The authors would like to thank our Cisco colleagues Brian Bresnahan and Wenkai Zhu for their valuable contributions to the paper.

Abbreviations

CM	cable modem
CCAP	converged cable access platform
CIN	converged interconnect network
DAA	distributed access architecture
DS	downstream
FMA	flexible MAC architecture
GNT	DOCSIS bandwidth grant
LLS	low latency scheduling
REQ	DOCSIS bandwidth request
R-PHY	remote PHY
US	upstream

Bibliography & References

[1] John Chapman, Gerry White, Hang Jin; *Impact of CCAP to CM Distance in a Remote PHY Architecture;* 2015 Spring Technical Forum Proceedings

[2] John Chapman, Jennifer Andreoli-Fang; *Mobile Backhaul over DOCSIS*; 2018 SCTE Technical Forum Proceedings

[3] "CM-SP-MULPIv3.1-I18-190422: MAC and Upper Layer Protocols Interface Specification", CableLabs, 2019

[4] Pawel Sowinski, Andy Smith, Tong Liu etc: *Remote PHY 2.0*; 2019 SCTE Technical Forum Proceedings

[5] CM-SP-R-UEPI-I05-170111: "Remote Upstream External PHY Interface Specification", CableLabs.

[6] CM-SP-R-DEPI-I06-170111: "Remote Downstream External PHY Interface Specification", CableLabs.

[7] CM-SP-R-PHY-I10-180509: "Remote PHY Specification", CableLabs