DYNAMIC STEERING OF POWER-STARVED CMs, DSG-STBs, & MTAs

Ayham Al-Banna, Tom Cloonan
ARRIS Group, Inc.

Abstract

Power starvation of DOCSIS client devices is a serious problem caused by the insertion of a large number of RF splitters within the subscriber's home coaxial network. This condition introduces significant attenuation in both the US and DS directions, causing the power-starved devices to suffer from degraded performance and service.

Several solutions that address this problem already exist. However, these solutions are either suboptimal or impractical. We propose a novel solution based on dynamic steering of power-starved devices that does not suffer from any of the drawbacks found in the existing solutions.

INTRODUCTION

Since Multiple Service Operators (MSOs) offer various services over the Hybrid-Fiber Coaxial (HFC) network, subscribers tend to have multiple devices in their homes, including Cable Modems (CMs), DOCSIS® Set-top Gateway Set-top Boxes (DSG STBs), Multimedia Terminal Adaptors (MTAs), etc. As subscribers decide to expand their cable access to multiple devices and rooms within their houses, it is common for many new Radio Frequency (RF) splitters to be added within the home coaxial distribution network. This practice can lead undesirable "power starvation" for the devices receiving signals that are passed through these many RF splitters. Power starvation of devices presents challenging problems for MSOs as they strive to offer good Quality of Experience (QoE) service to their subscribers.

Multiple existing solutions to address the above problem are listed in this paper. While these solutions can overcome the problem of power starvation, some of them suffer from serious limitations that make them expensive, suboptimal, or impractical. We propose a novel solution, which does not suffer from any of the limitations present in the existing solutions.

WHAT IS POWER STARVATION?

In this section, the power starvation condition is defined and described. Let us start with a normal and operational scenario where all subscribers' devices (CMs, DSG-STBs, or MTAs) are placed behind a small number of RF splitters as shown in Fig. 1.

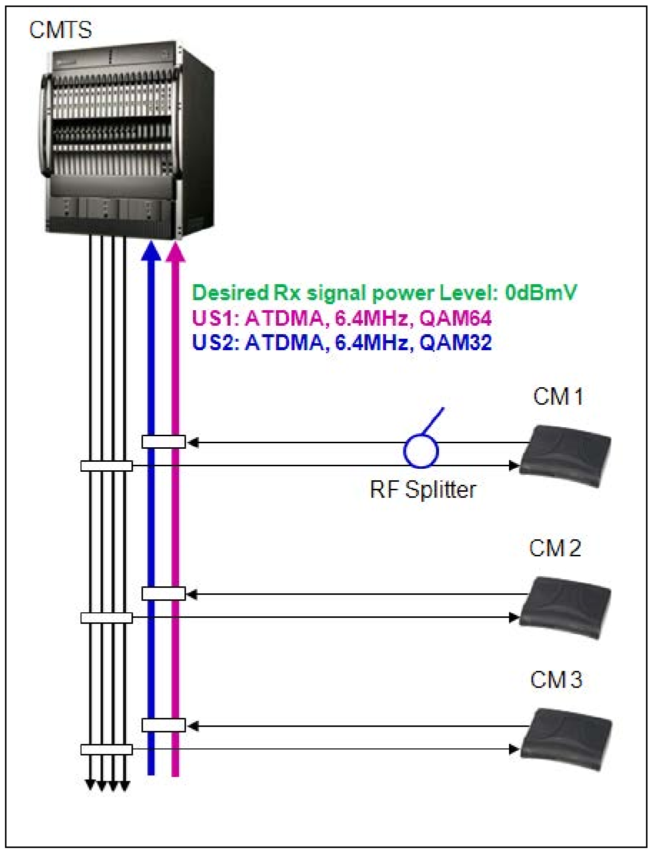

Figure 1. An operational Scenario, where 3 CMs are connected to a CMTS via 4 DS channels and 2 US channels.

All CMs are using Upstream Channel 1 (US1). There is a single RF splitter in the house 1 and no splitters in houses 2 and 3. The CMTS is configured with a receive signal power level of 0 dBmV.

In Fig. 1, we use CMs for illustration purposes while keeping in mind that the discussion also applies to DSG-STBs and MTAs. Herein, we assume that the distance between the CMs and the Cable Modem Termination System (CMTS) is small such that the receive signal power level is acceptable for all CMs when the number of RF splitters installed within each house is small. To simplify the discussion, we also assume that the houses are in close proximity to each other such that the propagation loss between these different houses is negligible.

Figure 1 shows that there is a single RF splitter inside house 1, while no splitters are introduced in the houses of the 2nd and 3rd subscribers. Observe that RF splitters not only introduce loss in the Downstream (DS) direction, but also in the Upstream (US) direction. In Fig. 1, we assume that there are two US channels configured on the CMTS and all CMs are using a single US channel. The two US channels are configured in the following fashion:

US1: ATDMA channel with 6.4 MHz Bandwidth and Quadrature Amplitude Modulation (QAM) 64

US2: ATDMA channel with 6.4 MHz Bandwidth and QAM 32.

All CMs are assumed to be using the US channel labeled US1.

During the ranging process, the CMTS instructs all CMs to adjust their Transmit signal power levels such that their signals arrive at the CMTS at the desired Receive signal power level, which is configured on the CMTS (0 dBmV in this example). This is shown in Fig. 2(a).

In Fig. 2(a), different CMs have adjusted their Transmit signal power level to compensate for the loss between them and the CMTS. We observe that CM1 is sending at higher signal level than that of CM2 and CM3 to compensate for the extra loss introduced by the RF splitter in house 1. In Fig. 2(b), all CMs are hitting the CMTS at roughly the same Receive signal power level. Assuming that there is no channel distortion, all CMs will have comparable Signal-to-Noise Ratio (SNR) values, as depicted in Fig. 2(c), because: 1) All CMs have comparable Receive signal power levels, and 2) The noise level experienced by all receive signals at the CMTS port is identical. Finally, Fig. 2(d) shows that all CMs have good performance because their operating points are well below the maximum acceptable Packet Error Rate (PER) value.

Next, we consider the more interesting scenario of power starvation. Assume that the subscriber that owns CM2 introduces two RF splitters on the cable before feeding it to CM2. Consequently, the CMTS will ask CM2 to increase its transmit signal power level such that the receive signal power level at the CMTS equals the desired value of 0 dBmV. CM2 responds by increasing its transmit signal power level and hits the CMTS at 0 dBmV. Assuming that the noise level did not change, observe that while the transmit signal power level is higher than the value in the operational scenario for CM2, the receive signal power level and SNR values are still similar to the corresponding values in the operational case. This is shown in Fig. 3, where CM2 is still operational with good service since its operating point is well below the maximum acceptable PER value.

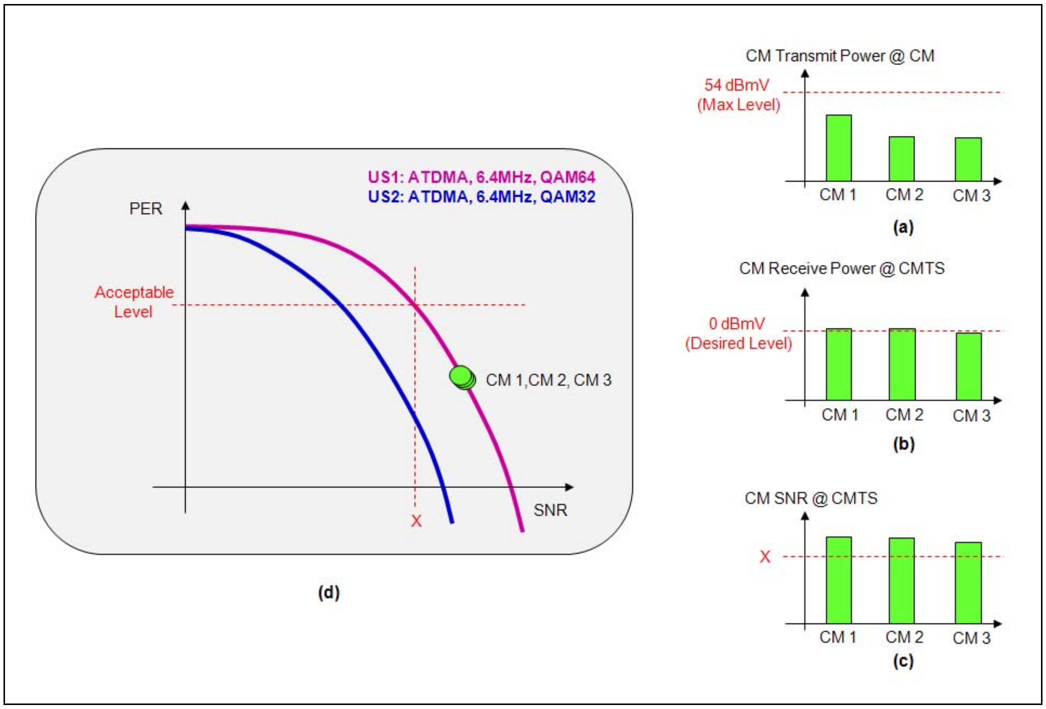

Figure 2. Different curves corresponding to the scenario in Fig. 1. (a) All CMs adjust their transmit signal power level differently to compensate for the attenuation along their way such that their signals hit the CMTS at 0dBmV. While CM2 and CM3 have comparable transmit levels, CM1 has a higher transmit power level to accommodate for the splitter loss. (b) The receive signal power level of all CMs is roughly equal to the desired level of 0 dBmV. (c) All CMs have high SNR values. (d) PER vs. SNR curve showing all CMs have large SNR values and therefore low PER values (below the maximum acceptable limit).

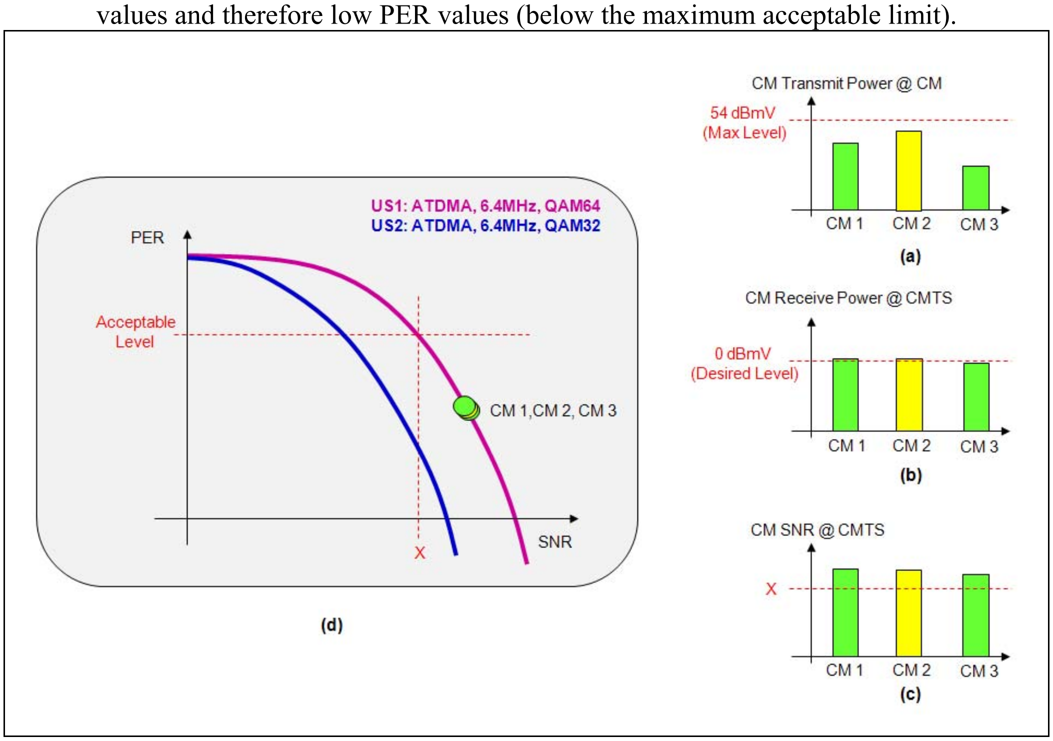

Figure 3. Different curves corresponding to the scenario in Fig. 1 but with 2 RF splitters added along the path of CM2. (a) CM2 increased its transmit signal power level to compensate for the loss introduced by the two splitters. (b) The receive signal power level of all CMs is roughly equal to the desired level of 0 dBmV. (c) All CMs have high SNR values. (d) PER vs. SNR curve showing all CMs have large SNR values and therefore low PER values (below the maximum acceptable limit).

Now, what happens if the subscriber that owns CM2 introduces a third RF splitter along the way, which will further increase the attenuation of RF signals that propagate over the cable. Once the CMTS measures a reduced received signal power level, it runs over the usual behavior of instructing the CM to increase its transmit signal power level such that the receive signal power level at the CMTS equals the desired value of 0 dBmV. As CM2 tries to increase its transmit signal power level to satisfy the CMTS's request, it eventually gets blocked by its own limitation to increase the level because every CM has a maximum permitted limit for the transmit power level (e.g., 58dBmV for QPSK, 54dBmV for QAM32, see Table 6-6 in [1] for more details.) This causes the receive signal power level of CM2 at the CMTS to be less than the desired value of 0 dBmV and therefore CM2 will experience a SNR value that is less than the SNR values of all other CMs on that US channel. Observe from Fig. 4 that all CMs on the US channel obtain good service except for CM2, which has some performance issues as seen from the PER vs. SNR curve. We refer to CM2 as a "power-starved" CM because it increased its transmit power level to its maximum level and yet was not able to hit the CMTS with the desired receive signal power level. Observe that the SNR value that belongs to a power-starved CM can be much lower than the average CM's SNR on that US channel.

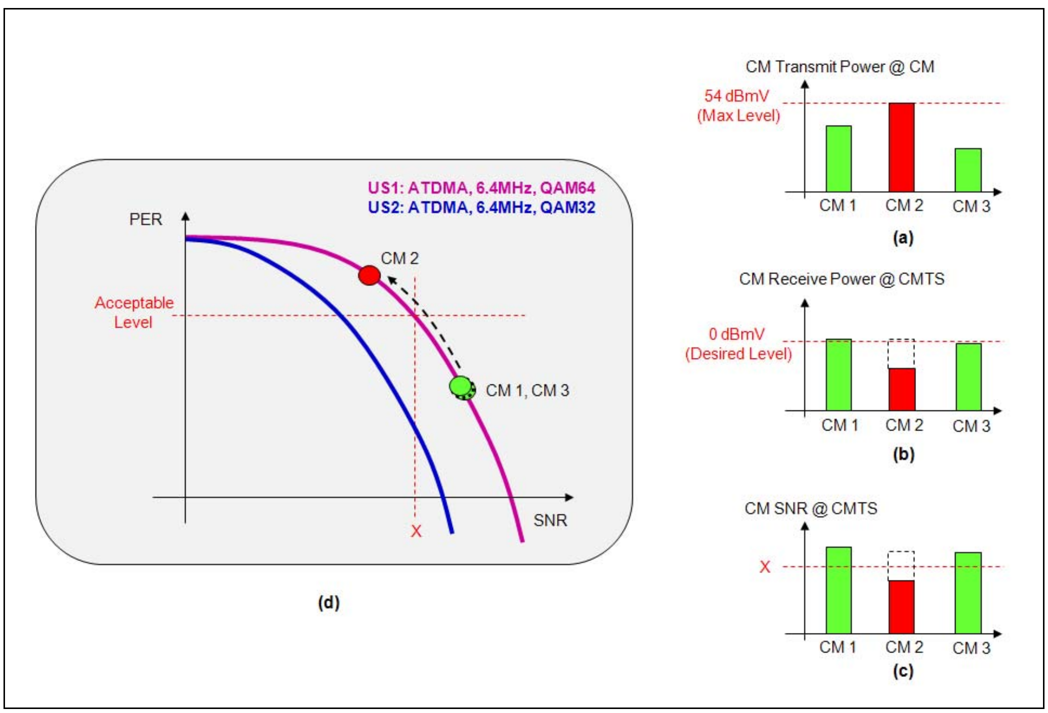

Figure 4. Different curves corresponding to the scenario in Fig. 1 but with 3 RF splitters added along the path of CM2. (a) CM2 increased its transmit signal power level to compensate for the loss introduced by the 3 splitters. However, it got clipped by the maximum limit that CM2 can transmit (54 dBmV in this example.) (b) The receive signal power level of all CMs is roughly equal to the desired level of 0 dBmV except for CM2, where the level is well below the desired value of 0 dBmV. (c) All CMs have high SNR values except for CM2 that has low SNR value. (d) PER vs. SNR curve showing all CMs have large SNR values and therefore low PER values except for CM2 that has low SNR value and hence large PER value (exceeding the maximum acceptable PER threshold).

EXISTING SOLUTIONS AND THEIR LIMITATIONS

There exist multiple solutions for the power starvation problem. However, none of these solutions is optimal as explained in this section. Normally, the MSOs have to choose from one of the several undesirable paths:

1. Do Nothing!

   Since the majority of the CMs on the US channel are receiving good service and only a small fraction of the CMs experience performance issues, one may think that this is acceptable. Unfortunately, this situation is not at all acceptable for the subscribers whose CM is power-starved, and could easily lead to customer churn. Thus, this is an expensive and impractical solution for the MSO!

2. New Modulation profile:

   Another solution to the power starvation problem is to design a new modulation profile (ex: more FEC correction, lower modulation order, narrower channel width, etc.) that can provide adequate PER values even in the presence of the low SNR values of the power-starved CMs on that US channel. This unfortunately yields lower throughputs for all CMs on the channel as shown in Fig. 5. This solution has several disadvantages including: 1) Degraded service (less throughput), for the non-power-starved CMs and 2) lower overall channel bit rates resulting in an upstream plant with lower efficiencies. This solution is not optimal!

3. The SCDMA MSC feature:

   The Synchronous Code Division Multiple Access (SCDMA) Maximum Scheduled Codes (MSC) feature is a good candidate solution for the problem of power starvation. The MSC feature limits the number of active codes used by the power-starved CM while keeping the transmit signal power level unchanged. This

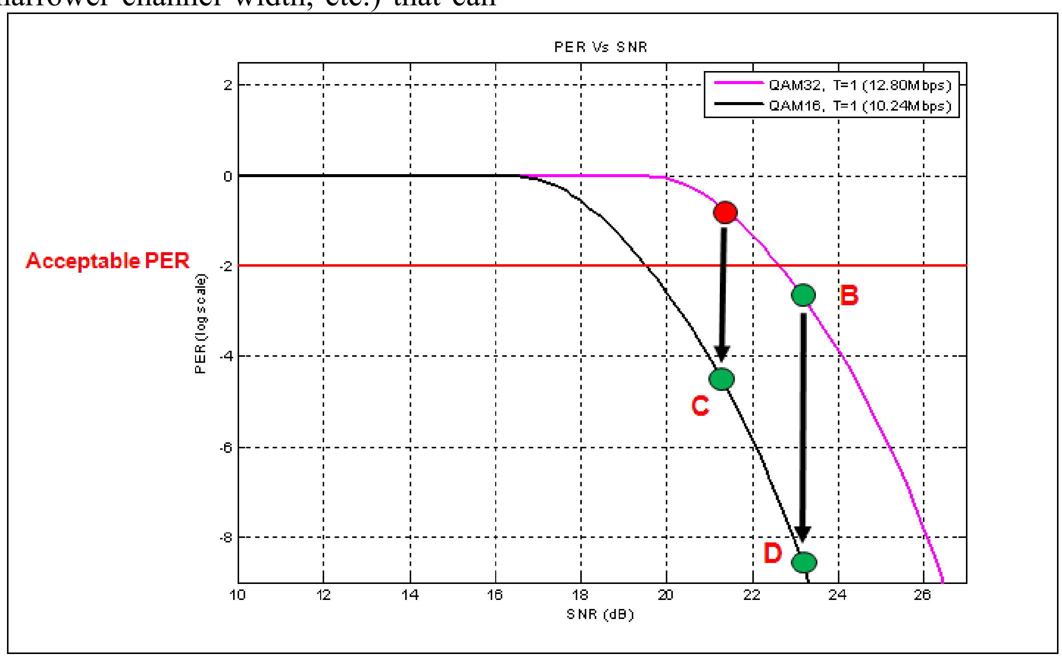

Figure 5. Creating a new modulation profile to accommodate the small SNR values of power-starved CMs is an existing solution. This results in less throughput for all CMs on that US channel. (Channel width is 3.2MHz in this example.)

results in an increased power per code as shown in Fig. 6, which essentially increases the SNR value for that power-starved CM and enables it to obtain good service without changing the modulation profile for the US channel. This solution, however, requires SCDMA-capable devices to be present at both the headend and subscriber's home. Therefore, this solution may not be optimal especially when either the CMTS or the power-starved CMs are SCDMA-incapable.

NOVEL SOLUTION

Power starvation of devices presents a challenging problem for the MSOs because it only affects a few CMs on the US channel. The desired solution needs to be optimal in the sense that it enables the proper operation of power-starved CMs while not degrading the service of non-power-starved CMs or affecting the overall efficiencies of the US channel spectrum.

In this section, we introduce a novel solution for the power starvation problem that does not suffer from any of the disadvantages of the above solutions. The solution is based on an intelligent algorithm that identifies low-SNR Power-Starved CMs and dynamically moves those CMs to channels with modulation profiles that can accommodate the limited SNR values of the power-starved CMs.

In particular, the power-starved devices are first identified using several metrics that can include: Transmit signal power level, receive signal power level, SNR, PER, Modulation Error Ratio (MER), channel noise, DS receive signal level, etc. Once the power-starved device is identified, the system scans through all US channels that are accessible by the power-starved device and identifies an US channel that can provide good performance at the low SNR value of the power-starved device. The power-starved device is then moved (via DOCSIS DCC commands) to the other US channel to obtain good service.

The above algorithm is illustrated in Fig. 7, where the identified power-starved CMs is moved to another US channel that requires smaller SNR values to achieve the same target PER value (Y < X). The target channel can be an US channel with smaller bandwidth (less noise in the passband), a channel with a lower order modulation profile, a channel with a modulation profile with higher Forward Error Correction (FEC) settings, or some/all of the above. Observe in Fig. 7 that the power-starved CM2 is moved from US1 (an ATDMA channel whose bandwidth is 6.4MHz and whose modulation profile is QAM64) to US2 (an ATDMA channel whose bandwidth is 6.4MHz and whose modulation profile is QAM32). Observe that moving CM2 to US2 results in an acceptable PER value (even though the modem is still power-starved, though!).

One important attribute of the proposed solution is the ability to identify power-starved devices and *dynamically* move them to US channels that are suitable for their low SNR values. The dynamic feature of this algorithm can be very beneficial especially when the SNR value of the power-starved device improves. This can happen in several ways, including: 1) when the US noise level decreases, or 2) when the subscriber fixes the problem inside the home by removing some of the previously installed RF splitters. Moving power-starved devices back to their original channels once their SNR values have increased helps to provide better subscriber service and easier optimal network management.

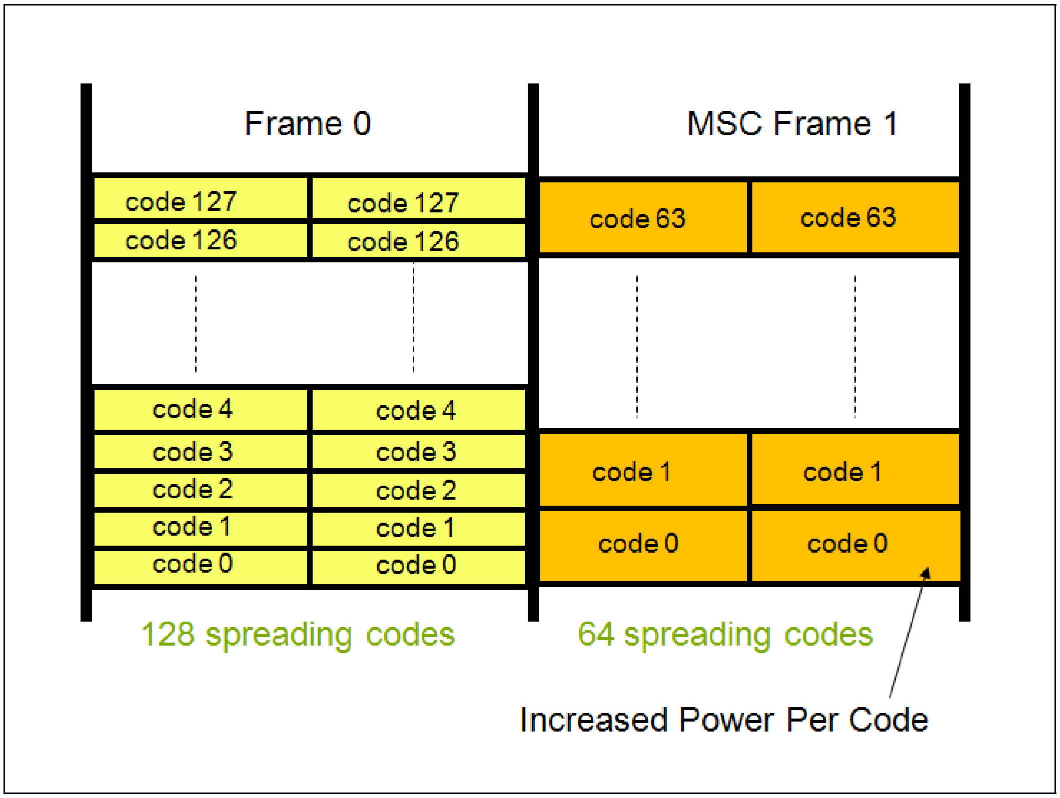

Figure 6. Increasing the code power in SCDMA through reducing the number of active codes is an existing solution for the power starvation problem.

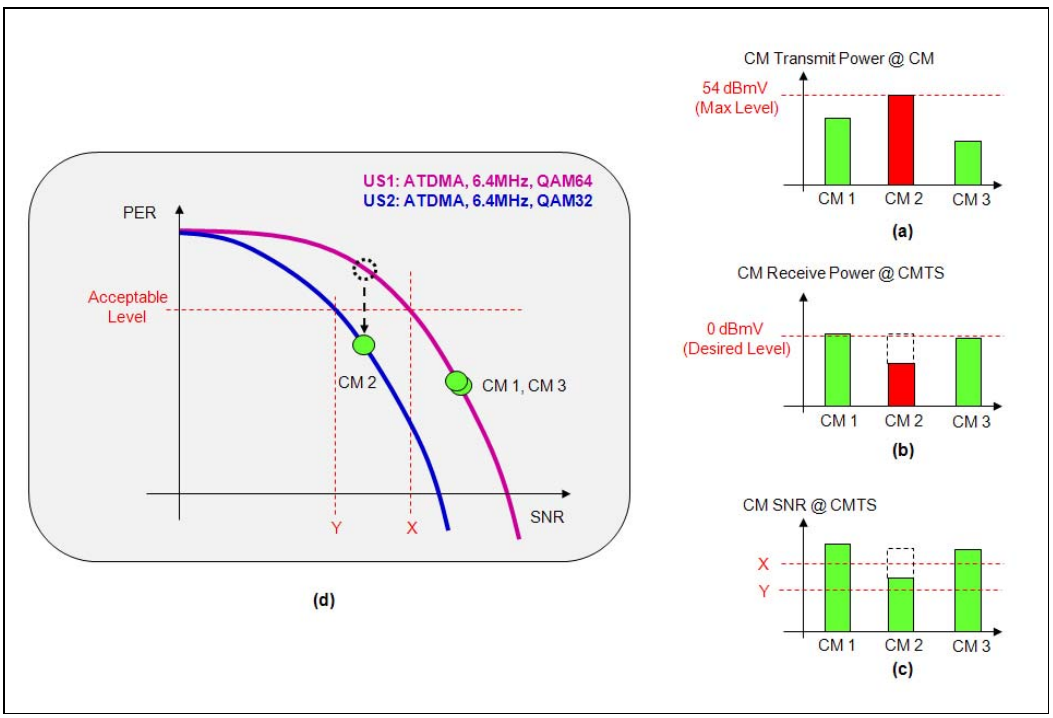

Figure 7. Proposed solution of Dynamic steering power-starved CMs. The power-starved modem, CM2, is moved from US1 (QAM64) to US2 (QAM32) which requires less SNR value to provide for the same target PER. In particular, note that US1 requires SNR=X to provide the desired PER value, while US2 requires SNR=Y (less than X) to provide the same desired PER value. CM2 is still power-starved but its low SNR value is properly accommodated by US2 and therefore no performance issues are encountered.

The network operator may wonder how to locate the additional bandwidth required for a second channel. Most MSOs do not want to consume a large amount of US bandwidth on their plant as they move into a future where the upstream bandwidth will become even more precious. Fortunately, the second US channel can be provided in several ways without wasting bandwidth or greatly affecting the spectrum efficiency. These techniques include:

1. Logical channels. A logical channel is an excellent mechanism to provide the second US channel because it is only used when the power-starved devices on that US channel needs to send data in the US direction. The bandwidth grants for logical channels are assigned dynamically.

2. Narrow channels with robust modulation profiles in the noisy band below 20MHz. This portion of the spectrum is lightly used and can be utilized for supporting the power-starved modems. The same principle also applies to "spectral holes" between other high-speed DOCSIS® channels.

3. Existing DOCSIS®1.0 TDMA channels. Some MSOs already have a TDMA channel present on their network to support legacy devices. This low throughput channel can also be utilized as a home for power-starved devices.

4. MSOs started to deploy DOCSIS® 3.0 US channel bonding which requires multiple US channels to be present. One of these US channels might be adequate to host power-starved devices.

Observe that once the low throughput US channel is identified and selected, it can be efficiently used to host all power-starved devices moved from different US channels within the MAC domain.

Dynamic steering of power-starved devices to other US channels (that are suitable for their low SNR values) is a general solution that has several advantages. These advantages include: 1) the ability to work with all DOCSIS® devices (DOCSIS®1.x, DOCSIS®2.0, DOCSIS®3.0), 2) the ability to work with TDMA/ATDMA/SCDMA channels, 3) the ability to improve the performance of power-starved modems without impacting the performance of non-power-starved modems, and 4) the ability to improve the performance of power-starved modems without impacting the efficiency of the spectrum.

CONCLUSIONS

The problem of power starvation within CMs, DSG-STBs, and MTAs was discussed in this paper. The article illustrates how this problem can occur whenever subscribers introduce many RF splitters into their homes. The existing solutions were listed along with their limitations. In general, the solutions were found to be either expensive, impractical, or suboptimal. A novel solution based on dynamic steering of the power-starved devices to other US channels that can accommodate their lower SNR values was proposed. The paper showed that the offered solution does not suffer from any of the disadvantages experienced by the existing solutions.

REFERENCES

[1] Cable Television Laboratories, Inc., *Data-Over-Cable Service Interface Specifications, DOCSIS® 2.0, Radio Frequency Interface Specification, CM-SP-RFIv2.0-C02-090422,* April 2009.

Authors' Contact Info:

ARRIS Group, Inc.
Address: 2400 Ogden Ave., Suite 180, Lisle, IL 60532, USA
Tel: 630.281.3009
Fax: 630.281.3362
E-mail: Ayham.Al-Banna@arrisi.com
E-mail: Tom.Cloonan@arrisi.com

Biography:

Ayham Al-Banna, Ph.D.: Sr. Systems Architect at ARRIS Group, Inc., Chicago. His research interests include RF Communication, Traffic Management, QoE, and QoS. Ayham has published a book and numerous publications in the area of Wireless and Cable Communications.

Tom Cloonan, Ph.D.: Chief Strategy Officer at ARRIS Group, Inc, Chicago. His areas of research include QoS for DOCSIS® systems, Traffic Management, routing architectures, and TCP Performance. He holds more than 20 U.S. patents and has written more than 80 articles for technical publications and conferences.