

SCTE CABLE-TEC
EXPO'15

OCTOBER 13-16 NEW ORLEANS

Managing DOCSIS Capacity over HFC Networks

DOCSIS 2.0, 3.0 and Future 3.1

John J. Downey

Consulting Network Engineer

Cisco Systems



Agenda

- Capacity Planning Questions
 - Capacity “Base-lining” and Overhead
- Factors that Affect Data Throughput & Per-CM Speeds
- Spectrum Allocation Thoughts
- Five Potential Migration Steps and Ideas
- JD’s 10 Rules for Capacity Planning
- Checklist, Closing Points & Resources

Number 1 Capacity Planning Question

- Which type of over subscription (OS) to consider?
 - Concurrency – monitor traffic patterns & determine avg bps
 - Track actual utilization & decide on “watermark”
 - Limit number of devices per service group (SG) or per US ch
- What about future requirements?
 - More video streaming; Hulu, NetFlix, (result of cord-cutters) ...
 - 4K TV, DVR multiple shows, live chs broadcast over Internet, ..
 - Slingbox, Periscope, home security, nanny cams, ...
- What devices must be accommodated?
 - DOCSIS 2.0, 3.0 & D3.1; D1.x still around, DSG, legacy STB
- Factors of 25:1 to as high as 100:1 have been used

Second Major Question

- Do I do node splits or allocate more spectrum?
- Competitive pressure for peak speeds is driving more spectrum allocation
- More node splits = more optics
 - More optics in HE = more US ports, which may require more US RF combining
- Symmetrical design of 1:1 (DS to US) is desirable
 - 1:2 ok, but 1:4 can cause some concern
- Capacity typically added for higher rates
 - Not usually to add more customers or make SGs bigger

DS Capacity Overhead

- Each US in MAC domain can decrease “usable” DS by .25 Mbps
- “Usable” refers to layer 2 speed reporting by CMTS/CM
- Layer 3 reporting by PC could be 5-10% less if avg frames < 1518 bytes
- More overhead: DOCSIS(6), BPI+(5), Bonding(6), VLAN (4), IPv6 = 20 B > IPv4
- Multicast sent down each “primary”
- The “Powerboost™” and peak-rate TLV affect

Powerboost™

- Name trademarked by Comcast
 - Other cable companies may call it by a different name
- Simple manipulation of when to rate-limit CM
- Need very large DS max burst setting in 1.1 CM file
 - IE. 5 MB vs default of 3044 B
 - Trace could appear oscillating slightly up and down
- D3.0 CMs support TLV for per-CM peak rates
- Can also apply US Powerboost

DS Baseline Capacity

- Annex B using 256-QAM, “raw” rate = 42.88 Mbps
 - “Usable” normally quoted as 37.5 Mbps
 - Actual and assuming “Primary” = ~36 Mbps
 - “Secondary-only” = ~37.5 Mbps (no signaling & MAP overhead)
- Annex A uses 8 MHz wide ch
 - ~ 47 Mbps usable rate for 256-QAM
 - ~ 50 for secondary-only
- D3.1, 192 MHz OFDM block (24 MHz licensing?)
 - @1024-QAM = ~ 1.5 Gbps (~ 47 Mbps / 6 MHz equivalent)
 - @4096-QAM = ~ 1.8 Gbps (~ 56.25 Mbps / 6 MHz equivalent)
- Remote-PHY may allow much higher D3.1 modulation schemes to be used

US Baseline Capacity

- US at 64-QAM, 6.4 MHz = ~ **27** Mbps usable
 - IPv6 may necessitate a-ugs mod profile changes
- 8-ch US bonding available, but limited spectrum
 - North America 5-42 MHz with most stopping at 40 MHz
 - $8 \times 6.4 = 51.2$ MHz needed
 - Euro-DOCSIS US uses 5-65 MHz
 - **48 dBmV** max output/ch unless Extended Pwr in CM implemented!
- D3.0 spec'ed to 85 MHz (not all CMs support it)
- D3.1; 6 MHz block at 256 or 1024-QAM = **36** to **45** Mbps
- Case could be made for 5-85 MHz leading to 80 MHz of D3.1 @ 256-QAM, which is ~ 470 Mbps
- Remote-PHY may allow higher D3.1 modulation schemes

Downstream Speed Affected By:

- MAP & DOCSIS overhead
- Frame size or MTU affects usable rate
 - VPN, wireless, ...
- Modem
 - Config file, CPU (PPS), ack suppression, & Ethernet
- Transport layer
 - TCP or UDP
 - US speeds & TCP windowing effect from delay/latency
- Max DS Burst (Powerboost) - perception is reality
- Computer OS and Windows® stack

Upstream Speed Affected By:

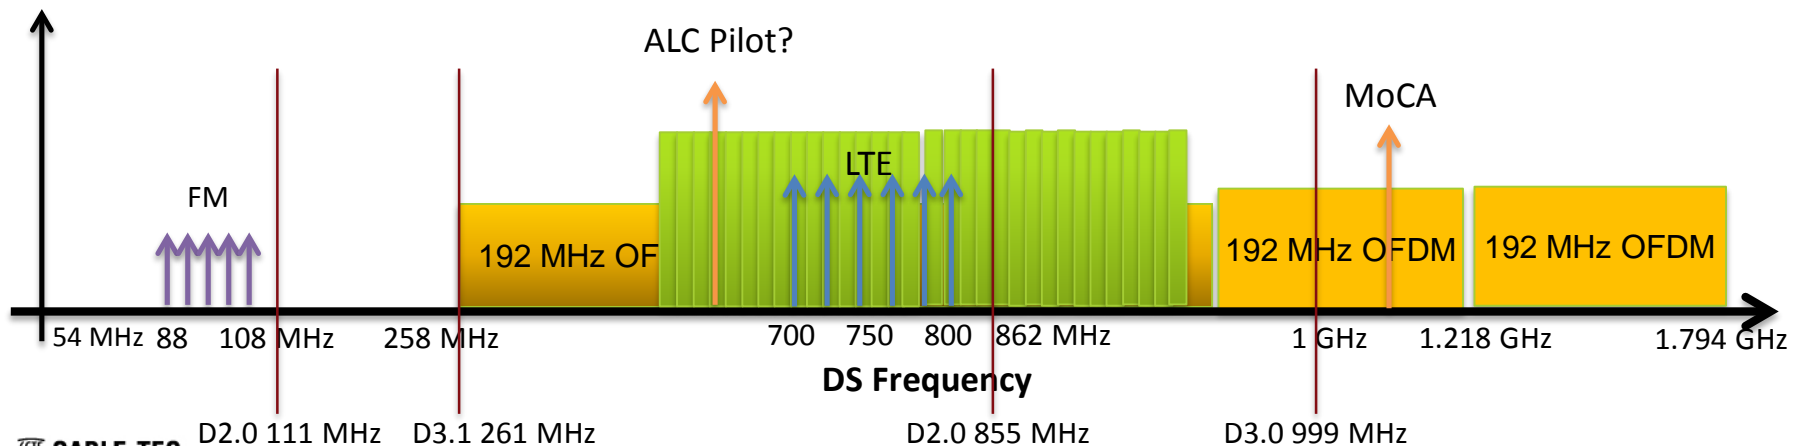
- Rate limiting and/or congestion
- Dropped frames and/or older CMs
- DOCSIS protocol – Req/Grant cycle
 - Downstream MAP Advance
 - DS modulation & interleaving, MAP advance safety, CM time offsets, US bonding, SCDMA, M-CMTS (CIN)
 - Concatenation
 - Max concat & traffic burst settings
 - D3.0 with CCF avoids this
- Fragmentation
- Modulation profiles and channel width
 - Bigger “pipe” = more aggregate speed
 - Packet serialization time (per-CM speed)

D3.1 Capacity Planning Concerns

- How will DS capacity planning work?
 - Block size
 - Subcarriers could be muted
 - Different modulation on subcarriers
 - CMs in same SG with different profiles
 - **Answer**: use highest profile calculation
- How will US capacity planning work?
 - Block size
 - Sharing time with D2.0 & D3.0?
 - Different modulation on many subcarriers
 - CMs in same SG with different profiles
 - **Answer**: rely on minislot allocation (ch utilization)

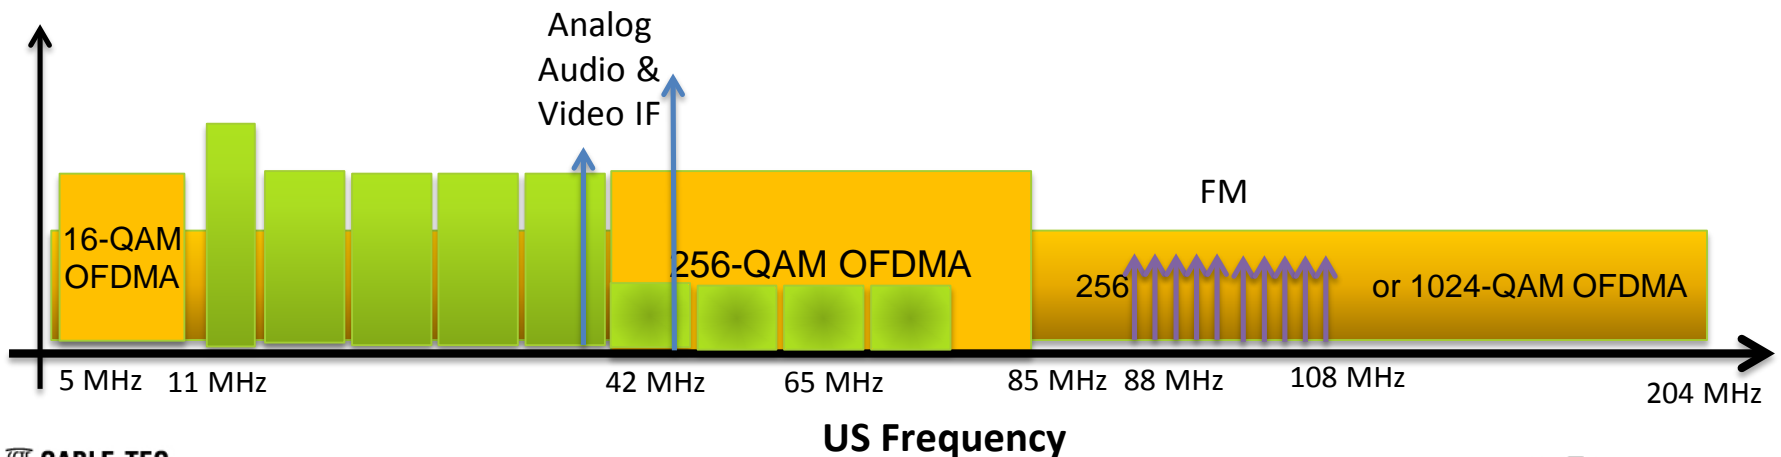
DS Spectrum Planning

- DS spectrum allocation to 1, 1.218, or 1.794 GHz
- DOCSIS 2.0 = 111 to 855 MHz; DOCSIS 3.0 up to 999 MHz
- D3.1, 192 MHz block(s) starts at 111 MHz optional & 261 MHz mandatory
 - **Note:** no DS spectrum sharing between D3.0 & D3.1
- LTE potential issues between 700 to 800 MHz, can't afford to avoid
 - Do not make DSs “primary” in this region and use RBGs
 - Lose D2.0 capacity and can “steer” VoIP to “primary”
 - Gain more speed since “secondary-only” chs
 - Utilize CM FBC to “see” ingress



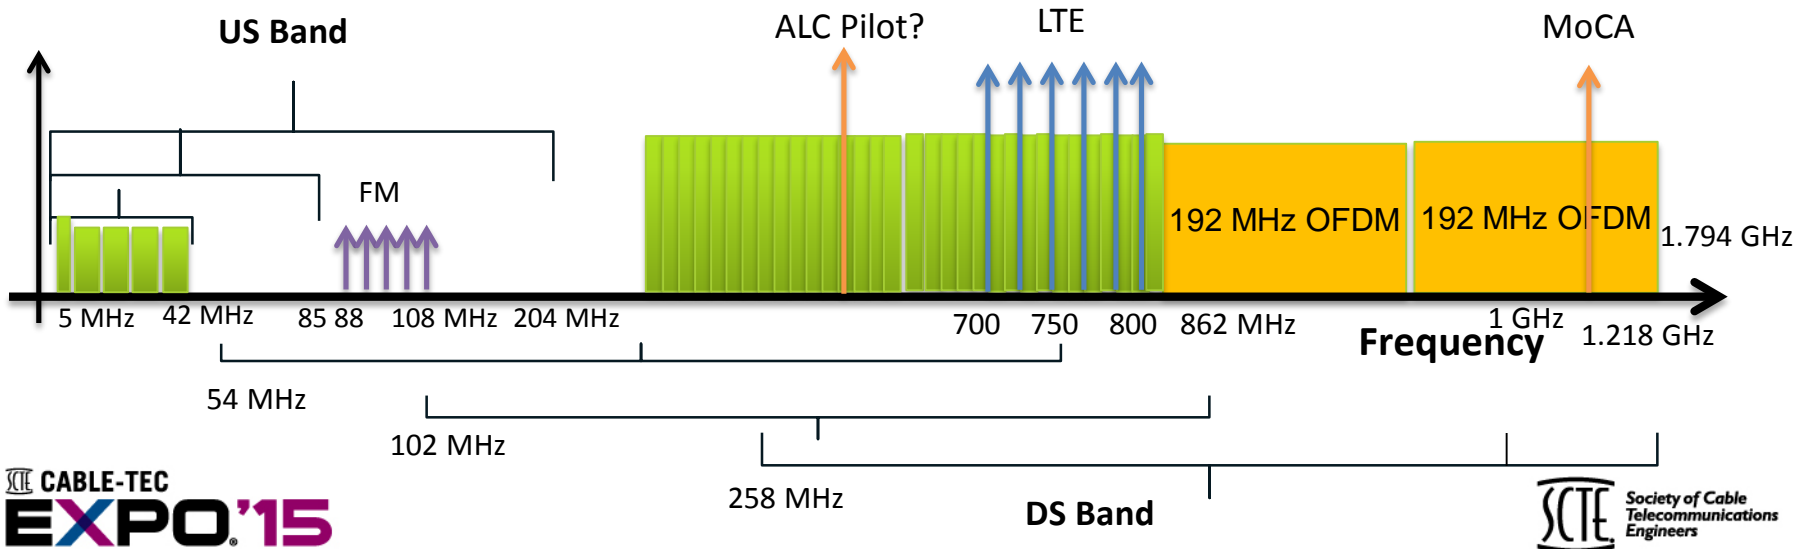
US Spectrum Planning

- D3.0 and 3.1 US spectrum sharing to 42 MHz
 - Synchronize D3.0 & D3.1 traffic to avoid scheduling on top of each other
 - What about licensing of time-shared spectrum?
- **Note:** SCDMA dropped & lowest ch width now 1.6 MHz for D3.0
- US spectrum allocation to 85 or 204 MHz
 - D2.0/3.0, 4-ch to 42/65 MHz; D3.0, 8-ch to 85 MHz
 - D3.1 OFDMA block size (6 MHz?)
- Don't forget about ADI



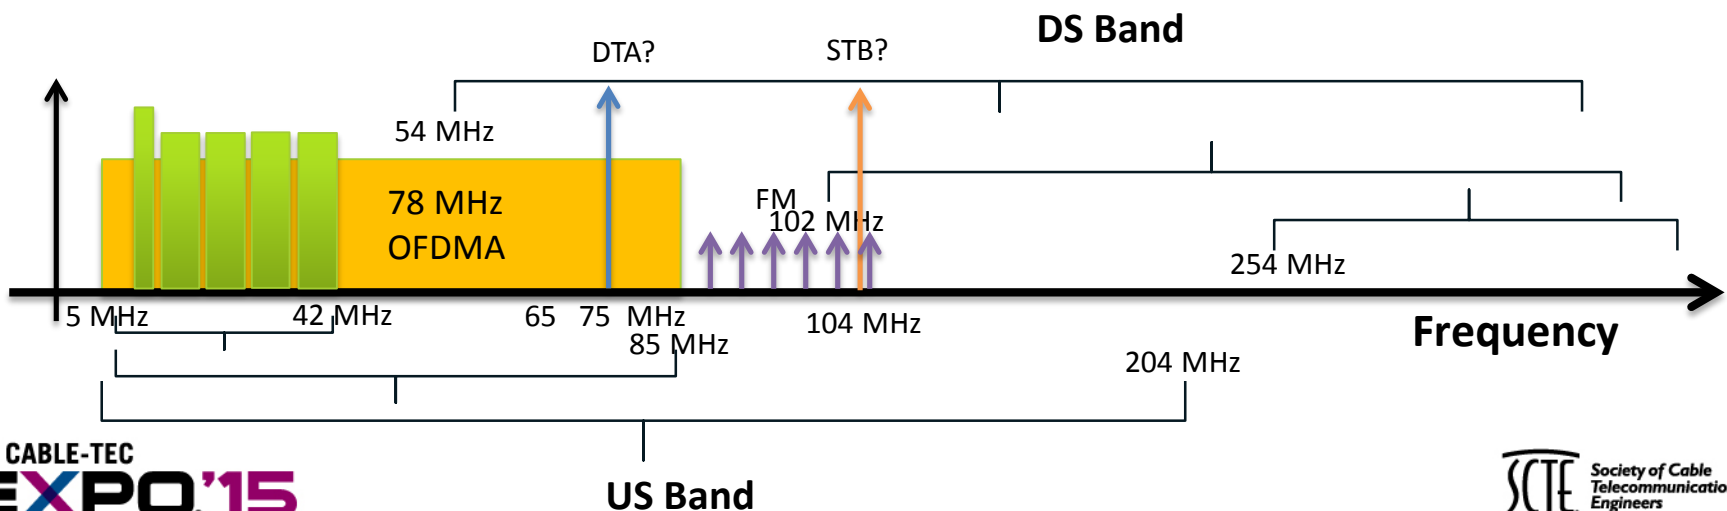
D3.1 DS Spectrum Thoughts

- DS out to 1.218 GHz for actives/passives with higher options (1.794 GHz)
- Upgrade to 1.218 GHz
 - Coax typically supports well past 1 GHz
 - Amplifiers using GaN instead of GaAs support module upgrades without changing housings (re-spacing)
 - Taps may need to be replaced (**good time to evaluate non-linear designs**)
 - MoCA interference (1150 MHz?)
- Higher freqs could create problems: actives, passives, coax attenuation
- US decision affects low-end of DS and what about DS ALC pilot



D3.1 US Spectrum Thoughts

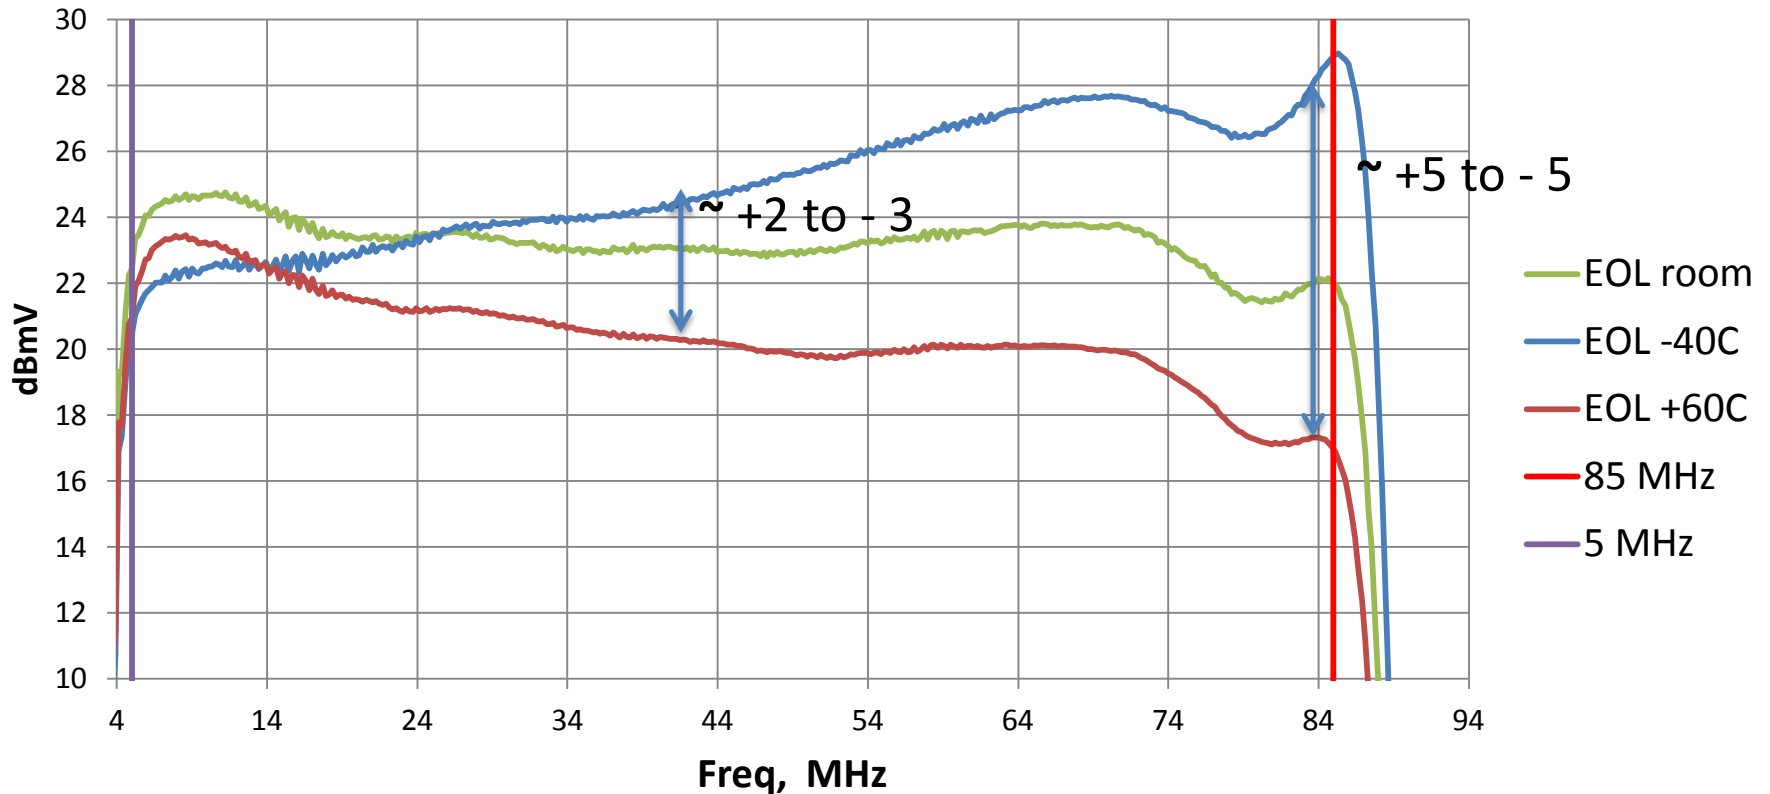
- 204/254 split forces STB upgrade
 - Could generate more features and services with IP STB
 - What about US AGC and pilots?
 - Maybe only N+0, Remote-PHY, or RFoG systems
- Higher options “eat” into DS spectrum
 - What about DTA hard-set at 75 MHz?
- Many may settle on 85 MHz (**removable duplex filters?**)
 - FM-band potential interference
 - Some legacy MPEG-2 STBs have DS addressability ch at 104 MHz
 - 65 to 85 MHz upgrade may not make sense



D3.1 US Spectrum Thoughts (cont)

- Increasing US spectrum creates much more coax attenuation from CM to CMTS
 - US Tx level issues (D2.0 = 54 dBmV; D3.0, 4-ch = 51 dBmV, 8-ch 48 dBmV)
- No US ALC/AGC
 - Relying on CM/CMTS long-loop-level control
 - Typical 6 dB swing @42 MHz (annual thermal fluctuations)
- Potential fixes (no coax 😊)
 - US thermal EQs to help stabilize negative fluctuations
 - D3.0 with Extended Pwr ECN = 3 - 6 dB more US Tx
 - D3.1 is ~ 2 dB more power per equivalent D3.0

US Level Swing over Temperature



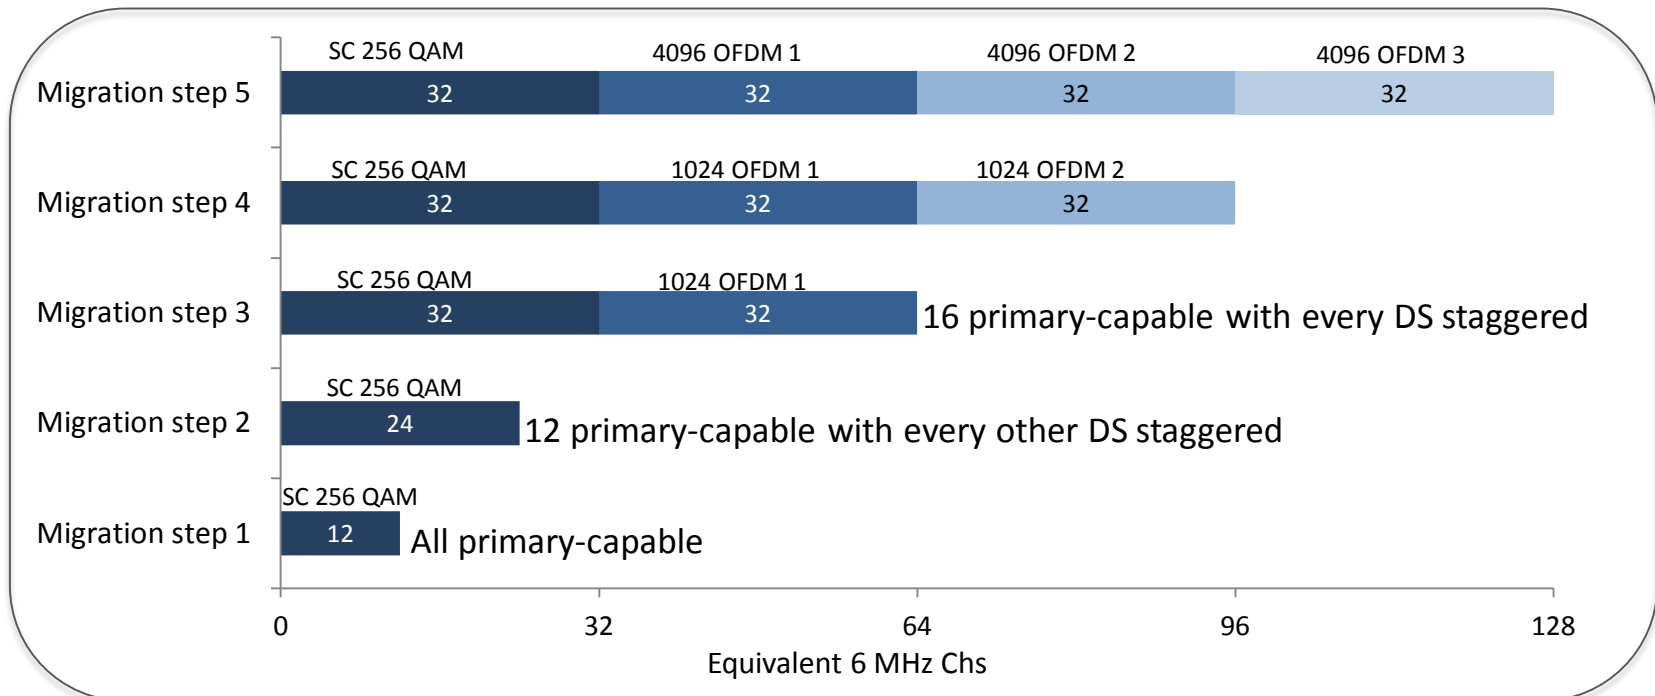
- Potential +5 to -5 dB swing @ 85 MHz
 - 5 amp cascade with ~ 7000' of .5" coax
- CM US Tx reserve & CMTS range config will help on + side
- Minus side creates higher noise floor, hence thermal EQ idea

Five Potential Migration Steps & Ideas

1. 12 DSs: all primary-capable
 - Accommodate DSG, D1.x, D2.0, & D3.0 CMs
2. 24 DSs: 12 primary-capable & all D1.x removed
3. 32 DSs: 16 primary-capable & D3.1 OFDM block
 - Step 3b: 32 + 32 DSs
4. 85 MHz US x 1.2 GHz DS with more D3.1
5. Remote-PHY: 204 MHz US x 1.2 GHz DS

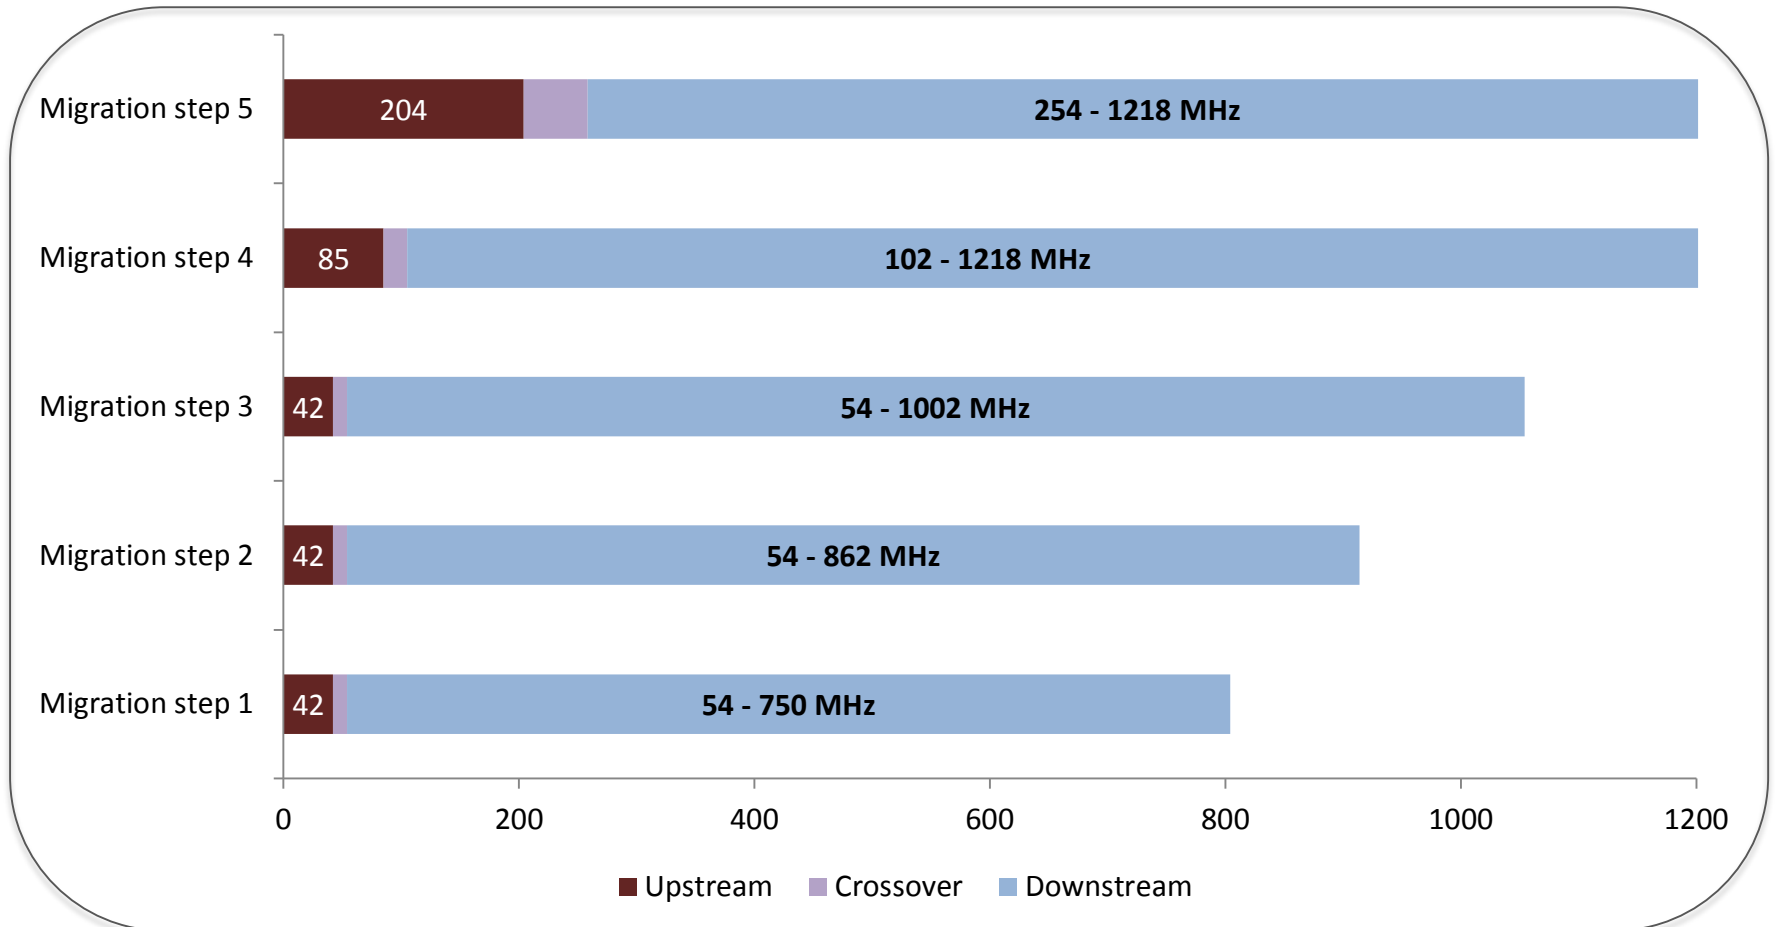
DOCSIS 3.0/3.1 DS Channel Allocation

- Channel allocation listed for all 5 steps
- D3.1 listed in 6 MHz equivalent
- First 3 migration steps can work in 750 MHz, 860 MHz or 1 GHz plant
- 1.2 GHz upgrade when desired, but shown for steps 4 & 5
- Remote-PHY w/ higher modulation in step 5



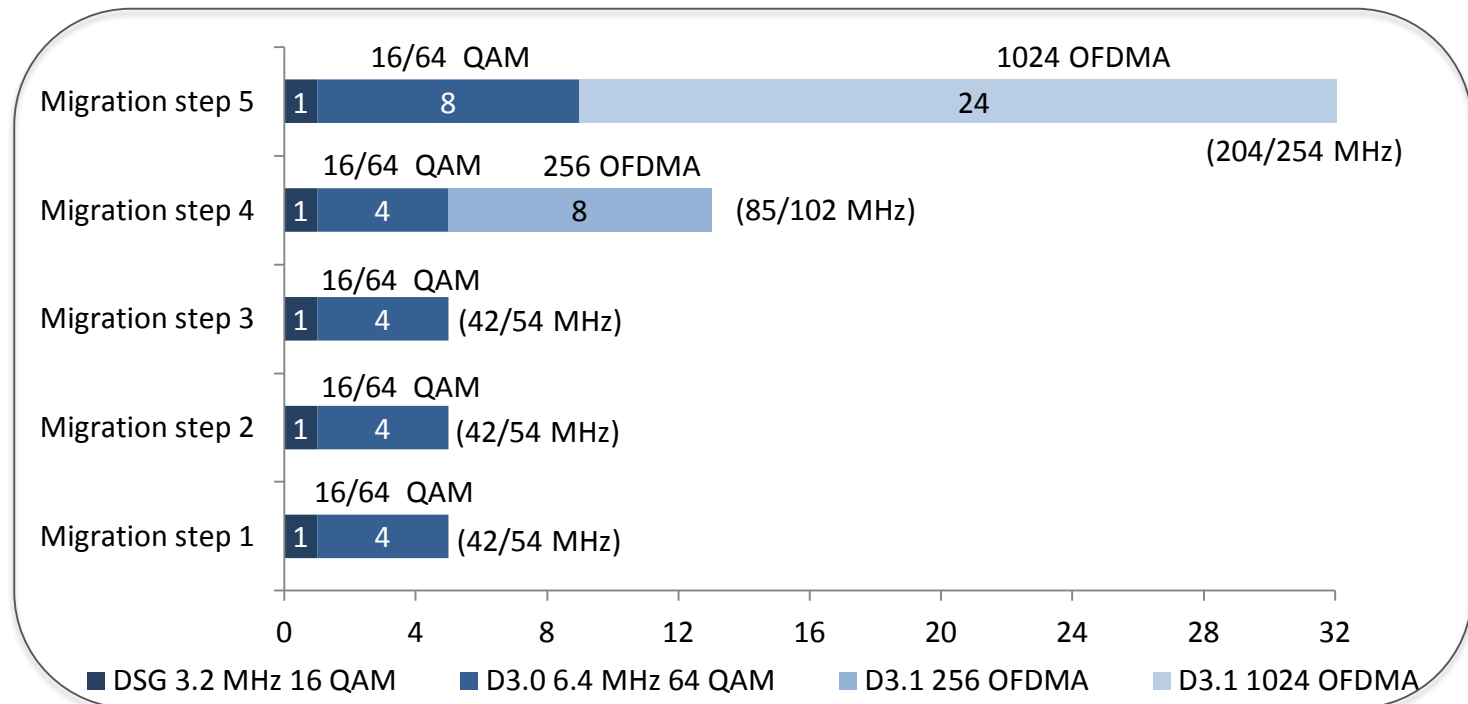
DOCSIS 3.0/3.1 DS Spectrum Allocation

- D3.1 can utilize most spectrum
 - Different modulation in some spectrum, sub-carriers and even specific CMs



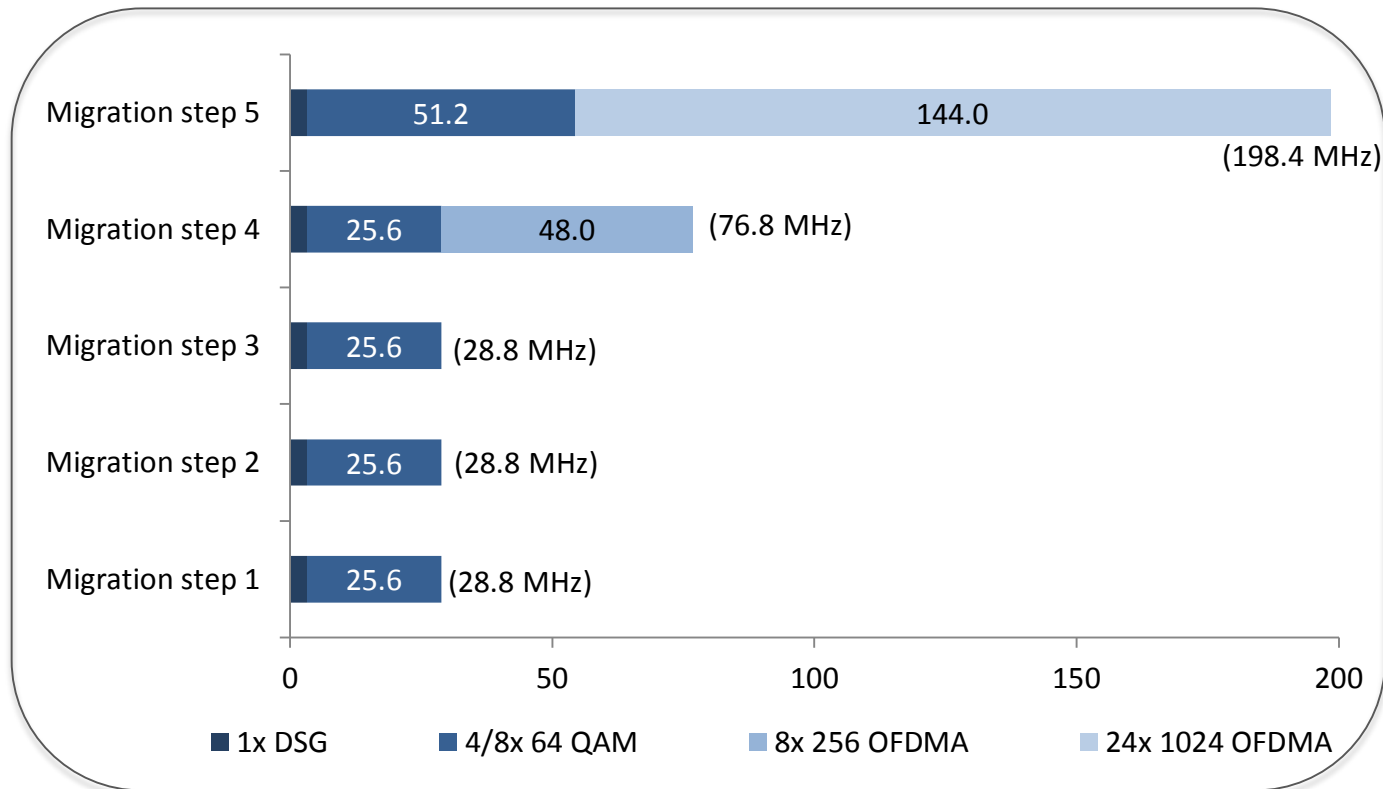
DOCSIS 3.0/3.1 US Channel Allocation

- Channel allocation listed for all 5 steps
- One quarantine channel for DSG & 4 - 8 ATDMA
- D3.1 listed in 6 MHz equivalent
- Diplex filter upgrade at step 4
- Remote-PHY with 204 MHz US for step 5 and higher modulation



DOCSIS 3.0/3.1 US Spectrum Allocation

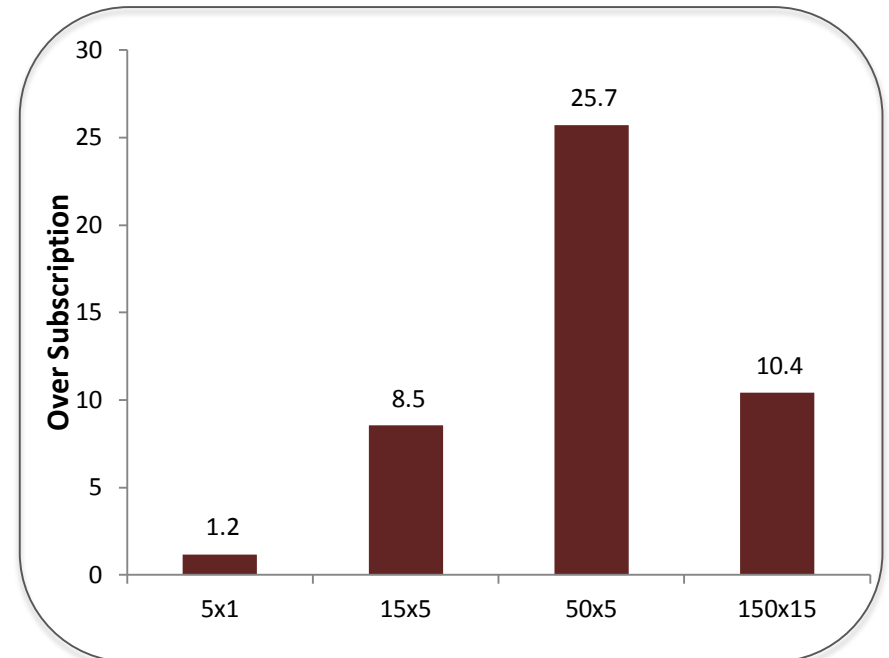
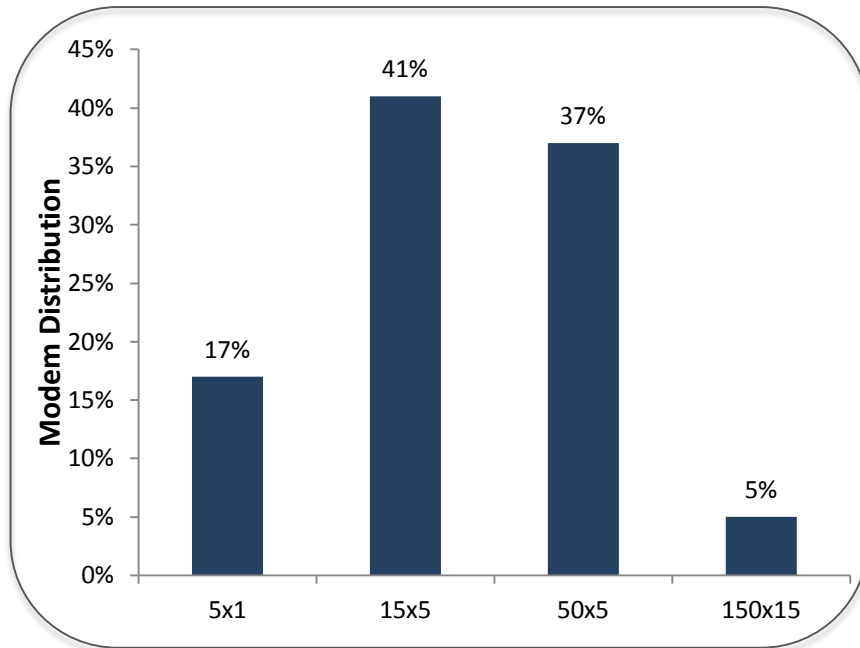
- All steps start DSG @11 MHz band edge
- D3.1 can utilize all spectrum, but maybe not needed
 - Different modulation in some spectrum, sub-carriers and even specific CMs



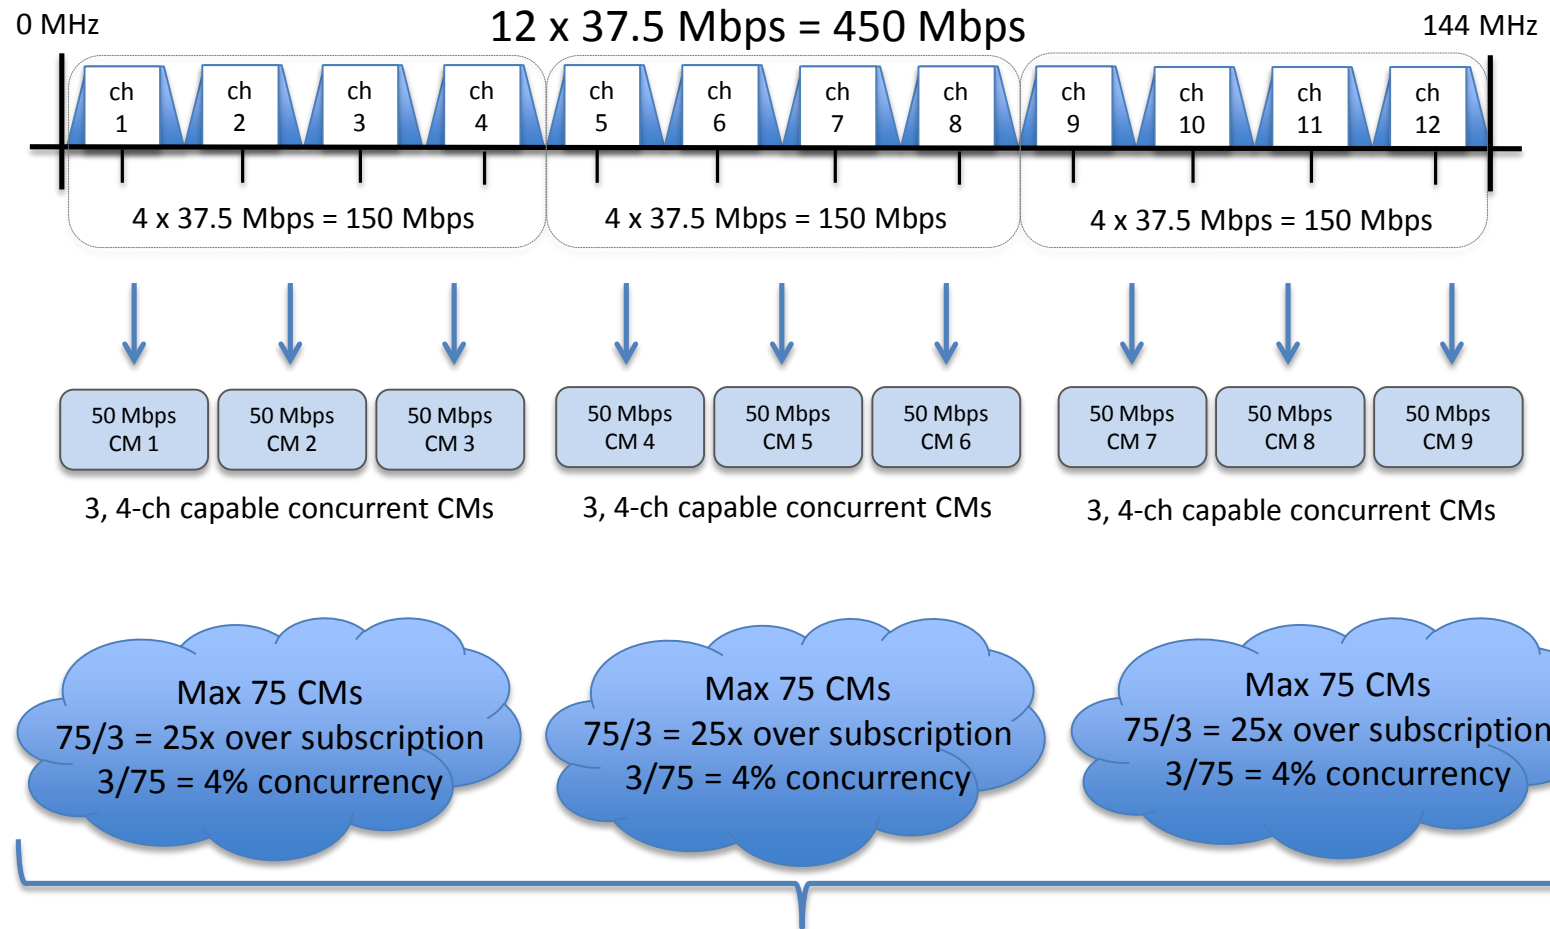
Step #1: 12 DSs

5-42 MHz US & 750 MHz DS

- 600 CMs & 1000 DSG STBs
- DSG, D1.x, 2.0, & 3.0
- Middle tiers make up bulk of subs
- US not limiting factor
- Highest offering is ½ of CM capability
- D1.x OS could easily be greater



25X Over Subscription Estimate

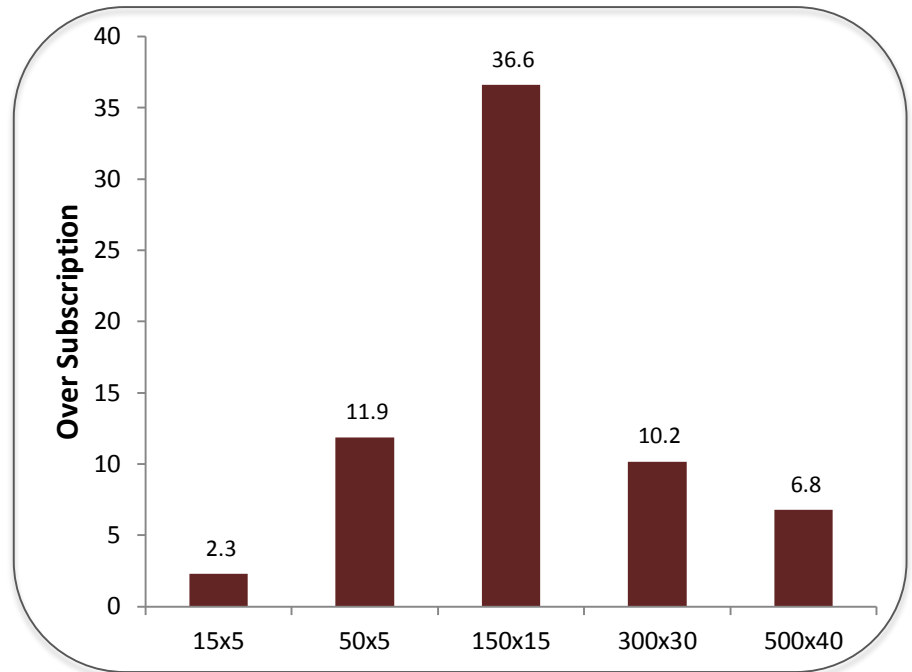
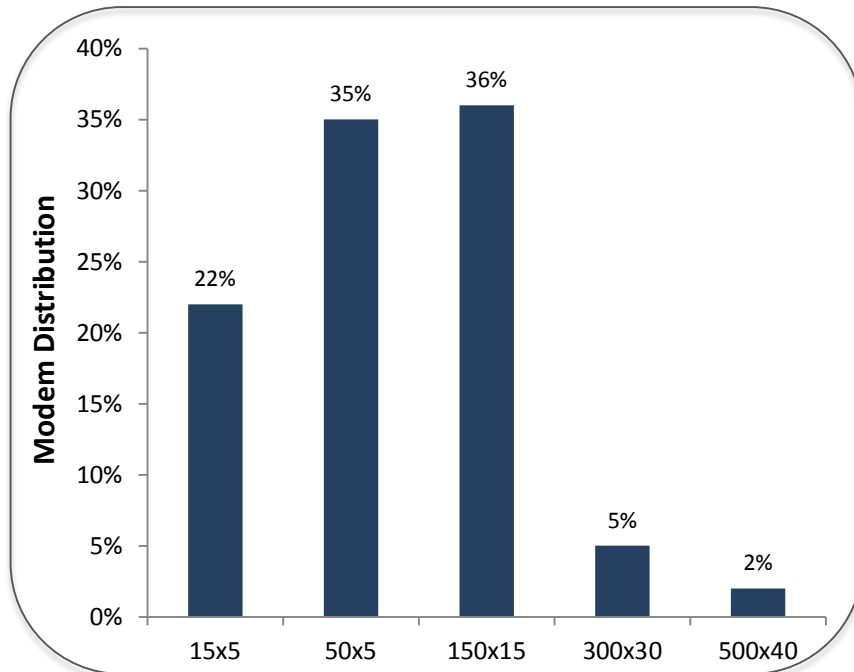


225 CMs each at 50 Mbps
 $11.25 \text{ Gbps} / 450 \text{ Mbps} = 25x \text{ OS}$

Step #2: 24 DSs

Extended DS to 862 MHz

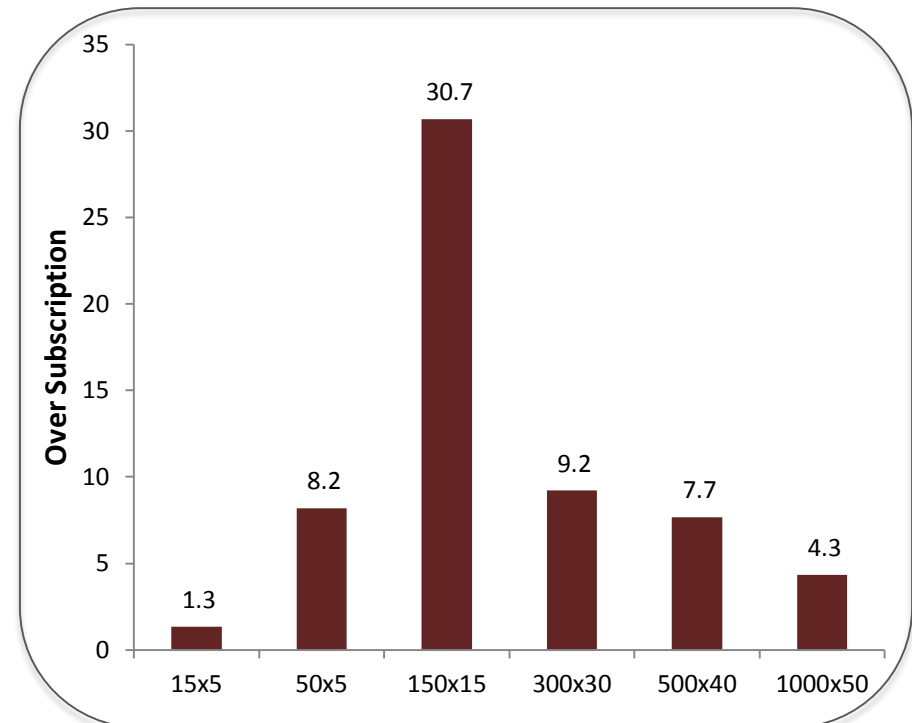
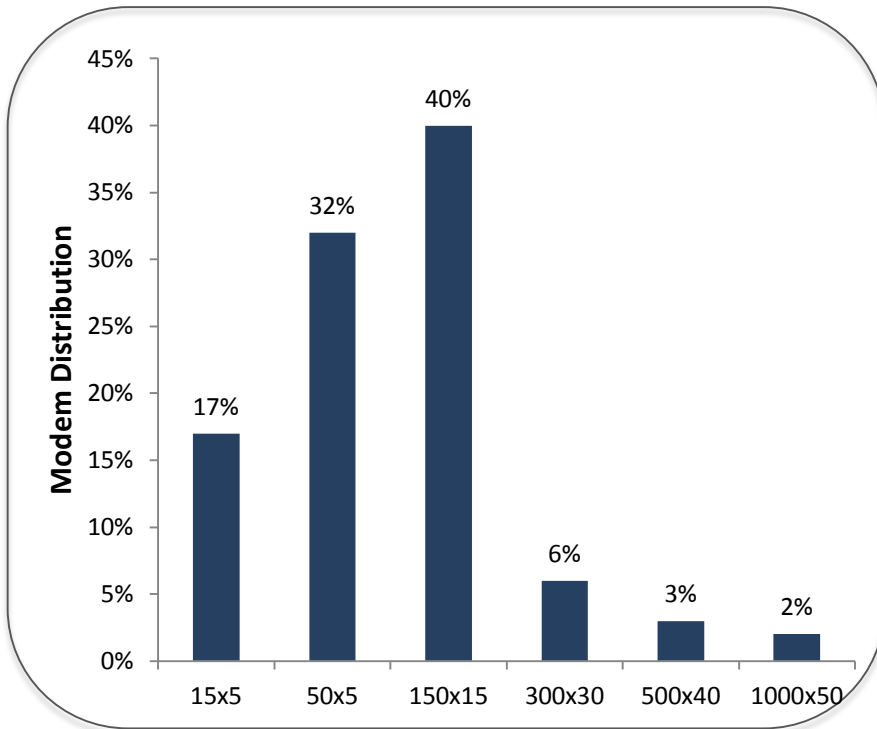
- D1.x CMs removed
 - Higher tiers offered
 - DS:US ratio of 10:1 slightly backed off
 - Subs migrate up in tiers
- D2.0 = 15x5
 - D3.0, 4x4 = 50x5
 - D3.0, 8x4 = 150x15
 - D3.0, 16x4 = 300x30
 - D3.0, 24x4 = 500x40



Step #3: 32 DSs & D3.1 OFDM Block

Advanced DS to 1 GHz

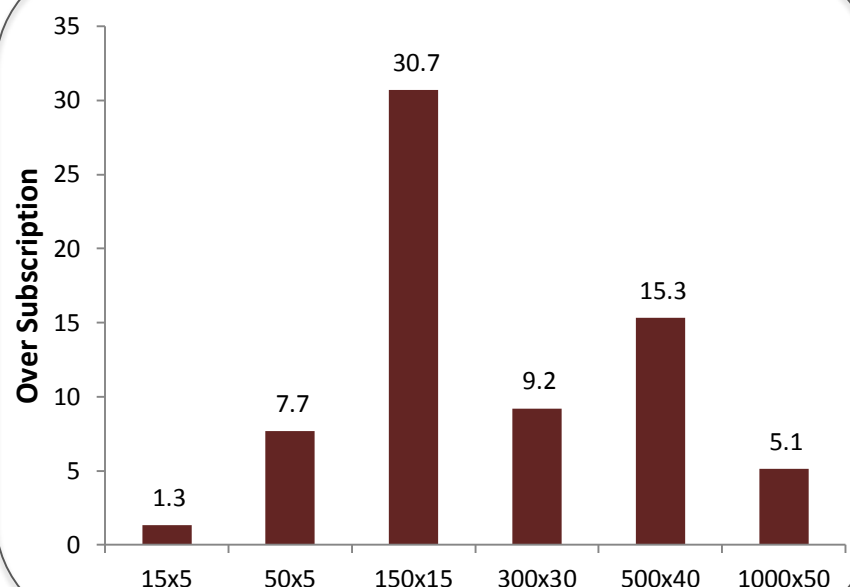
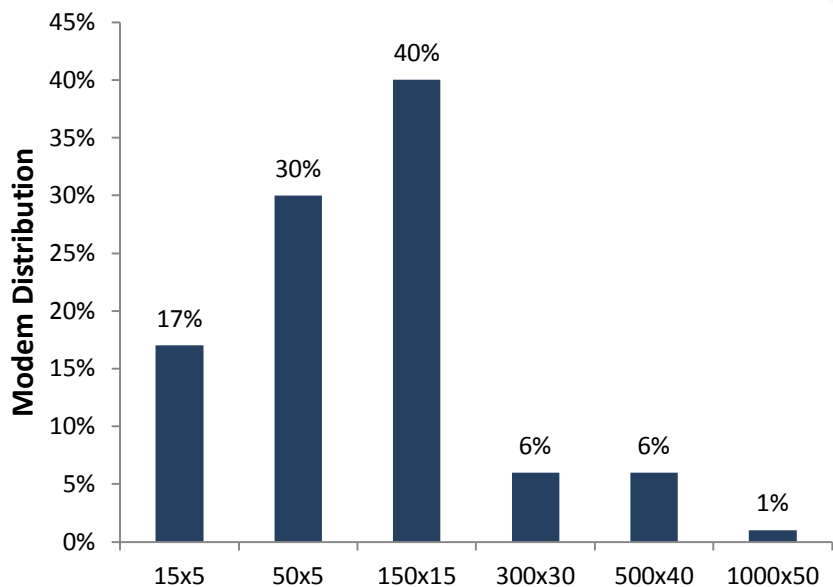
- 1 Gbps tier offered
- D3.0, 32x4 = 500x40
- D3.1 = 1000x50
 - D3.1 & D3.0 chs bonded
- All these speeds could be offered in previous steps
 - More spectrum needed may = more analog/digital reclamation
- Could deploy pre-D3.1 CMs



Step #3b: 32 DSs & 32 DSs

Stick with D3.0?

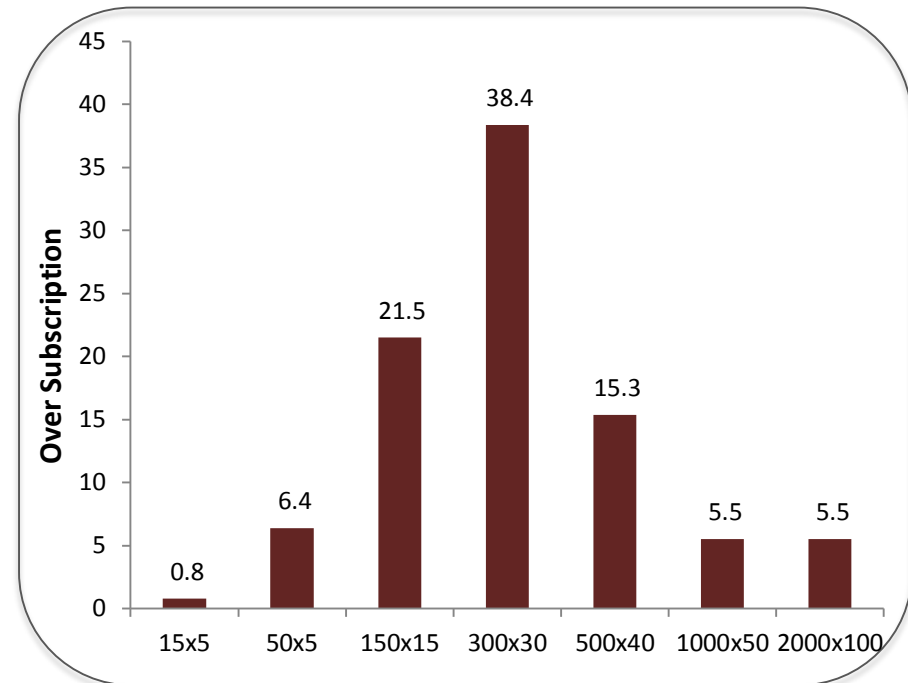
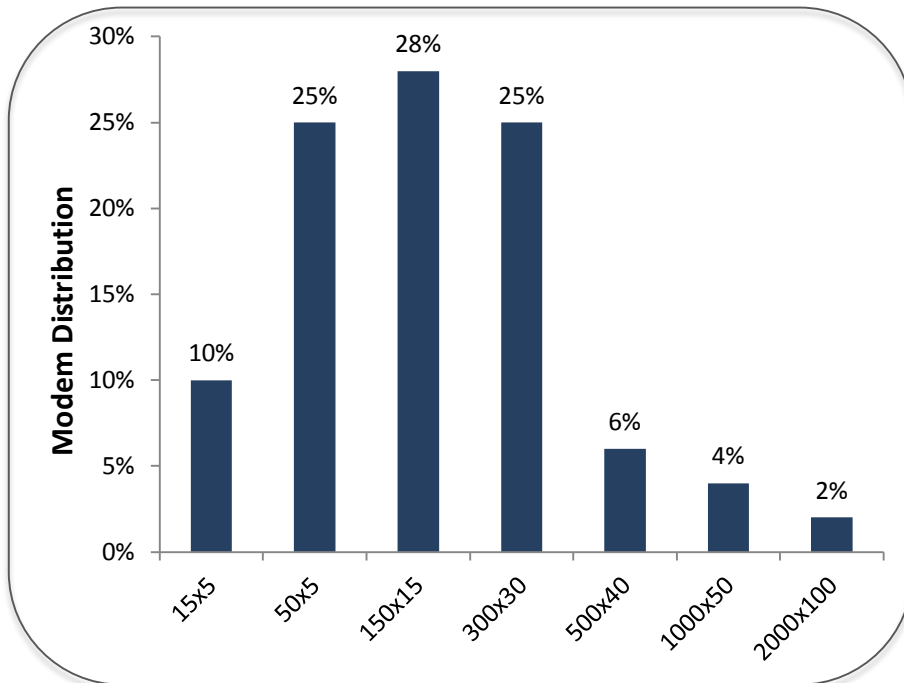
- 32-DS bonding narrowcast
 - D1.x, D2.0, D3.0 < 500 Mbps tiers
 - Extra 32-ch narrowcast or overlay for .5 to 1 Gbps customers only
 - Software-based or phy layer changes?
 - Amount of SGs/CMTS shrink?
- Mainly for peak-speed offering
 - Same spectrum allocation as D3.1
 - No D3.1 CPE to wait on or pay for
 - ~ 1.2 Gbps pipe for annex B
 - ~ 1.6 Gbps for annex A
 - 33% more spectrum



Step #4: 32 DSs & D3.1 OFDM Blocks

Goal of 85 MHz US & 1.2 GHz DS

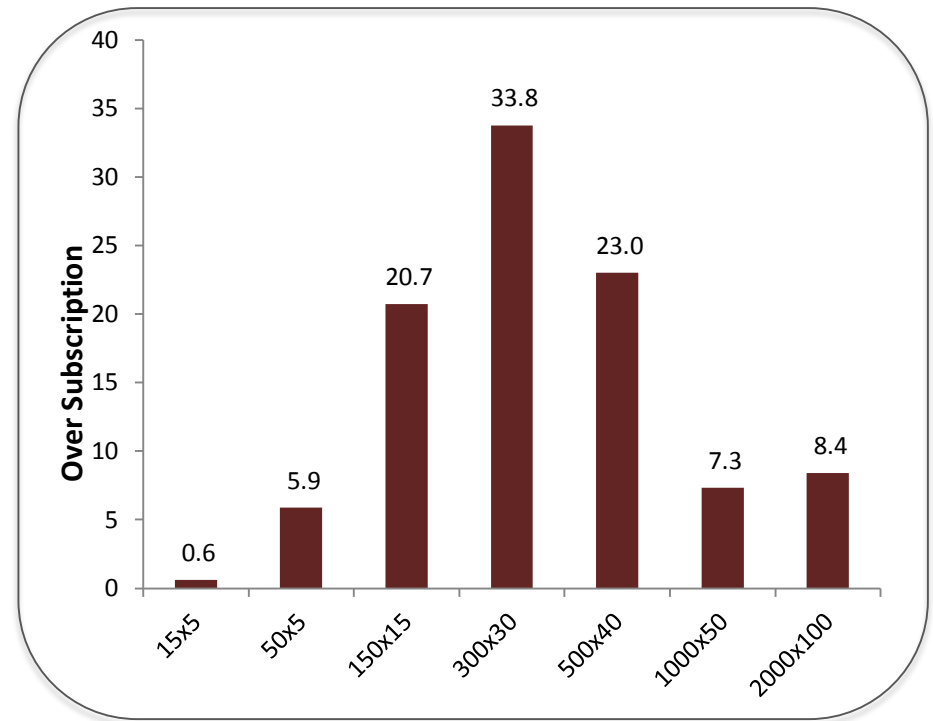
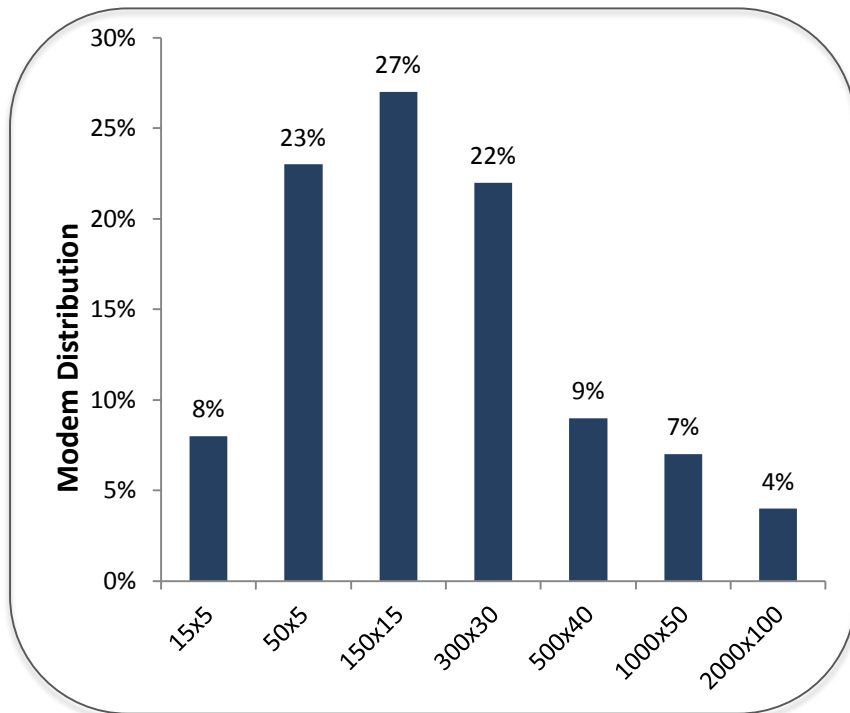
- No DS bonding of D3.1 w/ D3.0
- D3.1 US OFDMA added
 - Eight blocks of 6 MHz OFDMA using 256-QAM
- Four to 8 ATDMA US chs
- Higher US could be offered if time-sharing spectrum
 - D3.1 w/ 13, 6 MHz chs = ~ 470 Mbps



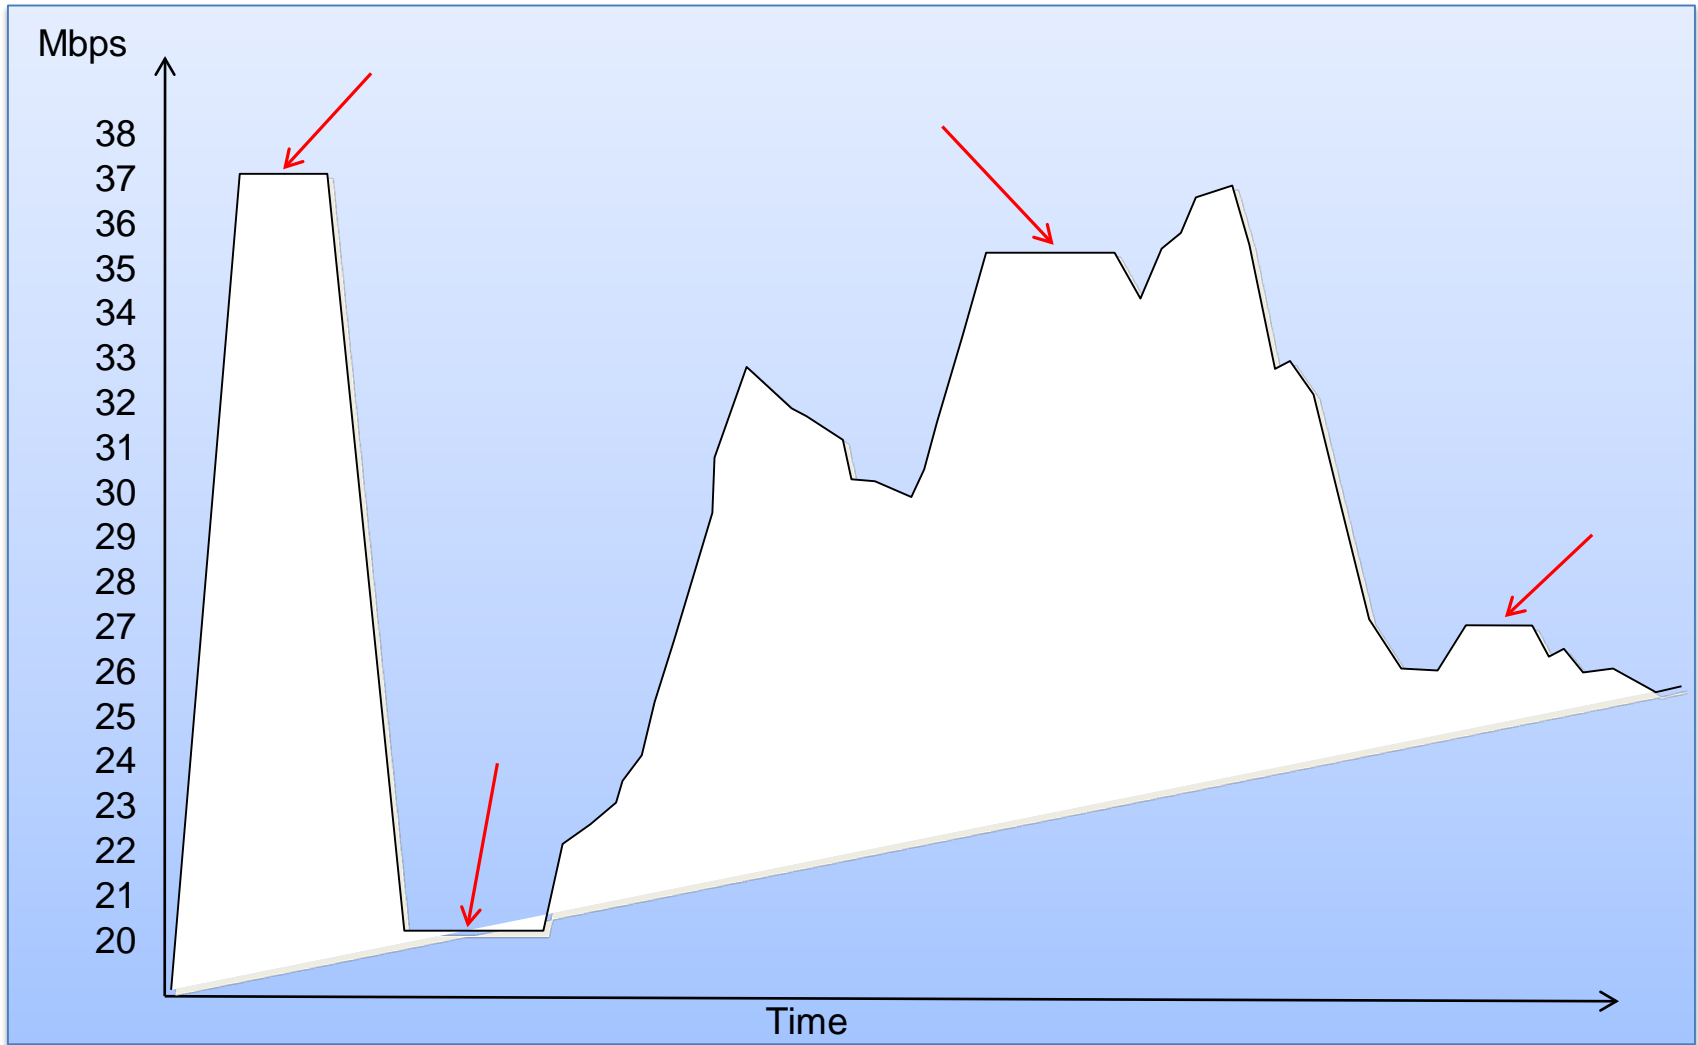
Step #5: 32 DSs & Multiple D3.1 Blocks

Remote-PHY

- Longer distance (2000 km!)
- MER close to 40 dB US & 50 dB DS
- DS OFDM blocks using **4096-QAM**
- US OFDMA blocks using **1024-QAM**
- Could offer symmetrical .5 x .5 Gbps service here as well
 - 1x1 Gbps could be limited by US
 - Aggregate US ~ 1.485 Gbps



What Does This Bandwidth Graph Represent?



JD's 10 Rules for Capacity Planning

1. Aggregate speed = 2 x highest offering
2. Configured speed = 1.1 x Marketed speed
 - Layer 3 vs layer 2 reporting
 - Minimum 1.05 multiplication factor
3. DS:US = 10:1
 - Account for US acks from DS TCP (ex: OTT video)
 - 20:1 minimum initially
 - Possibly as low as 30:1
4. Powerboost for median tiers
 - 5, 10, 15 Mbps
 - US Powerboost possible as well
5. Powerboost for high tiers to bypass rule 2
 - Be sure peak rate TLV used

JD's 10 Rules (cont)

6. Concurrency = oversubscription
 - Aggregate = 2x = 10-20:1
 - Aggregate = 4x = 25-50:1
 - Aggregate = 10x = 50-100:1
7. SG size = 1:1 or 1 DS x 2 US connectors
 - HHP per SG depends on DSs/SG
8. Base math on 768B frame size
 - Most traffic 64 and 1518 B frames
9. Monitor actual utilization
 - Powerboost can skew decision for node splits
 - QoE required – D3.0 CM FTP feature
10. Control “abusers” and prioritize some flows
 - DPI, byte counting with STM/IPDR, cloning,
 - Higher priority for gaming, video, call signaling, CM registration

10 Points to Ponder

1. Many speed sites report at layer 3 of OSI model
 - Configure cm file for 5-10% higher than marketed
2. No control over actual frame size (64-1518 B)
 - Frame size overhead 18/64 (28%) vs 18/1518 (1.2%)
 - MTU affected by wireless, VPN,
3. Small frames = small DOCSIS pipes
 - Only 35 Mbps when all frames are DS VoIP of 229 B
4. DS TCP requires US acks
 - US pipe could slow down DS speed tests
 - Small US acks make US pipe worth less
 - DOCSIS overhead usually 11 B per frame
 - 10.24 Mbps raw = 9 Mbps usable, but only 7.5 with acks!

10 Points to Ponder (cont)

5. PowerBoost can give perception of greater speed
 - Could cause issues when deciding to do node splits
 - How to control peak rate
6. More frames = more PPS = higher CPU usage
 - At some point CPU in modem could (will) be bottleneck
 - TCP (typically 2 DS per 1 US ack)
7. During congestion, you still want priority for VoIP signaling, maybe video acks, and CM registration
8. Load balancing is good, but what speed tier pushes customer to bonding?
 - Maybe >50% of linerate

10 Points to Ponder (cont)

9. Netflix/Hulu TV are using ABR, which is TCP-based
 - Will cause US traffic in form of acks
 - New CMs may have ack suppression on by default
 - Typical US to DS TCP ratio of ~2%
 - With ack suppression, that can drop below 1%
 - Ack suppression doesn't alleviate CM CPU
 - DS IP video of 3-7 Mbps and may make ack suppression inefficient
10. Many tweaks needed to get per-CM US speeds > 3 Mbps
 - Lots of concatenation leads to fragmentation
 - Fragmentation adds headers
 - Preamble & guardtime added to each fragment
 - D3.0 US bonding can do CCF and keep < 2000 B
 - May not require fragmentation, so less overhead

Checklist

- Highest offering < 1/2 of aggregate
- Use PowerBoost judiciously
 - Use peak rate TLV
- Utilize CMTS features for robustness & “self-healing”
 - Load balance (2.0 & 3.0), RBGs, dynamic modulation, ...
- Manage fair use of network (DPI, shaping, STM/IPDR?)
- Monitor actual traffic load and usage
 - Large rolling time average can mask results
- Monitor modem states and service flows are correct
 - Partial mode, registered-traditional-docsis, etc.
- Use periodic speed test to prove no congestion
 - Newer CMs have built-in FTP functionality

Closing Points

- Understand theoretical rate before testing
 - Math & new designs are beginning steps
- Bigger pipes are more efficient
 - But doesn't necessarily mean less latency
 - May need to “steer” latency/jitter-sensitive services to smaller, subset bonding groups
- Determine when additional capacity necessary & how long to implement
- Different service tiers on different ch-capable CMs could necessitate new LB rules
 - Don't move if traffic < 20% or > 60% of aggregate
- Don't assume DOCSIS is your only bottleneck
 - WAN, TCP, CPU, memory, ..

White Papers, References and Resources

- Understanding Data Throughput in a DOCSIS World
 - http://www.cisco.com/en/US/tech/tk86/tk168/technologies_tech_note09186a0080094545.shtml
- Troubleshooting Slow Performance in Cable Modem Networks
 - http://www.cisco.com/en/US/tech/tk86/tk89/technologies_tech_note09186a00800b123c.shtml
- Upstream FEC Errors and SNR as Ways to Ensure Data Quality and Throughput
 - http://www.cisco.com/en/US/tech/tk86/tk319/technologies_white_paper09186a0080231a71.shtml
- How to Increase Return Path Availability and Throughput
 - http://www.cisco.com/en/US/products/hw/cable/ps2209/products_white_paper09186a00801be4ad.shtml
- Understanding Map Advance
 - http://www.cisco.com/en/US/tech/tk86/tk89/technologies_tech_note09186a00800b48ba.shtml
- Carrier-to-Noise Ratio in Cable Networks
 - http://www.cisco.com/en/US/products/hw/cable/ps2209/products_white_paper0900aecd800fc94c.shtml
- Determining RF or Configuration Issues on the CMTS
 - <http://www.cisco.com/c/en/us/support/docs/broadband-cable/radio-frequency-rf-hybrid-fiber-coaxial-hfc/7137-rf.html>
- Cable Modem Provisioning Scenarios Doc
 - http://www.cisco.com/en/US/tech/tk86/tk89/technologies_white_paper09186a008025c169.shtml

SCTE CABLE-TEC
EXPO'15
OCTOBER 13-16 NEW ORLEANS

John J. Downey
jdowney@cisco.com

