

802.3bn Ad-hoc Multiple Modulation Profile (MMP)

January 2, 2013

Jorge Salinger, Comcast – Chair

Instructions for the WG Chair

The IEEE-SA strongly recommends that at each WG meeting the chair or a designee:

- Show slides #1 through #4 of this presentation
- Advise the WG attendees that:
 - The IEEE's patent policy is described in Clause 6 of the *IEEE-SA Standards Board Bylaws*;
 - Early identification of patent claims which may be essential for the use of standards under development is strongly encouraged;
 - There may be Essential Patent Claims of which the IEEE is not aware. Additionally, neither the IEEE, the WG, nor the WG chair can ensure the accuracy or completeness of any assurance or whether any such assurance is, in fact, of a Patent Claim that is essential for the use of the standard under development.
- Instruct the WG Secretary to record in the minutes of the relevant WG meeting:
 - That the foregoing information was provided and that slides 1 through 4 (and this slide 0, if applicable) were shown;
 - That the chair or designee provided an opportunity for participants to identify patent claim(s)/patent application claim(s) and/or the holder of patent claim(s)/patent application claim(s) of which the participant is personally aware and that may be essential for the use of that standard
 - Any responses that were given, specifically the patent claim(s)/patent application claim(s) and/or the holder of the patent claim(s)/patent application claim(s) that were identified (if any) and by whom.
- The WG Chair shall ensure that a request is made to any identified holders of potential essential patent claim(s) to complete and submit a Letter of Assurance.
- It is recommended that the WG chair review the guidance in *IEEE-SA Standards Board Operations Manual* 6.3.5 and in FAQs 12 and 12a on inclusion of potential Essential Patent Claims by incorporation or by reference.

Note: WG includes Working Groups, Task Groups, and other standards-developing committees with a PAR approved by the IEEE-SA Standards Board.

(Optional to be shown)

25 March 2008 (updated January 2012)



Participants, Patents, and Duty to Inform

All participants in this meeting have certain obligations under the IEEE-SA Patent Policy.

- Participants [Note: Quoted text excerpted from IEEE-SA Standards Board Bylaws subclause 6.2]:
 - “Shall inform the IEEE (or cause the IEEE to be informed)” of the identity of each “holder of any potential Essential Patent Claims of which they are personally aware” if the claims are owned or controlled by the participant or the entity the participant is from, employed by, or otherwise represents
 - “Personal awareness” means that the participant “is personally aware that the holder may have a potential Essential Patent Claim,” even if the participant is not personally aware of the specific patents or patent claims
 - “Should inform the IEEE (or cause the IEEE to be informed)” of the identity of “any other holders of such potential Essential Patent Claims” (that is, third parties that are not affiliated with the participant, with the participant’s employer, or with anyone else that the participant is from or otherwise represents)
- The above does not apply if the patent claim is already the subject of an Accepted Letter of Assurance that applies to the proposed standard(s) under consideration by this group
- Early identification of holders of potential Essential Patent Claims is strongly encouraged
- No duty to perform a patent search

Patent Related Links

All participants should be familiar with their obligations under the IEEE-SA Policies & Procedures for standards development.

Patent Policy is stated in these sources:

IEEE-SA Standards Boards Bylaws

<http://standards.ieee.org/develop/policies/bylaws/sect6-7.html#6>

IEEE-SA Standards Board Operations Manual

<http://standards.ieee.org/develop/policies/opman/sect6.html#6.3>

Material about the patent policy is available at

<http://standards.ieee.org/about/sasb/patcom/materials.html>

If you have questions, contact the IEEE-SA Standards Board Patent Committee Administrator at patcom@ieee.org or visit <http://standards.ieee.org/about/sasb/patcom/index.html>

This slide set is available at
<https://development.standards.ieee.org/myproject/Public/mytools/mob/slideset.ppt>

Call for Potentially Essential Patents

- If anyone in this meeting is personally aware of the holder of any patent claims that are potentially essential to implementation of the proposed standard(s) under consideration by this group and that are not already the subject of an Accepted Letter of Assurance:
 - Either speak up now or
 - Provide the chair of this group with the identity of the holder(s) of any and all such claims as soon as possible or
 - Cause an LOA to be submitted

Other Guidelines for IEEE WG Meetings

- **All IEEE-SA standards meetings shall be conducted in compliance with all applicable laws, including antitrust and competition laws.**
 - **Don't discuss the interpretation, validity, or essentiality of patents/patent claims.**
 - **Don't discuss specific license rates, terms, or conditions.**
 - Relative costs, including licensing costs of essential patent claims, of different technical approaches may be discussed in standards development meetings.
 - Technical considerations remain primary focus
 - **Don't discuss or engage in the fixing of product prices, allocation of customers, or division of sales markets.**
 - **Don't discuss the status or substance of ongoing or threatened litigation.**
 - **Don't be silent if inappropriate topics are discussed ... do formally object.**

See *IEEE-SA Standards Board Operations Manual*, clause 5.3.10 and "Promoting Competition and Innovation: What You Need to Know about the IEEE Standards Association's Antitrust and Competition Policy" for more details.

Attendance – 04Dec12

- Joe Solomon – Comcast
- Jorge Salinger - Comcast
- Hal Roberts – Calix
- Leo Montreuil – Broadcom
- Saif Rahman – Comcast
- Bill Powell – ALU
- Marek Hajduczenia – ZTE
- Wim De Ketelaere – Excentis
- Rob Howald – Motorola
- Mike Darling – Shaw
- Raanan Ivry – Wide Pass
- Duane Remein – Huawei
- Andrea Garavaglia - Qualcomm
- Steve Shellhammer – Qualcomm
- Satish Mudgere - Intel
- Tom Staniec, Cohere
- Curtis Knittle, CableLabs
- Matt Schmitt, CableLabs
- Paul Nickolich, IEEE
- Peter Wolff, Tital
- Charaf Hanna, ST
- Tim Brophy, Cisco
- Ron Wolf, Aurora
- Kevin Noll,
- Alan Brown

MMP Ad-hoc Logistics

- Conference Calls on Tuesdays at 9:00 – 10:00 AM ET
- We won't have a call on December 25th or January 1st
- Let's move the 1st call in 2013 to Wednesday, January 2nd at the same 11:00 AM – noon ET timeslot (no conflicts that day)
- We'll resume normal scheduled calls on Tuesday January 8th at 9 AM ET

Objective and Goal

- This ad-hoc will be a forum to discuss the merits and draw-backs of MMP for EPoC
- The ad-hoc will try to arrive to a recommendation on whether MMP should be used or not in EPoC for the next EPoC F2F meeting.
- While we may discuss approaches for implementing MMP for EPoC to facilitate the discussion on its merits or draw-backs, it is not the purpose of this ad-hoc to arrive to a recommendation on how MMP would be implemented even if it is deemed appropriate to use it.

Agenda – 18Dec12

- Review IEEE Patent Policy
- Attendance
- Recap:
 - We have gone over multiple presentations including data that shows why MMP would be beneficial: see presentations from Dave Urban titled "Comcast Multiple MCSv3.pdf" and "Comcast MTA data by nodev2.pdf", and slides included in the meeting materials from Rogers and Shaw. Other data provided separately by BrightHouse was considered although not reviewed in detail (see "mallette_01_1112.pdf" from the November F2F meeting in San Antonio) and additional data is expected.
 - We have also heard during this ad-hoc a presentation from Andrea Garavaglia explaining one way in which MMP could be implemented (see "Multiple_Profiles_for_EPoC_14Dec2012.pdf").
 - We can go over more data on the usefulness of MMP and on other ways to implement MMP within EPoC, but let's turn our attention to the complexities introduced by MMP should it be implemented.
 - There have been concerns raised about this within several presentations, especially the ones from Ed Boyd (see "boyd_02_1112.pdf") and from Eugene Dai (see "dai_01a_1112.pdf"), both from the November F2F in San Antonio.

Notes from 12/4/12 meeting

- Ad-hoc will focus on downstream; use of different MCS in the US is already possible given burst mode
- Reviewed slides from Dave Urban
 - Questions regarding node sizes
 - Questions regarding the accuracy of readings from CMs
 - Questions regarding readings showing low and high SNR values
 - Questions regarding one-time snapshot vs. continuous performance
 - Questions regarding network elements contained in HFC plant
 - Questions regarding correlation between US and DS level
- Multiple questions following meeting included in next pages

Questions/Comments from Hal Roberts

1. Impact of Analog Optics: It is often stated that the data we should be analyzing should subtract the impact of the analog optics, since this will not be present in some EPoC scenarios. Observation: The impact of the analog optics is negligible. The optics sets a baseline of about 42dB SNR for digital channels. However the problematic channels have noise that is 10dB worse than the optical baseline. Eliminating the optical noise will only improve the SNR of a 30dB modem by a fraction of a decibel. Therefore elimination of the analog optics will not substantially alter the problem or solution space. The higher SNR channels will, however, improve by elimination of the analog optics.

2. Source of Low SNR Modems: These are generally correlated to modems with low received signal power (RSSI) due to long in-home cable runs or high split ratios. Observation: However this leads to the following mystery when examining the Gateway Data.....

NOTE: Comcast will try to correlate SNR with receive power level

Questions/Comments from Hal Roberts

3. Gateway Data - Impact of Locating Modem at Home Entry Point:

Locating the modem here would seem to solve both low signal problems as well as reduce in-home ingress. Surprisingly, locating the modem here has a relatively negligible affect with only a 2dB shift in average SNR and a ½ dB reduced Standard Deviation. Observation: How do we reconcile #2 and #3?

NOTE: Agreed that the collected data indicates that placing the CM at the home entry point makes little difference on SNR vs. other home placements

4. General Conclusion: The only potential way of eliminating the low SNR outliers is to reduce those modems experiencing a low signal (eliminating the analog optics will not do it). Therefore, locating the CNU at the entrance to the home where the signal level should be close to 0dBmV (and any in-home ingress at lower levels) would seem to be the way to do this. However, current data shows this is not the case. The only possibilities to explain this seem to be: A) Low signal is not the cause of low SNR or B) the 'gateway' location mysteriously still experiences low signal strength even though it is not behind a long cable run or high loss splitters.

Questions/Comments from Marek Hajduczenia

- First, I think that looking at 50+ million modems and drawing conclusions based on such data is a mistake from the get go. Statistically, we will have much smaller service groups than 50 million modems and we should be observing behavior per node, and variations in such behavior, rather than averaging everything. The law of large numbers is there for a reason – in here, we should be looking at trees, and not at the forest.

NOTE: Agreed, but analysis of data from individual service groups shows the same SNR variation

- Second, I find it hard to reconcile data presented on the call yesterday with data shown by Ed from BHN at the last meeting. To me, it seems that either both operators have completely different network designs, use different equipment, or perhaps the measurement methodology is different. While one data set supports clearly multiple profiles, the other one puts that observation into question. Do we then follow the larger data set just because it is larger?

NOTE: It seems that the datasets from BHN and Comcast show the same SNR variations, not different results.

Questions/Comments from Marek Hajduczenia

- Third, in the SNR distribution per node shown on the yesterday's call, it seems that two profiles would cut it for a grand majority of the connected modems. The question then becomes: how distant are these modems from the node itself? Is there any correlation between distance to node and observed SNR or it is rather a complex function of cable, passives and any other sources of noise? Furthermore, are all modems connected to the given node exactly the same? We assume all modems are alike, but I think it is only fair to assume they are not, and under the very same conditions may behave slightly different, especially in terms of measured SNR values (seems that some discussions yesterday support this conclusion).

NOTE: Agreed that 2 profiles would cover the majority of CMs, but would not cover the high and low ends, and that all CMs are not the same (there are many different types)

- Fourth, and perhaps last for now. What is the premium (in relative terms, % wise if you like, taking a single profile equipment as a base value of 1.00) acceptable for equipment supporting multiple profiles versus single profile equipment? There is no free lunch as we all know and making devices more complex entails extra cost due to hardware and management complexity. I am trying to understand where the pain threshold is located and when it becomes simpler for the network operator to go and fix the coax problems rather than trying to address them through the use of super-intelligent and highly-adaptable equipment. Recall that Ethernet is first and foremost about simplicity and robustness, while adaptiveness and tons of options (and knobs to fine tune such options) does not sound very Ethernet-like to me.

Questions/Comments from Eugene Dai

First, some general comments on the statistics. 20+ million data is a very large sample space for cable modems; how to map it into a much smaller sample space such as the number of CNU's connect to one OCU which is in the order of two digits? The nice Gaussian sharp distribution will disappear. This may be the explanation of the difference of the two sets of data from two MSOs.

NOTE: Agreed, SNR distribution for a single node is not a nice Gaussian curve

Second, Many factors affect the SNR distribution; take one for example, the data from 20+ million CMs came from many different HFC plants, and even more, if we break them into cable segments – thousands of them. For a relative “bad” cable segment, all CMs may show lower SNR than a better cable segment. The correlations of the data have to be considered and analyzed.

NOTE: Agreed, grouping of CMs per node seems to show the same distribution, but we should analyze the data further.

Third, consider the above, for a much smaller sample space such as the CNU's connected to a given OCU, we may not see that kind of distribution; more true in N+1 or N+0 environment.

NOTE: Comcast is analyzing collected data from N+0 nodes.

Questions from 12/11/12

- Is there an assumption that plant conditions will change by the time that EPoC is deployed, such as will there be more nodes that would be deeper in the network?
 - We want EPoC to work in the network as we have it today
 - YES, there will be more nodes than we have today
 - We expect that cascades will be shorter which should allow for better SNR
 - However, we expect that SNR distribution will remain, and perhaps distribution spread will increase
 - Biggest contributors to SNR spread will likely remain constant: drop length variation, feeder cable and taps, equipment location in the home, difference in equipment implementation, etc., and perhaps there may even be more variation than there is today as we move to incorporate more commercial services, where EPoC and legacy services could be deployed in different ways.
- If we assume that there won't be variation, then we will have to use the lowest common denominator MCS and because of that capacity/throughput in the EPoC infrastructure will be lower than could be.

Previously presented related materials

Presentations related to data


- mallette_01_1112 from November IEEE F2F
- schmitt_01a_1112 from November IEEE F2F

Presentations containing analysis


- elbakoury_01_0912 from October IEEE F2F
- varanese_01_0912 from October IEEE F2F
- boyd_02_1112 from November IEEE F2F
- dai_01b_1012 from October IEEE F2F and expanded in dai_01a_1112 from November IEEE F2F
- garavaglia_01a_1112 from November IEEE F2F

Data collected by Rogers Cable represents D2 CMs, not including D3 and MTAs

SNR Distribution of Cable Modems



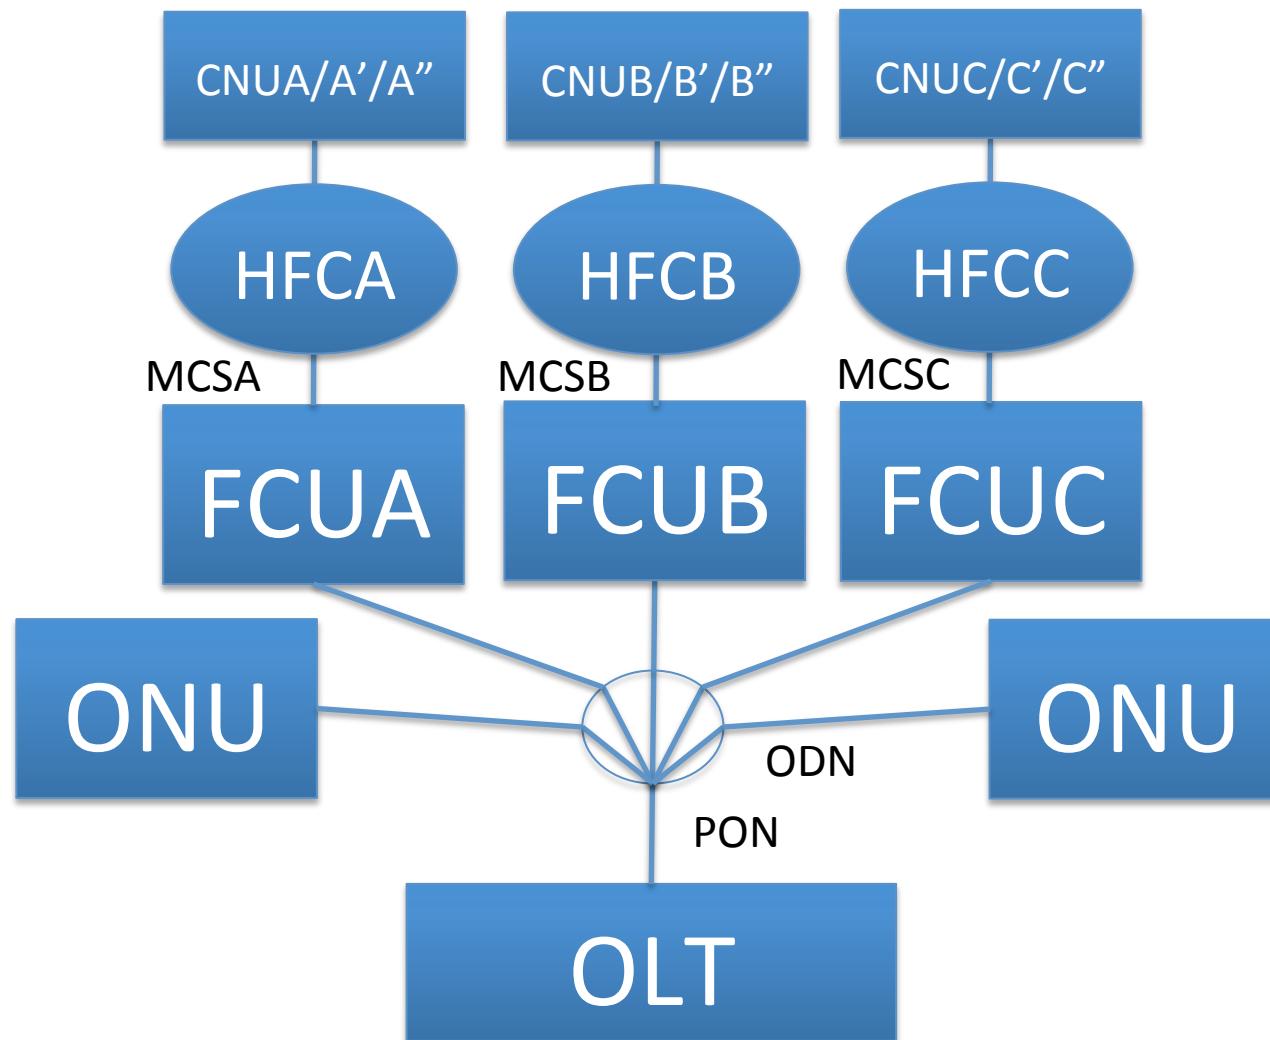
Data collected by Shaw Cable



Notes from 12/14/12 meeting

- See separate presentation from Andrea Garavaglia

Consider the following case...



- Assuming a single MCS per FCU, can the CNUs in FCUA operate on a different MCS than the CNUs in FCUB and/or FCUC?
- Assuming that the answer to the above is that it is possible, then, somehow, the DBA in the OLT would have to account for the different MCS in CNUs on the same PON, is that right?
- If that is the case, wouldn't this be the same whether the CNUs are in the same FCU or different FCUs?

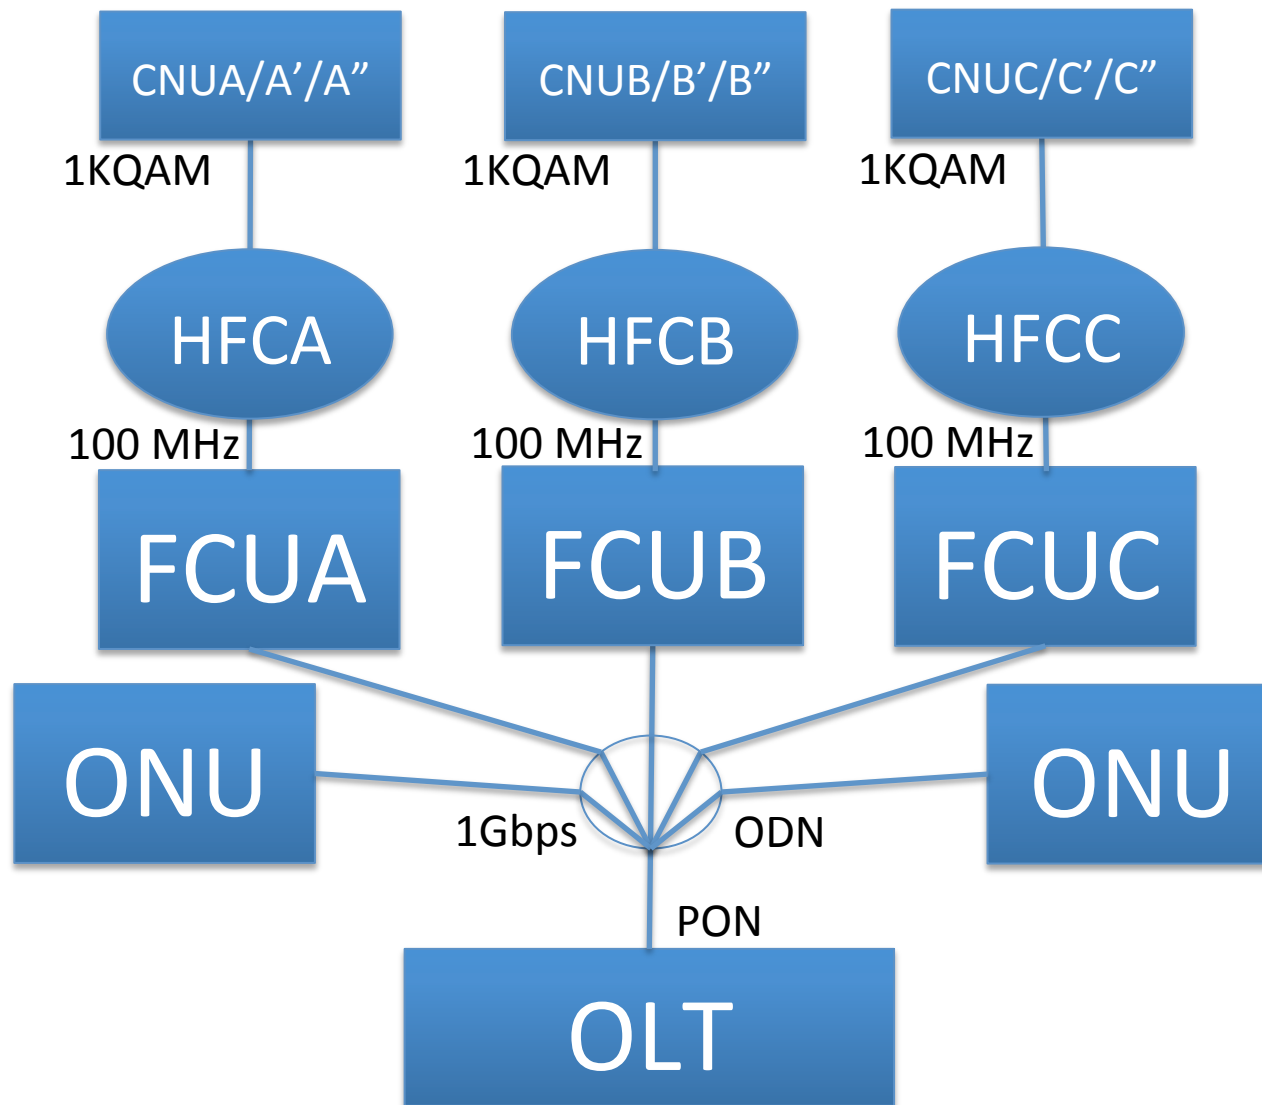
Notes from 12/18/12 Meeting

- We reviewed Ed Mallette's presentation (see separate attachment)
- The key message of Ed's presentation is that regardless of whether Multiple MMP is implemented in a single FCU or not, at the very least use case 1 (from Ed's slides), in which multiple CNU's would operate in different modulation profiles in a single PON or MAC domain, is expected to be supported by EPoC
- No one on the call was able to confirm that such use case would be possible or that it would not be possible. This question will be posed to the entire group to try to get an answer on viability
- We reviewed Dave Urban's additional slides, presenting data collected for only MTAs, including plots for individual nodes, and including SNR vs. US TX level and SNR vs. DS RX level
- The data shows that even in cases of a single node, and even in cases of demonstrated proper CM operation, there is a clear difference in SNR that spans at the very least between 3

Notes from 1/2/13 Meeting

- There weren't any presentations for today on the specific reasons for not being able to implement MMP.
- We discussed the diagram included in slide 22. There was much discussion about what would or not be possible. No specific conclusions were established. I attached a transcript of the discussion which was taken by Joe Solomon (see separate document).
- To facilitate the discussion during our next call I expanded the single diagram into 4 diagrams, included in slides 25-28.
- We will continue the discussion during our next call, on Friday 1/4/13 at 9 AM ET (see separate meeting invitation).

Simplest use case




Set-up:

- 1 Gbps PON
- Same spectrum on all HFCs
- Same MCS on all FCUs

Possible implementation?

- OLT sends full frames to ONUs, but padded frames to FCUs to make sure RF side of FCU is not overrun with data (FCU removes padding)
- OLT needs to know the MCS used by the FCUs (same for all)
- Frames to all CNUs contain the same amount padding

More complex use case



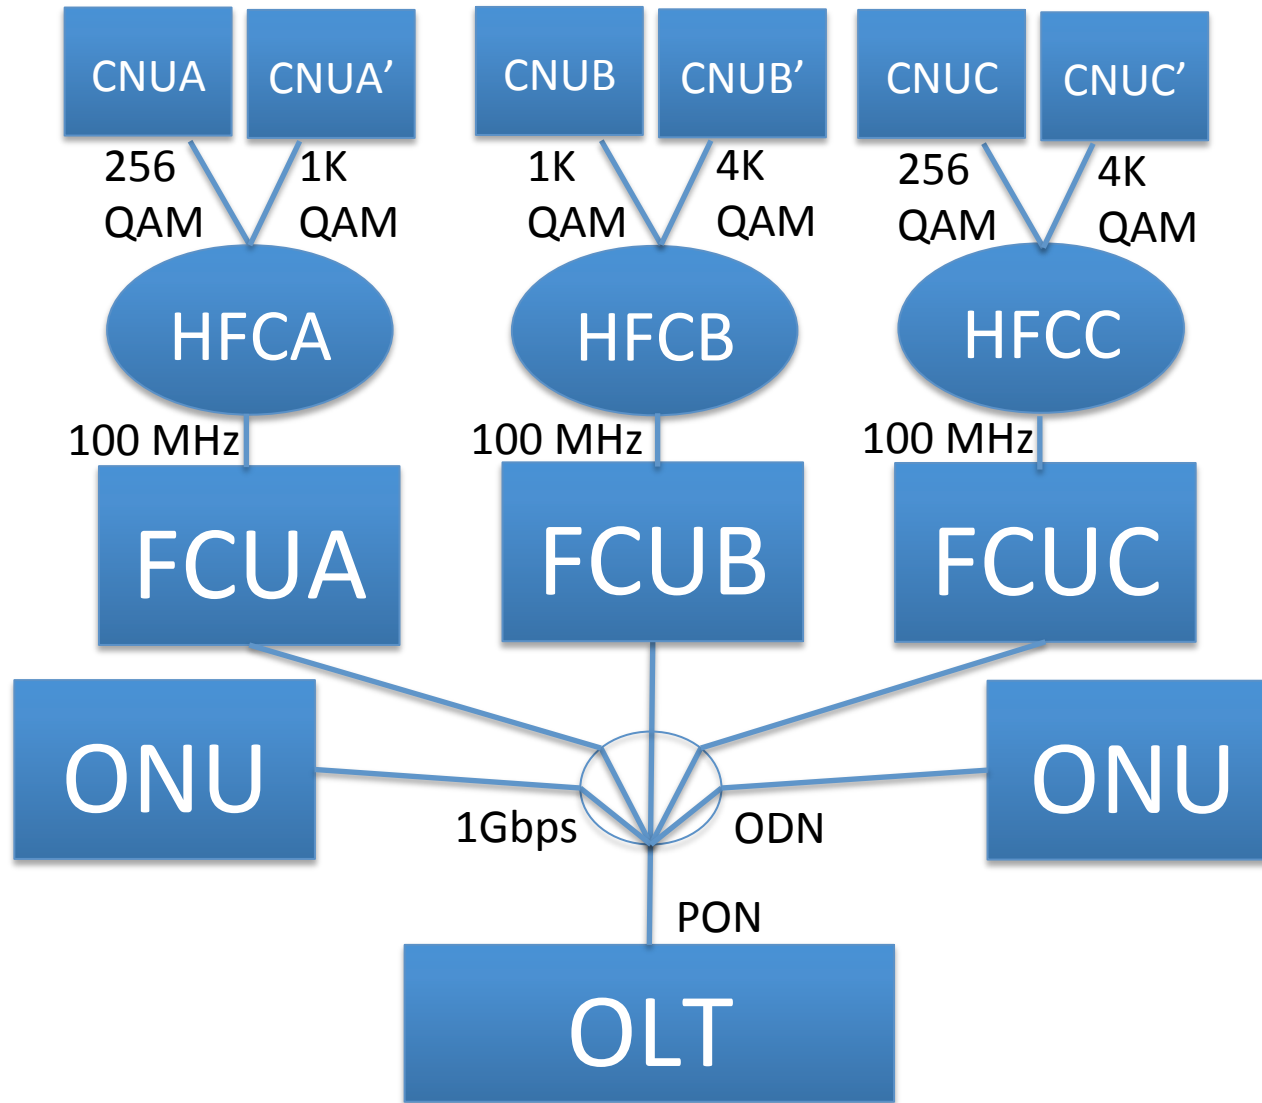
Set-up:

- 1 Gbps PON
- Same spectrum on all HFCs
- **Different MCS on FCUs**

Possible implementation?:

- OLT sends full frames to ONUs, but sends padded frames to FCUs to make sure RF side of FCU is not overrun with data (FCU removes padding)
- OLT needs to know **the MCS for each FCU**
- Frames to all CNUs **within each FCU** contain the same amount padding, **but within different FCUs contain different padding**

Even more complex use case




Set-up:

- 1 Gbps PON
- Same spectrum on all HFCs
- **Multiple MCS on FCUs**

Possible implementation?:

- OLT sends full frames to ONUs, but sends padded frames to FCUs to make sure RF side of FCU is not overrun with data (FCU removes padding)
- OLT needs to know **the MCS for each CNU**
- Frames to **each CNU contain different amounts of padding**

Most complex use case



Set-up:

- 1 Gbps PON
- **Different** spectrum on HFCs
- **Multiple MCS on FCUs**

Possible implementation?:

- OLT sends full frames to ONUs, but sends padded frames to FCUs to make sure RF side of FCU is not overrun with data (FCU removes padding)
- OLT needs to know the MCS for each CNU **and the allocated channel capacity**
- Frames to **each FCU contain different amounts of padding**