Cisco DOCSIS 2.0/3.0 Cable Modem Upstream;

Ranging, Levels, MER(SNR), & T3/T4 Timers

John J. Downey – Cisco Systems
CMTS Technical Leader
4/22/17

Initial Maintenance Ranging

A new CM will start ranging typically around 6 -9 dBmV in 3 dB steps until the CMTS "sees" it, which is about -20 dBmV at the CMTS. Once the CMTS "sees" it, the CM will report inti(r1) as they are doing initial maintenance (IM). This is contention time and CMs will backoff when, and if, they collide. The Cisco CMTS incorporates a backoff called; `range backoff 3 6`. This translates to a backoff between 2^0-2^3 (1-8 opportunities) and the next time between 2^3-2^6 (8-64 opportunities).

Note: The CM should quickly go from broadcast IM to unicast station maintenance (SM) for final ranging and report init(r2). This does not have to be 3 dB steps anymore.

The Cisco scheduler has dedicated time every 60 msec for IMs. We also have a command called `insertion-interval` that is `auto 60 480` on the 10K and `auto 120 1000` on the cBR-8. The amount of modems online and traffic utilization will make the insertion interval automatically change between those numbers.

In the example of a CMTS reboot, the insertion interval would be the lowest of every 120 msec and a bunch of modems would be in inti(r1). Assuming collisions are happening, they would backoff randomly that first time between 1-8 opportunities, this means between 1*120 and 8*120 = a backoff anywhere from 120 msec to .96 sec.

In regards to registration and potential contention, this will cause uncorr FEC because FEC T bytes are present in the IM burst modulation profile. This is why it's **not** recommended to track uncorr FEC till all the CMs are registered. Normal contention could cause incrementing uncorr FEC.

```
cbr8#sh cab modu 221
Mod  Mode   IUC     Type Pre Dif FEC  FEC  Scrmb Max Grd  Las Scrmb Pre Pre   RS
                         len enc T    k    seed  B   time CW        off Type
                                 BYTE BYTE       siz size sh
221  atdma  request qpsk 36  no  0x0  0x10 0x152 0   22   no  yes   0   qpsk0 no
221  atdma  initial qpsk 98  no  0x5  0x22 0x152 0   48   no  yes   0   qpsk0 no
```

Tip: Look for modems stuck in init(r1) as they could cause issues by "eating up" limited IM opportunities and cause uncorr FEC counters to appear high.

In regards to US ranging for D3.0 mtc-mode (US bonding), we suggest the cable interface command: `cable upstream ranging-init-technique 2.` To eliminate contention ranging on the other USs in the US BG once the first US has ranged. The default is tech 1, which means contention IM. Tech 2 is unicast, so basically SM burst. We had some issue with some modems with tech 2 in 3.18 code, but tech 3 or 4 could be tried as well. It also helps with RFoG systems and also for DS D3.0 LB.

Station Maintenance Ranging

A DOCSIS cable modem (CM) is polled at least every 20 seconds on the Cisco CMTS during station maintenance (SM). If linecard redundancy is configured, then SM is sent every 15 seconds. Refer to Appendix A for information on how to verify the SM period.

Note: This assumes 15-20 modems on the cable interface (mac domain). If < 15, then SM is every **n** sec where **n** is the number of modems. So, in a lab environment with only say 5 modems, the SM could be every 5 sec per CM. The automatic polling feature can be turned off for lab testing with the hidden, global test command; `test cable minimum-poll off`, then the rate can be set with the interface command `cable upstream ranging-poll interval <20000 - 30000>` in milliseconds. The default is 20,000 milliseconds.

Typically, three parameters are adjusted during SM; level, frequency, and time offsets. If upstream (US) equalization-coefficient is configured, then pre-eq taps are adjusted during the SM also.

- US levels can "swing" +/- 3 dB over days, months, years; but, within a 20 sec window it shouldn't be more than tenths of a dB unless customers are changing splitters or some other outside attenuation change.
- Frequency should not be drifting from the CM unless it has a bad US transmitter.
- Timing should not change by more than couple time offset ticks, but it has been observed on some faulty CMs to increase over time and cause serious issues with dynamic map advance.
- In regards to pre-eq, the Cisco CMTS has code to reset the CM pre-eq (called direct load) if the MER drops more than 3 dB in one SM period. The amount of direct loads initiated on a CM can be monitored with the `scm "mac address" verbose` command.
    - **Note**: This can be disabled with the cable interface command; `no cab upstream 0 equalization-error-recovery.`

Station Maintenance can be thought of as a 3-way "handshake". The first step of the 3-way "handshake" is where the CMTS will command the CM to adjust its level if it is outside a **threshold** from the **nominal** receive level. The **nominal** receive level default is 0 dBmV and can be changed with "`cable upstream n power-level y`**".** The value of "y" is controlled by the channel width selected and the range is listed in Table 1 below. The default **threshold** is +/-1 dB and can be changed with "`cable upstream n power-adjust threshold y`**"** The threshold range (y) can be changed from 0 to +/- 10 dB. A big threshold eliminates many adjustments, but will produce CMs with different carrier-to-noise ratios (CNRs). The flaplist also has a range of +/-1 to +/- 10 dB for reporting purposes and explained later.

Note: Never set the `power-adjust threshold` range to 0 because CMs will never set up unless they hit the CMTS exactly at 0 dBmV and ranging opportunities will be taken up by CMs continually changing levels. The Flaplist will be very active! The default range of +/- 1 dB should be sufficient, but a range of +/- 2 may be justified for temperature swings you don't want to track.

The CM will be commanded to change its power in units of .25 dB by the CMTS, but the CMs only have 1 dB of granularity. The CMTS may instruct the CM to change by 14, which means 14 * .25 = 3.5 dB, but the CM will only change by 3 or 4 dB.

The CMTS US level as stated by DOCSIS 2.0 and lower is listed in Table 1 below.

Symbol Rate (kHz)	Maximum Range (dBmV)
160	-16 to +14
320	-13 to +17
640	-10 to +20
1,280	-7 to +23
2,560	-4 to +26
5,120	-1 to +29

Table 1 – D1.0, 1.1, and 2.0 Sym Rate vs CMTS US Level Range

Symbol rate is multiplied by 1.25 to get channel width (CW). D3.0 only requires CWs of 1.6, 3.2 & 6.4 MHz to be supported. The CMTS D3.0 US extended range level is listed in Table 2 below.

Symbol Rate (kHz)	Maximum Range (dBmV)	Applicability
160	-13 to +17	CMTS MAY support
320	-13 to +17	CMTS MAY support
640	-13 to +17	CMTS MAY support
1,280	-13 to +17	CMTS MUST support
2,560	-10 to +20	CMTS MUST support
5,120	-7 to +23	CMTS MUST support

Table 2 – D3.0 Sym Rate vs CMTS US Level Range

Note: The extended US range listed in Table 2 above for D3.0 can be set if global modulation is activated on the CMTS, `cable modulation-profile global-scheme`. This is on by default on the cBR-8, but not the uBR10K.

The nominal level can be set depending on channel width, but it's easier to manage if left at 0 dBmV and less chance of total power overload from other ingress. DOCSIS 2.0 RFI specification 4.2.12 states that the total input power to the US demodulator must not exceed 35 dBmV for the entire 5-42 MHz band, but the DOCSIS 3.0 spec mentions < 29 dBmV total power from 5-85 MHz.

Step 2 of the 3-way "handshake" is where a CM will adjust properly when a CMTS commands it to. If the CMTS receives a level outside the threshold again, it'll poll the CM immediately within one second, "fast poll". This means Step 1 starts over. This could be because of incorrect readings outside the threshold, no response at all, bad timing, noisy US, etc.

Note: If the level of noise on the US is enough to distort the US level being received by the CMTS, then the CM and CMTS will go into "`power-adjust noise`" averaging mode. A "*****" will be displayed next to the receive level in the `show cable modem` command. When this occurs, CMs are polled using a one second interval. By default, the percentage of "noisy" ranging responses that cause a CM to enter "`cable upstream n power-adjust noise`" mode is 30%. This percentage may be increased to alleviate excessive power level adjustments in the presence of noise.

Step 3 of the 3-way “handshake” would indicate a success ranging sent from the CMTS to the CM.

Ranging Notes

`sta` = station maintenance (good state)
`cnt` = continuous ranging (ranging miss or adjustment is needed)
`dr` = down & recovery with extended ranging opportunities

Continue ranging happens when an us channel ranging miss, or adjustment is needed for a modem. In this condition, cmts will keep ranging polling against the modem with a very short interval (1 second) to recover the ranging miss or adjust the parameters; If 16 times continuous ranging polling failed for an us channel of modem, the us channel will be in "dr" state, which indicating the us channel is put to "down", but modem still has possibility to recover it, cmts keep sending ranging polling opportunity to this channel with a normal interval (not 1 second); If ranging miss recovered and parameters adjust success, the us channel will be in "sta" state.

Yes, "dr" means the communication failure and able to be recover.

If there's a communication failure and fixed it to normal, yes, the state changes is: dr->cnt->sta.

If previously the us channel is good, and then communication failure happened, the states change is: sta->cnt->dr.

Note that the "cnt" is a very short state stage, we usually see it only several seconds, it looks like a intermediate state. Most of time the modems us channel states are "sta" or "dr".

Levels

The CMTS US channel uses the equivalent of a bandpass filter that is tuned to the configured frequency for the US port. So it is measuring occupied channel power of the 3.2 MHz signal. If the US receiver is tuned to 30 MHz in the CMTS config with a 3.2 MHz CW, then it would only be looking at 28.4 MHz through 31.6 MHz. The adjacent channels energy is not a factor in the measurement. Technically, it would be symbol rate, so slightly less than ch width.

The HFC plant should be designed, so that in normal DOCSIS operation, CMs transmit between 40 and 48 dBmV. This gives enough headroom for customer induced losses, modulation changes, temperature related issues, and age factors along with US bonding limitations. By designing a lower bound of 40 dBmV, it helps alleviate noise induced from low-value taps (4, 8, 11 dB). Special equalized taps and feeder EQs with return frequency pads are available to achieve this.

Levels to keep a CM online are done during SM and each CM vendor may have implemented their preambles of the SM burst in the mod profile differently for QPSK versus 16-QAM. It is very possible that changing the SM burst to 16-QAM could make the CM appear to transmit 3 dB higher. Keep in mind that while the maximum US Tx power required by DOCSIS is +58 dBmV for a CM using QPSK, a CM using 16-QAM only needs to transmit at a maximum power of +55 dBmV. This may have an impact in cable systems where the total US attenuation between the CM and the CMTS is higher than 55 dB. Excessive US attenuation is usually related to subscriber drop problems or network misalignment.

When using DOCSIS–mode ATDMA, the preamble is either a QPSK0 or QPSK1 pattern leading to potentially 3 dB difference in amplitude vector on the constellation. Using 16-QAM and higher for the actual burst typically uses a QPSK1 pattern.

The reporting capability of CMs could be +/- a few dB. Who knows how accurate they really are and it could be different accuracy when doing QPSK vs 16-QAM. Some CMs report their transmit level based on the "long" or "a-long" burst even though they are "leveled" based on the SM burst. Using a mixed modulation profile could lead to erroneous results.

One issue when doubling channel bandwidth (BW) is MER affect. The CM is not told to increase its power to follow power/Hz calculations, but to keep a constant average power (Cisco Implementation). Some call this power/Hz or Power Spectral Density (PSD). This can have the negative result of 3 dB worse CNR since the noise BW has doubled, but the channel avg power has remained the same. This can also make the MER (SNR) drop much more than 3 dB because it will include linear impairments like group delay and micro-reflections in the wider chancel. Pre-eq is typically the answer to this issue. If doubling the BW automatically increased the expected US level at the CMTS by 3 dB to keep a constant CNR, then the CM could reach max transmit power in some situations. The work-around to this would be to ether increase the CMTS US configured receive level to make CMs transmit higher or add more padding inline, assuming max Tx is not going to be an issue.

1. If you change the CW from 3.2 MHz to 6.4 MHz, you may want the carrier to be 3 dB hotter. You can change the input level on the CMTS from 0 to +3 dBmV, which will make the CMs transmit 3 dB more and you'll have the same CNR you had to begin with.

2. If you install 3 dB more padding in front of the CMTS US port, the CMs will transmit 3 dB higher, but the total noise and interference on the front-end of the CMTS US port will be 3 dB lower. This is the preferred method with a side benefit that configurations do not change. It's advantageous to use the US combining networks made by PCI technologies, QRF, ADC, Cisco (SA), etc. They have plug-in pads that can be changed without disconnecting cables. This is difficult to achieve if multiple US chs of different widths are on the same connector.

3. Another scenario is installing 3 dB padding in front of the US laser in the node. This provides less chance of laser clipping, less noise in the node, and the same overall affect at the CMTS. The pitfall is you'll be adjusting the node receive levels for that one service. What about Settop signals, telephony, etc.? It also takes more work!

Remember to observe the transmit levels from the CMs to see if they are reaching their maximum transmit power, especially the ones connected to high-value taps. You can do `sh cab modem remote-query` (if configured) or `sh cab modem phy` to display the CMs' outputs and `sh cab modem` to see if there is a "`!`" displayed next to the receive column indicating the CM is at its maximum output power.

Max Transmit Levels

A CM's max Tx is dictated by the US modulation scheme used and the actual vendor implementation. Max output is listed in Table 3 below. If the CM has reached its maximum transmit power, a "**!**" will be displayed next to the receive level in the `show cable modem` command. The CM will be permitted to stay online if it is within a certain range. This range has a default of 4 dB from nominal and can be changed with; "`cable upstream n power-adjust continue y`". The value of y can be increased up to 15 dB, but many systems have settled with the default of 4 or increased it to 6 dB.

Note: A max transmit CM will be commanded to change level every 20 seconds during its SM, optimally only once. Some CMs have been observed to go into the fast polling mode for 5-10 times before moving on. This is a good reason to make sure < 5% of your CMs are in this state.

There is no special message sent by the CM when it reaches max transmit power. If the CMTS keeps telling the CM to increase its transmit power several times in a row, the CMTS concludes that the CM cannot increase its transmit power any higher and flags it with "**!**".

Note: There are some corner conditions where the flag may be erroneously set. For example, if the CMTS commands the CM to adjust its power by 1 dB and the CM only adjusts its transmit power if the request is above 1 dB, the CMTS may erroneously conclude that the CM is out of transmit power. Thus, besides looking at the "!", the customer should also look at the actual reported power level and see if it is outside the target power level. The real intent of the "!" is for troubleshooting and identification.

By increasing the "continue" command to 6 dB, the CM will be permitted to stay online if the CMTS receive level is between -6 dBmV and 0 dBmV. If the level is above -1, you won't see a "!". If the level is below -6 dBmV, the CM will go offline. For systems that still have high-value taps (29 & 26 dB), this helps keep the CM online, but will produce CMs with different CNRs & MERs.

Warning: Allowing a large power-adjust continue to be configured can lead to CMs having a large range to overcome isolation and appear, potentially, on US ports where they should not! It could also allow CMs located off low value taps to range very high and create intermittent laser clipping.

DOCSIS 3.0 Level Issues & Initial Ranging

Activating multiple frequencies per US connector on a D3.0 CM has different maximum power per channel vs. a D2.0 CM as shown in Table 3 below.

Modulation	Max Tx – D2.0 (dBmV)	Max Tx – Single Ch D3.0 (dBmV)	Max Tx - D3.0 2 Chs (dBmV)	Max Tx - D3.0 4 Chs (dBmV)
QPSK	58	61	58	55
8-QAM	55	58	55	52
16-QAM	55	58	55	52
32-QAM	54	57	54	51
64-QAM	54	57	54	51
SCDMA	56	56	53	53

Table 3 – Modulation vs Max Transmit Level

As explained above, it can be seen that the max power for one channel has been raised by 3 dB over a D2.0 CM, but max transmit per channel for four frequencies stacked using 64-QAM, ATDMA is only 51 dBmV & 53 for S-CDMA.

In existing plant (D2.0 and <), CMs range and the CMTS sends deltas, but never really knows the actual CM Tx power. The Cisco CMTS has code that allows a CM to register and/or stay online if it is within the "`power-adjust continue`" window. This does not happen with D3.0 CMs for initial ranging and registration. The spec says the CM Tx level will be reported to the CMTS on the first ch the CM ranges on. This does not take into account the "`power-adjust continue`" window except after registration for on-going station maintenance ranging.

D3.0 increased CM Tx by 3 dB, but when we stack 4 chs, it drops by 6 for an overall affect of 3 dB less output for specific modulation schemes. Hence, 54 dBmV for D2.0 (64-QAM) equals 51 dBmV for D3.0 (64-QAM) with 4-ch stacking. When a D3.0 CM registers, it does so on a single channel, a reference channel, and relays its Tx level back to the CMTS. The CMTS can determine if that level will be adequate for 4-ch bonding. It isn't using the "`power-adjust continue`" range for that determination.

Note: Cisco has a feature that will drop from 4-ch to 2-ch (if configured) and finally single-ch and/or non-mtc mode. This depends on the Tx level supported plus it adds in the `max-channel-power-offset` calculations. This command has a default of 3 dB, but a value of 6 is recommended. If that level is not adequate for all options, then the CM resets itself. Example; CM ranges on US0 and reports 55 dBmV, CMTS wants to do 4-ch bonding and determines that 64-QAM for 4-ch US bonding has a max output of 51 dBmV + 3 max-ch-offset = 54, CM drops to 2-ch BG, if configured.

Note: A DDTS was implemented to extend this range to 12 dB and also drop to single-ch mtc-mode. D3.1 CM DS operation requires US mtc-mode or the CM may not register.

If the CM is fine with the reference channel, it still needs to finalize initial maintenance/ranging with the other US channels. A 12 dB dynamic range is allowed to accommodate different losses at different frequencies. This is referred to as Dynamic Range Window (DRW). Example; CM ranges on US0 and reports 48 dBmV, CMTS wants to do 4-ch bonding and determines that 64-QAM for 4-ch US bonding has a max output of 51 dBmV. The CM registers and the dynamic window is set for 42-54, but 54 is beyond max, so the window is reset to 39-51 dBmV dynamic window. US1 needs 50 dBmV and succeeds with ranging. US2 needs 52 dBmV, which is outside 51. It will report 51 Tx, and the CMTS will report -1 dBmV. US3 needs 56 dBmV, it will max out at 51 and hit the CMTS at -5 dBmV. The dynamic window may adjust multiple times.

Note: This difference in CM Tx levels is typical because of US frequencies and cable attenuation.

If the D3.0 CM does successfully register and the levels change slightly during normal periodic ranging, the "`power-adjust continue`" will be taken into account and the CM will stay online and be marked with "!". Example; CM ranges on US0 and reports 50 dBmV, CMTS wants to do 4-ch bonding and determines that 64-QAM for 4-ch US bonding has a max output of 51 dBmV, the CM registers. Temperature goes up, which adds more attenuation to cable, and CMTS reports -3 dBmV instead of normal 0. CMTS sends SM to CM to change by 12 steps of .25 dB and CM only has 1 dB more range left. CM transmits 51 dBmV (its max), CMTS reports -2 dBmV and goes into fast poll mode to change again. The CM cannot change and the CMTS determine that -2 is within the `power-adjust continue` of 4 and allows the CM to stay online and it is marked with "!".

Tip: Cablelabs released an ECN allowing CMs to transmit ~ 3-6 dB higher than what is listed in table 3. Some CMs will support this and some may not.

Note: There may be a case for dynamic modulation and some channels could drop to QPSK. This modulation scheme allows 4 dB more power over 64-QAM! If D3.0 CMs happen to register on that channel first using QPSK, then the logic to allow full bonding is based on that channel's power and not the other chs.

Potential Minimum Transmit Issue

Minimum transmit has never been seen as a big issue. DOCSIS 3.0 has set the minimum level much higher than it used to be with D2.0, which was as low as 8 dBmV. This minimum CM transmit level is related to channel width and displayed in Table 4 below.

Symbol Rate (kHz)	Minimum Tx (dBmV)	Applicability
160	17	CM MAY support
320	17	CM MAY support
640	17	CM MAY support
1,280	17	CM MUST support
2,560	20	CM MUST support
5,120	23	CM MUST support

Table 4 – Sym Rate vs Minimum Transmit Level

One potential issue can arise when using a test modem in the headend and not providing the proper attenuation. If this D2.0 test CM ranges and comes online with a transmit level of 20 dBmV, which is legal, then it also has plenty of range left in its transmit. If it ever re-registers and ranges on the wrong upstream channel descriptor (UCD) it can cause packet collisions on other US ports. If this example were using 16-QAM with a max output of 55 dBmV, then it could range all the way to 55 dBmV, which is 35 dB more than it needs. Depending on the headend wiring and architecture, this signal could find a path to "bleed" across because of isolation issues or lack thereof. Because the CMTS can typically "see" a CM around -20 dBmV, this re-ranging CM could be detected on the wrong US port even with 35+20 = 55 dB of isolation design. This issue is exacerbated with a large `power-adjust continue` configuration.

Another issue pertains to lab setups and D3.0 and/or SCDMA. Many lab environments will not have attenuation to force CMs to transmit in the 40-50 dBmV range. The min range when using a D3.0 CM, 6.4 MHz ch is 23 dBmV. This is specific to mtc-mode being activated for US ch bonding. When testing with SCDMA, that min could be as high as 42 dBmV! According to D3.0 RFI spec, Min Tx = max[17 - Gconst] + 10 log (number_active_codes / #_of_codes_per_minislot), where the max is over all burst profiles used by the CM in ch x. The bottom line is, CMs may not register because of min Tx requirements not previously seen with D2.0. One example would be SCDMA using 6.4 MHz ch with QPSK for ranging and 64-QAM for data. Depending on active codes & codes per minislot, the min could be as low as 24 dBmV or as high as 42 dBmV!

MER(SNR)

As explained earlier about changing channel width and the affect on CNR, it can be even worse for MER(SNR). MER is modulation Error Ratio and SNR is signal-to-noise ratio from the US chip, whether from Broadcom or TI. This is after the carrier has been eliminated to give a pure baseband, signal-to-noise ratio. CNR is carrier-to-noise ratio, which is before it reaches the chip (RF domain) and measured with a spectrum analyzer. Doubling the CW could make MER much worse because of group delay and other linear impairments like micro-reflections, which do not show up on a spectrum analyzer in the frequency domain. Also, the Broadcom–based linecards could report better MER because it's reporting after internal EQ takes place, whereas the TI-based cards do not.

Note: A command can be configured on the US port to activate pre-eq and achieve better MER; `cable upstream n equalization-coefficient`**.** This allows the CM to "pre-distort" its signal and "hit" the CMTS US port with a "clean", flat signal. DOCSIS 1.1 CMs support something called an 8-tap equalizer while D2.0 & 3.0 CMs support a 24-tap equalizer, assuming US docsis mode is atdma. D1.0 CMs have no support for pre-eq and may need to be excluded from pre-eq to avoid potential issues.

There are several factors that contribute to the inaccuracy of the MER calculations and why monitoring multiple variables is recommended. Refer to Table 5 below for other monitoring ideas and how the type of impairment can affect those readings. The MER is averaged for all CMs, so this could be subjective. Because of the way MER is calculated, impulse noise and transient sources of noise like laser clipping cannot be detected. The actual payload/utilization could affect the average reading. Because CMs range on the SM burst, but may report their levels on the Long burst, you could have differences in readings. Also, D2.0 CMs support a preamble of QPSK, but 2 versions; QPSK0 and QPSK1. This slight difference can result in 3 dB difference. The MER formula is better associated with additive white gaussian noise (AWGN).

Impairments	CNR	MER(SNR)	Corr FEC	Uncorr FEC
AWGN	Bad	Bad	Bad	Eventually Bad
CW Carrier	Bad	Ok	Ok	Ok
Impulse Noise	Bad	Ok	Ok	Bad
Group Delay / Micro-Reflections	Ok	Bad	Bad	Eventually Bad
Laser Clipping	Ok-ish			
No traffic		Ok	----	----
Distorted traffic		Bad	Bad	Eventually Bad

Table 5 – Impairments vs Error Tracking

Uncorr FEC is the bottom-line for customer quality of experience (QoE). MER is the good indicator of a degrading US channel, but as mentioned earlier, impulse noise will not be tracked by MER properly, so other parameters should be verified as well.

- If the impairment were a steady-state, relatively narrow carrier like a continuous wave (CW) carrier, the linecard's built-in ingress cancellation would digitally "erase" it and report good MER even though CNR was very bad. Ingress cancellation will also cancel some common path distortion (CPD), but CPD resembles AWGN when all DSs are digital.

- Impulse noise would not be “seen” by the US chip, so MER, which is averaged, would report fine when in fact many packets or codewords could be dropped (uncorrected FEC). Dynamic US interleaving in the atdma modulation profile may help.
- In the case of group delay, this is an impairment that would not show itself as poor CNR, but if would affect MER, and eventually FEC codewords. Pre-eq can help alleviate this.

Note: Many systems rely on the Flaplist to indicate plant stability. The Flaplist relies on the SM burst which could be misleading since it uses QPSK while the data grants use 16-QAM. Also, troubleshooting with ping (layer 3 uses a “short” or “long” grant at 16-qam) could give misleading results when compared with ping docsis, which uses a SM burst at QPSK. Refer to Appendix B for Ping DOCSIS information. Using 16-QAM for SM keeps ping docsis the same as regular ping and the Flaplist as vulnerable as regular data. Using 16-QAM for SM will allow the CM to flap.

T3 & T4 Timers

There are two timers in the CM that can be utilized for troubleshooting, but they reside in the CM and must be queried. The T4 timer is a 30 sec timer that is reset in the CM every time step 1 of the 3-way “handshake” occurs. If a CM doesn’t receive a SM message within its T4 timer, it will go offline and restart scanning DS. If a DS issue is present, such as disconnected DS fiber link, the CM will never get its SM. The CMTS will go into its fast, 1 second polling mode and try 16 times total before giving up and quit sending DS sync messages and SM messages.

The T3 timer is a 200 msec range response timer and increments in the CM every time the last 2 steps of the 3-way “handshake” fail. This can usually be analogous to “misses” in the CMTS US Flaplist. US issues contribute to T3 timeouts, but 3 or so T3 timeouts every 20 seconds is not cause for alarm. This could be from normal US noise or timing. Typically, T3 timeouts are caused by a noisy US like laser clipping, bad timing (could be CM issue), and/or disconnected US path. If, for example, a technician removed an US fiber and the CM was just about to get a SM burst, it would still have DS connectivity and be able to reset its T4 timer. The T3, 200 msec timer for the last 2 parts of the 3-way “handshake” would not occur and it would increment. The CMTS would go into fast polling mode and the CM would report another T3 timer. This would occur 16 times total. The CMTS would quit sending DS sync and SM and then the CM would fail its T4 in 30 sec. and start to rescan for DS.

Note: T3 is only tracked for init maintenance of 1.0 CMs, but for IM and all SM bursts on D1.1 and above. This means a 1.0 CM could be causing the issues, but it may not report many T3 problems.

Once a CM is placed in the Flaplist, a “hit” should be recorded for a successful 3-way “handshake” every 20 seconds. DOCSIS 3.0 added a configurable T4 timer or multiplier since SM is done on each US channel that is used for bonding. This may be advantageous to alleviate many SM bursts.

D3.0 CM Periodic Maintenance, Partial US Bonding Mode, & T4 Multiplier

If the modem does successfully register and the levels change slightly during normal periodic ranging, the “`power-adjust continue`” will be taken into account and the CM will stay online and be marked with “!”. Example: a CM ranges on US0 and reports 50 dBmV, the CMTS wants to do 4-ch bonding and determines that 64-QAM for 4-ch US bonding has a max output of 51 dBmV, so the modem registers correctly. The temperature goes up, which adds more attenuation to cable, and the CMTS reports -3 dBmV instead of the normal 0. The CMTS sends SM to the CM to change by 12 steps of .25 dB and the CM only has 1 dB more range left. The CM transmits 51 dBmV (its max); the CMTS reports -2 dBmV and goes into fast poll mode to change again. The CM cannot change and the CMTS determines that -2 is within the power-adjust continue of 4 and allows the CM to stay online and it is marked with "!". If any of the US channels are outside the continue window, the modem with go into US partial mode and be marked with "p" instead of "UB". Review "`sh cab modem mac-address ver`" for specifics about the CM in question.

CMs have a T4 timer that is 30 seconds. The CMTS typically sends a SM message every 20 secs per CM. If the CM does not get a SM message within that time, it will go offline and start to re-scan DS. Now with multiple US chs for a D3.0 CM, it can become excessive to send every one of those chs a SM burst every 20 seconds. This is where the T4 multiplier comes into play.

The T4 multiplier value is only applicable to CMs operating in mtc-mode and sent in every RNG-RSP message. The T4 timer is for the CM, and not per US.

```
10k(config-if)#cable upstream ranging-poll interval 20000 ?
  t4-multiplier  T4 multiplier ranged from 1 to 10, 0 as auto
```

The multiplier value default is 0 (auto), which is the sum of all chs assigned in the US transmit ch set for that CM. With no MTC-mode, it is one.

Example:
CM is doing 4-ch US bonding (USs 0-3) and the CMTS uses a 4x multiplier by default. We tell the CM that its T4 is 30 * 4 and internally we are scheduling at 4 x 20 sec intervals.
CM has 120 sec T4 timer (30*4).
CMTS sends SM burst to US0 every 4*20 = 80 sec, US1 every 4*20 = 80 sec, US2 every 4*20 = 80 sec, and US3 every 4*20 = 80 sec.

Impairment example: If US0 has noise issues, the CMTS goes into fast polling. The CMTS tries 16 times total and eventually quits sending grants and it reports the US channel down. The CM should go into US partial mode. Partial mode is a CMTS concept where the CMTS stops scheduling grants on the failed channel. You will see a lot of T3 timeouts from the CM perspective or misses in the CMTS flaplist. In this case, you are not going to see T4 timeouts, since there are SM opportunities on other operational channels.

Warning: Because of this T4 multiplier, it could be 2 minutes between MER, level, and pre-eq updates! It may be advantageous to set this multiplier to 2 instead of auto.

The number of ranging misses before a channel is marked off is now configurable by this:

```
10K(config-if)#cable upstream resiliency modem-offline-detect ?
  <4-16>  Number of consecutive ranging misses
```

`Modem-offline-detect 8` is the default. When this criteria is met, it triggers the fast polling on all channels, instead of waiting for the normal polling cycle. This means that after 8 consecutive misses, the CMTS will automatically do a quick poll on all channels. This will speed up marking modems offline instead of waiting for the normal polling cycle. (check this)

The action taken is configurable by this:

```
10K(config-if)#cable upstream resiliency on-failure ?
  disable-channel       disable channel
  extended-ranging      continue to range
  reset-modem           take modem offline
```

The default is “`extended-ranging`” when a ch is down. This allows the CMTS to continue scheduling SM opportunities at the normal interval to give the CM a chance to recover the failed ch.

Question: What happens when there are 2, 2-ch US BGs and there are multiple flows - some associated with one group and some with another? Does this only send fast poll for all chs for that group the flow is on? Or, does it poll all chs in all US BGs which the modem has admitted flows?

Answer: The spec requires the CMTS to mark the CM offline whenever all chs associated to the primary flow are down. Otherwise, the CM remains online in partial mode. We mark a CM offline only when all USs that a CM’s primary US SF is associated are down. Ranging does not care about individual flows in that sense. Ranging cares about ranging all the individual chs. Upon detection of ch going down, ranging informs the scheduler for any flows associated to the bonding groups, in which its ch membership has the same impaired ch, to make sure that the scheduler stops scheduling grants on that US for any of the flows. More Q&A is located at end of this document.

Summary

About 90% of DOCSIS problems are attributed to Physical layer or IOS config issues. Phy layer issues such as US levels, ingress, noise, etc. must be understood and controlled. Having knowledge of CM ranging processes and IOS configuration features can help mitigate these potential problems.

Appendix A
Using "`Show Cable Hop`" to Calculate Current SM Period

```
uBR#sh cable hop
   Upstream     Port         Poll Missed Min    Missed Hop   Hop     Corr    Uncorr
   Port         Status       Rate Poll   Poll   Poll   Thres Period  FEC     FEC
                             (ms) Count  Sample Pcnt   Pcnt  (sec)   Errors  Errors
   Cable3/0/U0 33.008 MHz 800  * * * set to fixed frequency * * * 0       9
   Cable4/0/U0 down          1000 * * *   frequency not set    * * * 0       0
```

Take the "Poll Rate" number, divide by 1000, and multiply by the number of CMs in that mac domain. Example; `Show cable hop (SCH)` says 800 and I have 25 CMs on C3/0. 800 / 1000 * 25 = 20, ~ 1 SM every 20 seconds per CM. If the CMTS sends a SM to each CM once every 20 sec, this would be 1/20 * 25 CMs = 1.25 SMs per second. If the CMTS sends a SM once every 20 sec for 1500 CMs, this would be 1/20 * 1500 CMs = 75 SMs/sec being generated by the CMTS. **Note**: This is another reason why we don't want too many CMs in the max transmit state as they may generate even more SM bursts.

Appendix B
Using "`Show Cable Modem "mac address" flap`" to Calculate CM SM Period

The T4 multiplier can affect the frequency of SM bursts for each US. To verify the current SM period for a specific CM, find a CM in the flaplist. Once a CM is in the flaplist, you can look for the amount of time it takes to register a "Hit". A "Hit" is a successful "3-way handshake" for SM.

```
cBR8#scm 1859.3353.09f0 flap
Time source is NTP, 19:47:38.320 EST Sat Feb 13 2016
MAC Address    I/F            Ins    Hit    Miss CRC    P-Adj  Flap   Time
1859.3353.09f0 C2/0/2/U0      0      2401   15   0      5      5      Feb 13 19:22:49
1859.3353.09f0 C2/0/2/U1      0      2399   19   0      5      5      Feb 13 19:22:59
1859.3353.09f0 C2/0/2/U2      0      2354   16   0      4      4      Feb 13 19:05:19
1859.3353.09f0 C2/0/2/U3      0      2591   20   0      3      4      Feb 13 19:40:44
```

Find the time of a new "Hit".

```
cBR8#scm 1859.3353.09f0 flap
Time source is NTP, 19:47:40.040 EST Sat Feb 13 2016
MAC Address    I/F            Ins    Hit    Miss CRC    P-Adj  Flap   Time
1859.3353.09f0 C2/0/2/U0      0      2402   15   0      5      5      Feb 13 19:22:49
1859.3353.09f0 C2/0/2/U1      0      2399   19   0      5      5      Feb 13 19:22:59
1859.3353.09f0 C2/0/2/U2      0      2354   16   0      4      4      Feb 13 19:05:19
1859.3353.09f0 C2/0/2/U3      0      2591   20   0      3      4      Feb 13 19:40:44
```

Identify the time when a "Hit" increments by 1. It can be seen below that it's exactly 20 sec. This is expected for the T4 multiplier being used in this config. `cBR8#sh run int c2/0/2 | in t4` `cable upstream ranging-poll t4-multiplier 1` (default is 0 which means auto and would be 4 for this example since CM is doing 4-ch US bonding).

```
cBR8#scm 1859.3353.09f0 flap
Time source is NTP, 19:48:00.479 EST Sat Feb 13 2016
MAC Address    I/F            Ins    Hit    Miss CRC    P-Adj  Flap   Time
1859.3353.09f0 C2/0/2/U0      0      2403   15   0      5      5      Feb 13 19:22:49
1859.3353.09f0 C2/0/2/U1      0      2400   19   0      5      5      Feb 13 19:22:59
1859.3353.09f0 C2/0/2/U2      0      2355   16   0      4      4      Feb 13 19:05:19
1859.3353.09f0 C2/0/2/U3      0      2592   20   0      3      4      Feb 13 19:40:44
```

Appendix C
DOCSIS Ping

A novel method of forcing the CM and the CMTS to exchange Station Maintenance keepalives is to use the Cisco CMTS command **"`ping docsis {modem-mac-address | modem-ip-address} [count] [verbose]`".** This command will order the CM to send Station Maintenance Ranging Requests to the CMTS in quick succession in order to verify DOCSIS connectivity to the CMs. By using the "verbose" option, the power, timing, and frequency adjustments sent by the CMTS can be viewed for each Ranging Response message. Time adjusts between +/- 3 ticks aren't tracked.

```
uBR# ping docsis 10.2.2.36 verbose
Queueing 5 MAC-layer station maintenance intervals, timeout is 25 msec:
Reply from 0001.9659.43fd: 4 ms, tadj=0, padj=0.25, fadj=0
Reply from 0001.9659.43fd: 4 ms, tadj=-1, padj=0.25, fadj=0
Reply from 0001.9659.43fd: 4 ms, tadj=-1, padj=0, fadj=0
Reply from 0001.9659.43fd: 4 ms, tadj=-1, padj=0, fadj=0
Reply from 0001.9659.43fd: 4 ms, tadj=-1, padj=0, fadj=0
Success rate is 100 percent (5/5)
```

Note: If regular ping packets (layer 3) generate errors, but ping docsis packets don't, it could be because of the modulation profile. If using a "Mix" profile, the Short and Long data packets for ping will use 16-QAM, while the Station Maintenance packets for ping docsis will be QPSK. Verify the modulation profile to be sure. Ping from the CMTS will default to 100 B and does not include Ethernet headers, so it is really 118 B. This will typically be a short grant. Pinging with 1500 B will lead to 1518 and be a long grant.

I know DOCSIS ping only sends a SM ping to the first US it ranged on and we have a DDTS to address that: CSCuh41631 - Ping DOCSIS & Reporting while CM doing USCB.

I also use a little trick to see which US an US bonded CM ranges on and that is to look at the time offsets in the scm "mac" ver output. If the initial time offset and current time offset are exactly the same, that's usually the first US it ranged on.

One can execute the ping from the linecard console:

```
Slot-1-0#ping doc 0000.dead.1234 ?
  <0-15>   upstream channel
  repeat   Number of maintenance intervals to queue
  verbose  Line-by-line output
  <cr>     <cr>
```

Q & A

From: Dan Neamtu (dneamtu)

What is the impact of “power adjust continue” between the moment the CMTS has decided the TCS size based on the first US the modem has ranged on and the moment the modem is registered? Is it taken into consideration? Example: modem ranging on US0 at Tx=50 dB, CMTS allows a 4-ch TCS and then US1 Tx = 52 dB (US2 and US3 < 51 dB). Is this allowed due to power adjust continue or it will the modem be in partial service and this US in “dr” mode?

It is good you brought up the dynamic modulation profile corner case, maybe an example would make it easier to understand. I get that the decision of the TCS set is taken into consideration purely on the first US the modem has ranged on and this is dependant on its a-long burst mod profile. How about after registration, I assume each of the 4 US will have an upper limit of Tx power based on their own mod profile max Tx + power adjust continue, right?

Example: US 0 has a more robust mod profile which is a mix of QPSK/16QAM. US1-3 are a mix of 16QAM/64QAM.; max power offset = 3, 4ch and 2ch BGs defined. Modem ranging on US0 at 55 dB, CMTS will assign 4ch TCS (based on a-long burst of 16QAM + 3 db max offset).

The modem registers on 4 channels. Further on, the decision to put the USs in “dr” mode is based on the max Tx of each US depending on their mod profiles. Meaning that the max level US0 can transmit at before going to “dr” would be 56 dB (52 max Tx for 16-QAM + 4 dB from power adj continue) as opposed to USs 1,2,3 that will be allowed to transmit at 55 dB (51 max Tx for 64-QAM + 4 dB from power adj continue). Correct?

For MTC-mode modems, the cont command should not be taken into consideration on any US chs until SM polling. The other USs that are still ranging after the first one and the bonding decision tis done, will be based on the DRW.

Modems cannot go to partial mode until they get past the initial decision of which bonding group they can utilize. In DS, modems must fully register and then send cm-status messages back to CMTS to get put in partial mode and/or a resilient bonding group, RBG.

In US, it’s different. After the US BG decision, which should be based on power level, modulation, chs, US setting, max-ch-pwr offset, DRW, etc; I suspect we can place in partial mode if the other USs have an issue with timing, MER, anything that makes SM fail 8 times in a row. The 8 is default and can be changed.

For your example: Example: modem ranging on US0 at Tx=50 dB, CMTS allows a 4-ch TCS and then US1 Tx = 52 dB (US2 and US3 < 51 dB). Is this allowed due to power adjust continue or it will the modem be in partial service and this US in “dr” mode?
US1 would only Tx 51 dBmV and the CMTS would report -2 dBmV, but will be within DRW and I suspect SM would occur at this point and the cont would be used to decide if that US is allowed to stay “STA”.

Your next question: I assume each of the 4 US will have an upper limit of Tx power based on their own mod profile max Tx + power adjust continue, right?

I believe so and think I have seen this before where it ranged first n 16-QAM and reported 52 dBmV and when the other USs at 64-QAM ranged, the modem reported only 51 dBmV max, but still within DRW and power adjust cont and fully bonded.

For last question: The modem registers on 4 channels. Further on, the decision to put the USs in "dr" mode is based on the max Tx of each US depending on their mod profiles. Meaning that the max level US0 can transmit at before going to "dr" would be 56 dB (52 max Tx for 16-QAM + 4 dB from power adj continue) as opposed to US1,2,3 that will be allowed to transmit at 55 dB (51 max Tx for 64-QAM + 4 dB from power adj continue). Correct?
I would say yes, but we need to say it a different way. The modem can overcome X dB of attenuation, where X is the max Tx + the power-adjust cont. Or we say, the modem can hit the CMTS X dB below the optimum and still stay "STA".

Let's not forget that the modem could report 4-ch bonding, but it's really 2, 2-ch BGs. We could even have a case where no BGs created, but MTC-mode activated and the modem reports 4-ch because it decided to use all 4 default single-ch BGs for 4 different service flows.

From: Dan Neamtu (dneamtu)
Hi John, all,

After implementing US Bonding at my customer, they are reporting an issue with USs not going into "dr" mode if received power is below the 4dB default threshold of "power adjust continue".

As you can see in the example below, a 4-ch modem keeps US3 in "sta" mode even if the Received Power is -11 dBmV, whereas I would expect it to go to "dr" mode once the received power gets below -4 dBmV.

I checked and there are multiple modems behaving like this. The power adjust continue and threshold is kept at the default value. In case where the SNR is low, like in this case, the issue translates to customer complaints. What do you think of the below, am I missing something regarding the expected behaviour when US channels are below the power adjust continue threshold?

Dan,

Keep in mind that power-adjust cont is from the power-level ref. The power-level default is 0, but if they made it -3 with power-adjust cont 6, that would allow -9 to -3 dBmV. It seems your scenario is still at 0 with default power-adjust cont of 4, which would/should mean -4 to 0 and will stay online and be marked !. I wonder if your example below is cosmetic or based on old info and needs more time to update?

I agree that modems outside the power-adjust cont would eventually get a T4 timeout, but because of the T4 multiplier we have and using 4-ch US bonding, that could be 30*4 = 120 sec (2 min) before it actually records a T4 timeout and changes to dr.

In regards to low MER causing the SM to drop, then that could be missing the first SM (miss in Flaplist or T3 timeout in CM), then the CMTS should go into fast polling (every 1 sec). With the default detect of 8 and the fact we poll every 20*4 for T4 multiplier = 80 sec+ 7 more for a total of 8, then we mark it down.

Keep in mind that the SM has its own burst in the mod profile and that can exacerbate the issue when still using QPSK while data bursts are 64-QAM or 16-QAM.

Thanks for the reply. I confirm that the power-level is 0 and power-adjust is 4. I have checked now that some entries I was looking at were stay longer than 90 seconds at levels lower than -4 dB, so probably we can rule out outdated info

The customer is using QPSK for SM and 64-QAM for data (they know it is not recommended), that could explain poor service in the US, but is this related to fact that the US is not dropping to “dr”?

Check the attached short log. I picked a 4-ch modem that had two channels below -4 dB, one of them with SNR at the lower limit. Ping docsis works 100%, but normal pings encounters 3% packet loss (after collecting the logs attached I tried again and SNR on the worst US dropped even more and packet loss was 5%). I assume this is perfectly explainable by the modulation profiles difference and low SNR, but this would be avoided if the CMTS would put the two US channels in “dr” and only use the remaining ones.

What is the best way to troubleshoot why CMTS does not put US in “dr” once they drop below -4 dB? And again, can this be related to modulation profiles (QPSK for SM and 64-QAM for data)?

Maybe we’ve been led astray this whole time or something has changed. Maybe this is really basing the decision on the DRW (dynamic range window) and the power-adjust cont is being used only on the first ch it ranged on when it registered? I know DOCSIS ping only sends a SM ping to the first US it ranged on and we have a DDTS to address that: CSCuh41631 - Ping DOCSIS & Reporting while CM doing USCB

IRT the mod profile issue and when a modem is marked other than sta for partial-mode, we have another feature to address that as well: US Bonding Partial Mode - New Feature EDCS-1234521, which should be in SCI2, I think. For now, I suggest 16-QAM on maintenance bursts when using 16-QAM and higher for data bursts. Verify with your modems though and definitely settop boxes (STBs).

Typically, we wouldn’t have modems with such disparate levels, but I have seen the last US in the BG be 3-5 dB delta because of roll-off and so close to the diplex filter and dependent on how many amplifiers between the CM and CMTS and how well pre-eq works, assuming it’s configured.

I assume you are testing by having it register properly with 4-ch US bonding, then adding attenuation or setting CMTS to higher level, which would invoke a max US Tx issue. Can you wait a little longer and see if it finally reaches “dr”; > 2 minutes. BTW, I think newer IOS even reports the DRW values under verbose command.