The World Is Flat
Capacity Optimization in a Coaxial Network, Constrained by Total RF Power

Karl Moerder PhD, Futurewei Technologies Inc.
Fred Harris PhD, San Diego State University

Abstract

As data traffic continues to grow, the cable industry is reaching the point where downstream capacity in a coaxial network is completely limited by the total integrated RF output power available from a node or amplifier. MSOs today are faced with the problem of increasing the peak network capacity to meet marketing and customer demand. Pushing fiber deeper into the network is one solution, but it is expensive. Working with the existing coaxial network and efficiently utilizing the RF power for maximum capacity can be a cost effective alternative.

Increasing the system bandwidth to 1.2 or 1.8 GHz with a constant quadrature amplitude modulation constellation size and pre-emphasis, while maintaining amplifier spacing based on a 750 or 860 MHz plant design requires a positive slope in the power spectral density (PSD) of as much as 22 dB across the 1.2 GHz band and is far from optimal. Transmitting a signal centered at 1.2 GHz costs over 100 times as much RF power as transmitting the identical signal centered at 100 MHz. It is hard to imagine that this could be sustained as the system bandwidth is increased.

We show that when constrained by the total RF power available, the power spectral density at the output of the amplifier should be nearly flat. The small amount of ripple in the optimal power spectral density is due to the discrete set of QAM constellation sizes.

THE PROBLEM

As data traffic continues to grow, the cable industry is reaching the point where downstream capacity in a coaxial network is completely limited by the total integrated RF output power available from a node or amplifier. MSOs today are faced with the problem of increasing the peak network capacity to meet marketing and customer demand. Pushing fiber deeper into the network is one solution, but it is expensive. Working with the existing coaxial network and efficiently utilizing the RF power for maximum capacity can be a cost effective alternative.

Increasing the system bandwidth to 1.2 or 1.8 GHz with a constant quadrature amplitude modulation (QAM) constellation size and pre-emphasis, while maintaining amplifier spacing based on a 750 or 860 MHz plant design requires a positive slope in the power spectral density (PSD) of as much as 22 dB across the 1.2 GHz band. Coaxial loss increases with frequency, and the pre-emphasis compensates for the increased loss at higher frequencies. Using pre-emphasis with a constant constellation size to maintain a constant signal-to-noise ratio (SNR) over the entire system bandwidth at the receiver is far from optimal.

As an example, transmitting a 6 MHz 256 QAM signal centered at 1.2 GHz costs over 100 times as much RF power as transmitting the identical signal centered at 100 MHz. It is hard to imagine that this could be sustained as the system bandwidth is increased.

Any combination of coaxial loss, amplifier spacing and system bandwidth, results in a negative gain slope across the system bandwidth. Figure A illustrates the typical cable loss and PSD found in a traditional cascade of coaxial amplifiers with both analog and digital television. We express this slope in dB/100 MHz and as a simplification, approximate the slope as constant over the system bandwidth. For any particular slope and system bandwidth, this paper addresses the problem of selecting the QAM constellation size at each frequency to (1) minimize the RF power required to achieve the desired capacity, or similarly to (2) maximize the capacity achievable with a limited amount of RF power.

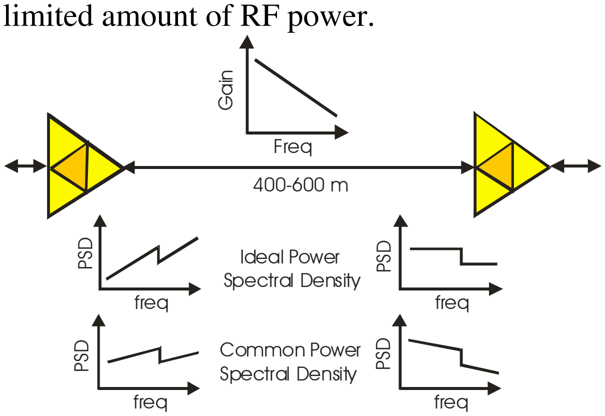

Figure A Traditional Amplifier Cascade

Figure A shows a typical span between amplifiers of 400 to 600 Meters. The total RF power at the output of each amplifier ranges from 73 to 76 dBmV. The frequency dependent loss of the cable is shown directly above the cable. The ideal PSD for a constant QAM constellation size, from an SNR point of view, at the input to each amplifier is flat across the band. This requires pre-emphasis at the amplifier output to offset the slope of the coaxial cable loss. The step drop of about 6 dB in power spectral density between analog television channels at lower frequencies and digital transmission at higher frequencies protects the analog television channels from interference by the digital transmission.

Many older amplifiers were incapable of providing sufficient linearity with a large slope in the PSD across the band. Figure A also shows a second example where the pre-emphasis is split equally between the input and the output of the amplifier. That is, each amplifier provides only half of the pre-emphasis at its output required to offset the gain slope of the coaxial cable. The remaining pre-emphasis is provided at the input of the next amplifier in the cascade. While the amplifier linearity is improved, a SNR penalty is paid.

In many hybrid fiber coaxial (HFC) networks around the world, the elimination of analog television signals is either underway or completed. This eliminates the need for protection of the analog television channels and simplifies the design of the HFC network.

In this paper, we use system capacity to mean the total system throughput in bits per second. We show that the PSD out of the amplifier and into the coaxial cable should be nearly flat over the system bandwidth for maximum system capacity when the total RF power out of each amplifier is limited. So, the PSD at the input to each amplifier should be decreasing with frequency. The total system capacity is maximized with higher-order constellations at lower frequencies and lower-order constellations at higher frequencies. This can be accomplished by filling the downstream bandwidth with DOCSIS 3.1 OFDM channels or by using a larger range of constellation sizes than DOCSIS 3.0 provides. Constellations ranging from 16384 QAM at low frequencies down to 64 QAM at high frequencies, or 4096 QAM at low frequencies down to 16 QAM at high frequencies, maximize capacity for most coaxial networks.

Figures B and C illustrate the PSD into the amplifier, and out of the amplifier and into the coaxial cable for two cases. Figure B shows the PSD for a constant 1024 QAM constellation size from 100 MHz to 1.2 GHz. Figure C shows the PSD for QAM constellation size that decreases from 16384 QAM at 100 MHz to 64 QAM at 1.2 GHz. The QAM constellations are scaled to maintain the same Euclidian Distance before pre-emphasis. Maintaining the same Euclidian distance, maintains nearly the same symbol error-rate performance with respect to SNR or modulation error ratio (MER). The small difference in symbol error-rate performance is due to the changes in the fraction of edge and corner points for different QAM constellation sizes. The capacity in both figures is 8.8 Gb/s, but the total RF power consumption in Figure C is 56% less than that in Figure B.

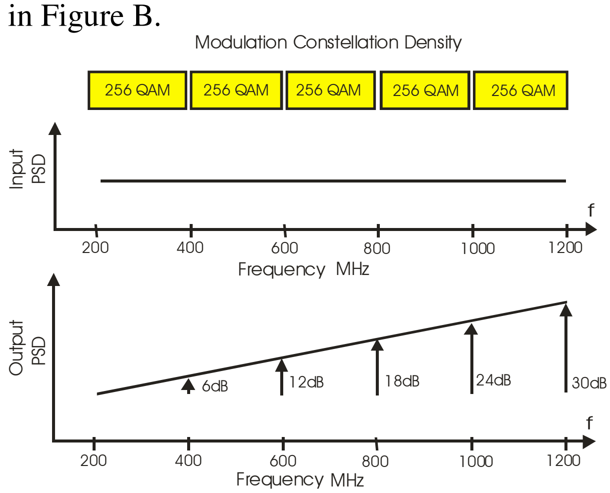

Figure B 1024 QAM System Capacity

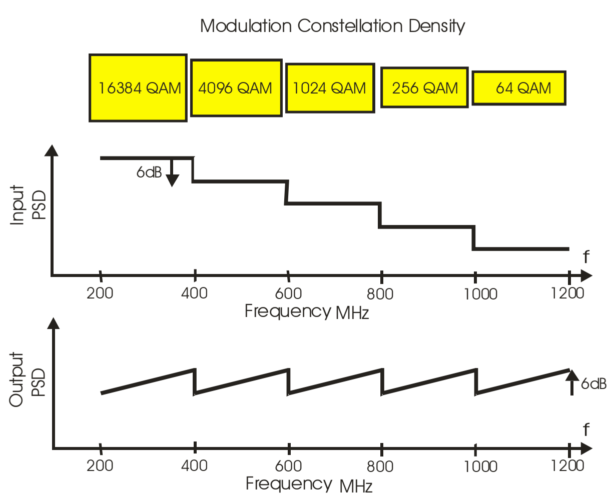

Figure C Optimal System Capacity

While new amplifier designs using state of the art second generation GaN technology can produce more total RF output power than earlier designs, total RF power and the resulting DC power consumption for each amplifier remains an issue for HFC network planning. The distortion in the amplifier depends on more than just the total RF output power. A large positive slope in the PSD increases the amplifier distortion. While distortion is more severe at higher frequencies, the magnitude of the difference in PSD between lower and higher frequencies further stresses the amplifier dynamic range. Distortion components from the higher frequencies can fall back on the lower frequencies that are disadvantaged by 22 dB or more.

A VIEW OF RF POWER AND CAPACITY

We view RF power in relative terms and capacity discretely in bits per second. That is, the cost of each increment of capacity is the RF power required at the output of the amplifier relative to that required for transmitting 4 QAM at the lowest frequency. This simplifies the discussion by eliminating specific link budgets and absolute power levels.

We divide the system bandwidth into equal sized blocks of frequency and consider square QAM constellations, so one unit of capacity is two bits per symbol for a fixed size frequency block. Figure D shows a log plot of the relative power, starting at $\log_{10}(1) = 0$, required to transmit a QAM constellation from 4 QAM to 16384 QAM at any frequency across the system bandwidth. We assume that the code rate, coding gain and any other performance parameters are independent of constellation size. These simplifications can be relaxed as discussed in the conclusion. In Figure E, there are three curves representing the linear RF power required for 4 QAM, 16 QAM and 64 QAM, as a function of frequency. The frequency range and QAM constellation sizes in the figure are limited to make the size of each cell visible. Each cell in the graph with a width of 50 MHz, below the first curve, or from one curve to the next represents one unit of capacity. So, one unit of capacity can be obtained by transmitting 4 QAM in any frequency block. A second unit of capacity can be obtained either by transmitting 4 QAM in a second frequency block, or by changing from 4 QAM to 16 QAM in the first frequency block. A third unit of capacity can be obtained either by transmitting 4 QAM in a third frequency block, or by changing from 16 QAM to 64 QAM in the first frequency block. Allocating capacity can be viewed selecting cells in the figure. Each cell represents the cost (RF power = power spectral density * frequency) to transmit the corresponding unit of capacity. An optimal solution to (1) is a minimum cost selection of cells with the required number of cells. An optimal solution to (2) is the maximum number of cells selected without exceeding the allowed total cost.

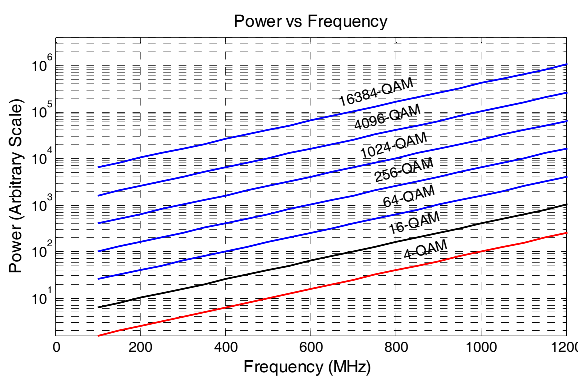

Figure D Power Cost

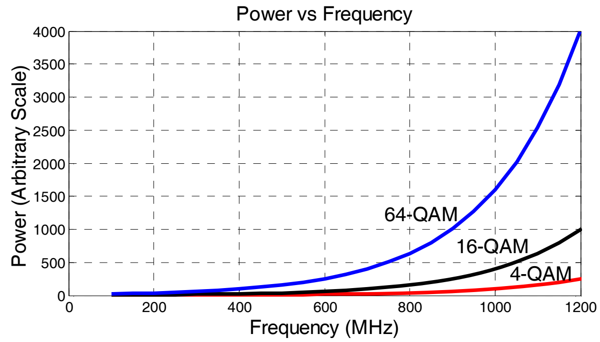

Figure E Power Cost

OPTIMAL CAPACITY AND RF POWER

We now address constructing a solution to (1) and (2) by adding capacity until obtaining the desired outcome. Each unit of capacity has equal value because it provides the same increase in b/s throughput. The cost of each unit of capacity is different. At each step in the process, we add the least expensive increment of capacity.

An optimal solution, S, to (1) is obtained by starting with zero capacity and repeatedly allocating the globally least cost unit of capacity until the desired capacity is obtained.

Suppose that the above solution, S1, were not optimal. That means that there is another solution, S2, with lower total cost than S1. By construction, the cost of every unit of capacity in S1 is less than or equal to the cost of every unit of capacity not in S1. Both S1 and S2 contain the same number of units of capacity, so any other solution to the problem, such as S2, may be obtained by removing some number of units, k, of capacity from S1 and replacing them with k units of capacity from not S1. But the cost of each unit of capacity removed from S1 is less than or equal to the cost of each unit of capacity selected from not S1. So the sum of the costs of the k units of capacity selected from not S1 is greater than or equal to the sum of the costs of the k units of capacity removed from S1. Therefore, the cost of S2 is greater than or equal to the cost of S1. This contradicts the assumption that S2 had lower total cost than S1. Therefore S is an optimal solution to (1).

An optimal solution, S, to (2) is obtained by starting with zero capacity and repeatedly allocating the globally least cost unit of capacity until all available RF power is used.

The proof of this is similar to that of (1).

In order for a solution to (1) or (2) to be feasible, it must satisfy the following constraint, as shown in Figure E: If any cell is allocated, then all of the cells below it must also be allocated.

As an example, it makes no sense to allocate 16 QAM for a particular frequency block without previously allocating 4 QAM for that same frequency block. The cost of 16 QAM is 4 times the cost of 4 QAM in the same frequency block. However our approach would only add the incremental cost, three times the cost of 4 QAM, under the assumption that 4 QAM has previously been allocated.

Fortunately, the incremental cost of 16 QAM is three times the cost of 4 QAM in the same frequency block. So selecting the globally least unit of capacity will always select 4 QAM in any frequency block before 16 QAM in the same frequency block. Therefore the above condition will be satisfied and the solution obtained will always be feasible. Although 4 and 16 QAM are used in this example, the same argument holds true for each increasing QAM constellation size.

THE ALGORITHM

With the above discussion, a simple algorithm for an optimal solution to (1) can now be described.

For any particular frequency block, the cell under the lowest curve is the RF power cost for 4 QAM, the cell between the first and second curves is RF power cost for 16 QAM, and the cell between the second and third curves is the RF power cost for 64 QAM.

Consider the following algorithm.

1. For each frequency block and for the curves for each size QAM constellation, create a list, L, of the height under the first curve, or between consecutive curves for that frequency block.
2. Sort the list, L, from least to most cost. If two or more cells have the same cost, the order is arbitrary.
3. Select the first n cells, where n = (desired capacity) / (capacity of a cell).
4. For each frequency block, we interpret the constellation size as the largest constellation size selected for that block.

This is, in fact, an optimal solution to (1). After sorting, the first cell in L is the globally least cost unit of capacity. Removing the first cell, leaves the second cell as the globally least cost unit of capacity in L. And so on. Therefore this algorithm satisfies the condition stated above for an optimal solution to (1).

CONCLUSIONS

We have shown that when constrained by the total RF power available, the power spectral density at the output of the amplifier should be nearly flat. The small amount of ripple in the optimal power spectral density is due to the discrete set of QAM constellation sizes. Extending the allowed set of constellations to include non-square QAM constellations will reduce the magnitude of the ripple. Care must be taken that we do not violate the feasibility constraint.

Using a fixed size frequency block allows for the simplest algorithm for both optimizing system capacity and demonstrating the algorithm's correctness. However, the fixed size frequency block is only a simplification; arbitrary frequency block sizes are possible.

A flat power spectral density at the output of the amplifier minimizes the total RF power required and thus the DC power required for the amplifier. It also simplifies the design of the amplifier. We mentioned that a large slope in the power spectral density makes the amplifier design more difficult. A flat power spectral density can eliminate the compromise of splitting the pre-emphasis between the input and output of the amplifier.

The effects we have shown in this paper are much less dramatic for a narrower total system bandwidth. It is only as we push the bandwidth to 1.2, 1.8 or even 3 GHz that a new look at total RF power, system bandwidth and system capacity becomes necessary.

RELATIONSHIP TO PREVIOUS WORK

Claude Shannon [2] developed Information Theory during World War II. From an information-theoretic standpoint, he solved the problem of selecting the optimal PSD for communication over an additive, colored Gaussian noise channel in his original paper. Robert Gallager [1] provides a description of the Water-Pouring algorithm for this selection process. Our algorithm provides a discrete calculation for the capacity of a coaxial channel limited by total availableRF power.

In a CableLabs internal white paper, *Shannon's Limits Applied to Cable Networks above 1 GHz*, Tom Williams [3], Gregg White and Alberto Campos suggest using "wide bandwidth transmissions above 1 GHz for best performance when transmit power is limited." Using a larger number of smaller QAM constellation sizes to achieve the same capacity is one way to use wider bandwidth transmissions at higher frequencies.

REFERENCES

[1] R. C. Gallager, Information Theory and Reliable Communications, Wiley, 1968.

[2] C. E. Shannon, "A Mathematical Theory of Communication," Bell System Technical Journal, Volume 27, 1948.

[3] T. Williams, G. White and A. Campos, "Shannon's Limits Applied to Cable Networks above 1 GHz," CableLabs, 2014.