Brownouts in the Brownfields?

A Technical Paper prepared for the Society of Cable Telecommunications Engineers
By

Dr. Robert L Howald
(Formerly) Fellow of the Technical Staff
ARRIS
rob.howald@arrisi.com

Stuart Eastman
Senior Network Architect
ARRIS
stuart.eastman@arrisi.com

Abstract

An important characteristic with both positive and negative consequences to operators is distributed plant powering. On the negative side, powering or outdoor actives costs money, supplies are subject to failure and maintenance, electrical storms can prevent service delivery on effected sections of plant, and time and resources must be spent in engineering this important subsystem for capacity and robustness. However, a key *advantage* of plant powering, and an advantage that HFC networks have over fiber-to-the-home (FTTH) alternatives, is the flexibility it offers in enabling new services. Powering enables MSOs to extend existing services (reach), enhance existing services (segment plant), as well as create entirely new services (Wifi™ Access Points).

Powering efficiency is climbing as a priority for operators, for operational expenditure (Opex) reasons, a desire to proactively address "green" initiatives that connect operators to the interests of the communities in which they operate, and to stay out of the spotlight with respect to any new regulatory burdens that may arise if government institutions become more involved with mandates. Unfortunately, many key business growth objectives tend to be in opposition to this goal. Such objectives take advantage of the powered plant to do more things that generate revenue. Multi-screen delivery, higher speed Internet, converged media experiences, on-the-go access, smart home services, and commercial services, for example, are all focus areas for operators, and many of these service evolutions have implications for the powering subsystem. And, the repercussions are almost exclusively negative to plant power consumption.

Many examples of initiatives can be cited with this characteristic effect:

- Distributed Wifi™ and Wifi™/4G Access Points
- Metro-Ethernet Services
- Distributed PHY Layer Architectures
- Spectrum Migration and Expansion
- Fiber Deeper and FTLA
- EPoC

MSOs are already executing on several of the above strategies and solutions. Others are in planning phases. Some are years away, thus ensuring continuing obstacles to power savings for the foreseeable future. And, of course, some may not happen.

In this discussion, we will assess and quantify how these various strategies and solutions impact the growth of plant powering requirements. We will project at what pace this might occur based on market trends and implementation projections. We can compare this power consumption "CAGR" against increased efficiencies expected from architecture and technology initiatives and even from potential alternative sources over time, and evaluate the bend trajectory of the curve. The discussion will offer key evolution insights that will allow operators to balance the opposing objectives of service growth and power efficiency.

The Now and The Next

Cable operators are engaging in many parallel initiatives aimed at maintaining their leadership as the preferred supplier of choice for consumer video and broadband services. Key themes in the evolution of the network architecture are simplification, new revenue-generating services, and capacity growth. Fortunately, there are many technologies being developed to address these themes, and vendors are building product solutions to enable operators' objectives. Unfortunately, from the standpoint of energy consumption, Opex, and "Green" initiatives, every one of these service, technology, and architecture developments negatively impact plant powering at a time when the opposite is becoming a focus. We can group the architecture evolutions into two broad buckets – things happening now, and things emerging as possible future phases of architecture evolution.

The Now

In the "happening now" category shown in Figure 1 are things like fiber deep…deeper…deepest (perhaps N+0) evolution, plant-based Wifi™ Access Points (APs), and Ethernet Services via point-to-point fiber connectivity, typically for business services.

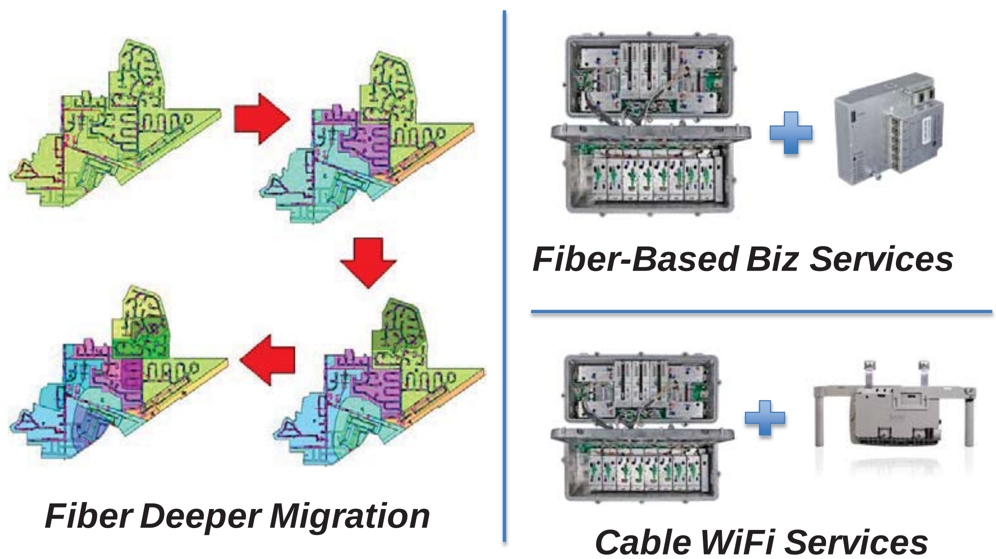

Figure 1 – Power-Hungry Plant Activities Happening Now

In these cases, with generally available solutions on the market, we can calculate with reasonable precision what power consumption penalty is incurred. Amplifier products that convert to nodes are standard products, and Ethernet Switch plug-in modules and Wifi™ APs have known powering requirements. Armed with these "knowns" we can easily calculate the implications to powering if they are deployed in the plant.

Business-as-Usual (BAU) Segmentation

The persistence of aggressive compound annual growth rates of downstream bandwidth has been dealt with by operators with targeted node segmentation. The flexibility of HFC allows this approach to deal with average capacity growth on an as-needed basis. The left side picture in Figure 1 shows a geographical example of the service group size impacts of node segmentation of what begins (upper left, green) as a single node serving area. Over time (clockwise), the service group is segmented into smaller physical regions with fewer homes passed per region. A given amount of spectrum correlates to a given amount of capacity that can be shared, and as this capacity gets shared with fewer subscribers, all subscribers obtain a higher average bandwidth. This is the essence of BAU fiber-deeper segmentation.

Node segmentation itself comes in two flavors – virtual segmentation and physical node splitting. In virtual segmentation, a multi-port node is architected such that each of its coaxial legs can be driven with its own unique RF lineup. In this case, unique does not necessarily refer to every spectrum slot (or occupied bandwidth as we move to a DOCSIS 3.1 world) being unique content on every port, but just that the composite line-up is unique to that port because of narrowcast service delivery. Common node platforms today of this style enable what is called up to “4 by 4” or 4x4 segmentation, meaning each of the four ports can have its own downstream receiver and upstream transmitter (or transmit path depending on the return optical technology approach used).

In actual node splitting, a new node is installed alongside (logically) the existing node, and, ideally, half of the connected subscribers are moved to the new node. This is a much more intrusive process with respect to the touch of the plant than virtual segmentation – and thus the attractiveness of the segmentable node. However, it is a very well-understood task that operators have optimized and established sound practices for over many years of growth and network migration.

Both approaches to node segmentation have negative consequences for plant powering, to a different degree. A virtual node split generally involves adding optical modules to the node, increasing its power consumption. An actual node split, of course, adds an entirely new node to power, although it may involve replacement of what was an RF amplifier such that the net impact is lessened. In both cases, the power consumption effects can be quantified with confidence, as these are generally available products.

In Figure 2, the resulting power consumption effect is shown assuming a methodical fiber deeper segmentation for a typical suburban architecture that would include both approaches (virtual and actual) over the progression of fiber deep. Each step incurs a powering penalty, as would be expected. The original 4x4 segmentable node (but yet not segmented) is the first bar on the chart and is the “100%” reference point for subsequent relative comparisons. Virtual segmentation increases the node’s power consumption with the addition of optical modules, resulting in approximately a 60%

increase on the node itself, but when considered in the context of the node serving areas as a whole, since none of the other components have changed, the effect is only about 2%.

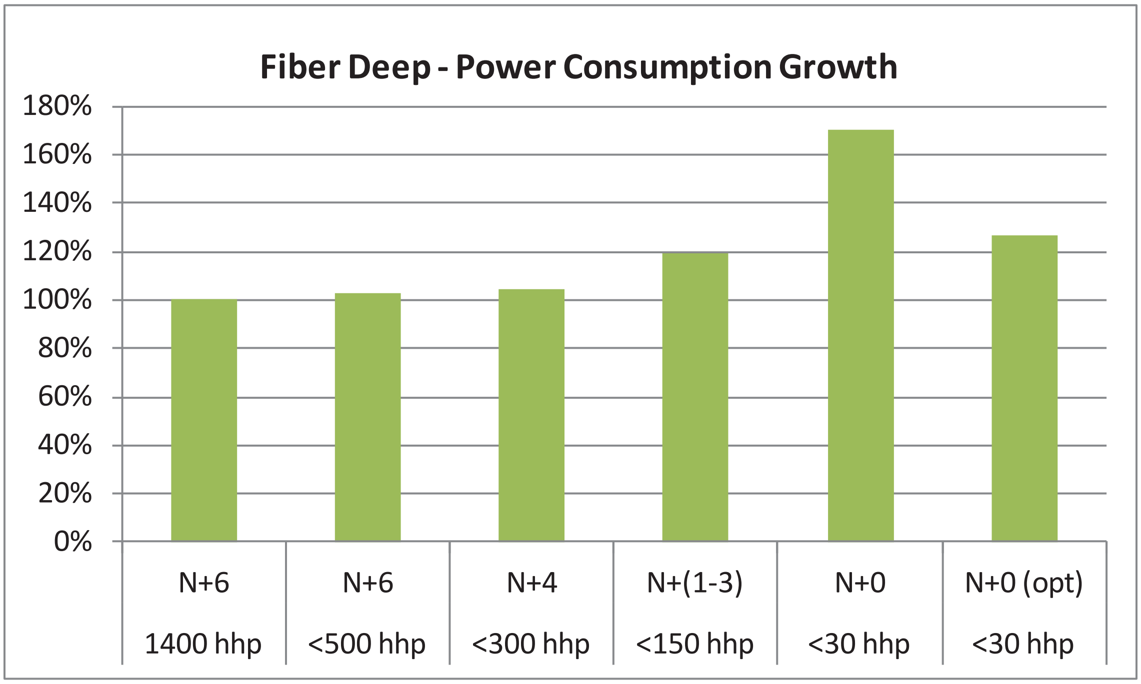

Figure 2 – Fiber Deep Implications to Power Consumption

After the virtual segmentation, fiber is physically pulled deeper and new nodes installed. In this real-life example, the node count increased to three in the first step to the "N+4" (average) case, and to eight in the ""N+small" cases of a 1-3 amplifier cascade. In both cases, more segmentable nodes were added where it made sense to do so in the design.

The evolution to N+0 does not look overly burdensome at about a 26% increase in power consumption of the node serving area for the case of "Optimized" N+0 as shown in Figure 2. However, of course, "optimized" is akin to achieving a greenfield placement of nodes whereby they would most efficiently cover the physical serving area. In a brownfield evolved from what was originally a >1000 hhp architecture, avoiding re-plumbing entirely (assuming same spectrum range) would mean that every active turns into a node location. This is the first bar for "N+0," in which 69 nodes would be required to be installed, increasing the serving area power consumption by over 60%. In practice, if an MSO were to implement an N+0 passive coaxial architecture, the result would fall somewhere between total inefficiency and total efficiency. Some consolidation would likely be possible and taken advantage of.

Business Beyond Usual

Operators have been offering business services to mostly small and medium-sized businesses through their DOCSIS network for many years. Some operators have also been serving larger business customers with fiber connectivity – mostly Ethernet- based services. This has been accomplished largely through parallel, dedicated fiber infrastructure. The synergy associated with fiber deep evolution of residential services, fiber-based business services, modular, segmentable nodes, and advances in wavelength division multiplexing (WDM) technologies has led to an interest in point-to-point Ethernet solution emanating from the residential HFC plant itself. Ethernet switches that provide point-to-point connectivity, aggregation, and MEF compliance, and which plug into node housings in the outside plant, are now available to meet the needs of larger business customers.

Most modern, segmentable nodes achieve the goal of making the network easy to segment by enabling pluggable module slots inside of the node housing for placing fiber optic receivers, transmitters, and accessories. Ethernet switches have been developed that plug into these slots so as to enable easy-to-deploy business services solutions. Of course, these additional modules result in an increase in power consumption of a node of the same configuration, but without switches. As with residential segmentation, Ethernet switch solutions are generally available products and therefore known quantities when it comes to power consumption impact. These effects can be easily calculated.

As modules that exist within a node, example devices may increase power consumption of the node they are housed in as shown in Table 1. The increase ranges from 12% to 42% for the range of possible configurations in a representative 4x4 node. Note that a fully segmented node, already chock full of optical receivers and transmitters, is constrained to a single Ethernet switch.

Table 1 – Power Consumption Impacts of Ethernet Switching – Modular Node

	Single SW 2 WAN/4 LAN	Two SW 4 WAN/8 LAN
1x1	+21%	+42%
2x2	+17%	+34%
4x4	+12%	

At the level of the serving area of the node, the impact is significantly lessened. Across a 500 hhp node serving area, for example, it amounts to less than 1% up to 4.4% as shown in Table 2, depending on the businesses in the serving area and the number of switches required to serve them. Each switch contains multiple ports so that several businesses can be served by a single switch.

Table 2 – Power Consumption Impacts of Ethernet Switch – 500 hhp Serving Area

Single SW/SG	Three SW/SG
1x1: +1.5%	1x1: +4.4%
4x4: +0.8%	4x4: +2.5%

To make footprint-wide estimates in the analysis model, we refer to sizing analysis done in the industry around the number of medium-to-large (i.e. non-DOCSIS candidates) business that may exist in a particular headend footprint. Mid-sized business will tend to occur in clusters in particular areas of a footprint and be more likely to exist together in an area served by the same node, whereas other parts of a footprint will be 100% residential or perhaps mixed-use with small businesses that can be served via DOCSIS. Because we are looking to understand the impacts for plants likely to be stressed for powering, we are obviously interested in cases where business services modules will be deployed so as to estimate the additional burden.

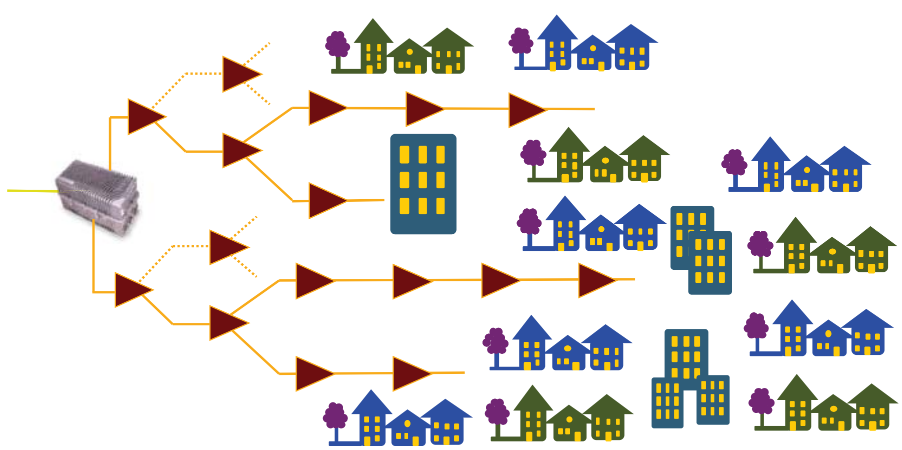

Figure 3 –Sizing the Mix of Business and Residential Services is Part of the CCAP Architecture Development and Specification

The estimate of switch and port count required was based on the sizing of business services support as engineered for the CCAP specification. In the capacity and port sizing of CCAP, a particular ratio of RF ports (nodes) to EPON ports is established, where each EPON port represents up to 32 termination points that could serve businesses. In the model, we simply assumed that the fiber connected businesses were instead served by Ethernet switch ports, and determined the number of ports from the CCAP port relationships.

Wifi™ Access Points

To meet the needs of on-the-go customers, operators are finding way to ensure that their HSD services can be accessed beyond the home. This offers customers higher data speeds than typically available through their cellular carriers, and provides a means for customers to potentially avoid data service charges on their wireless bill. Operators, of course, like the stickiness aspect, and generally try to target Wifi™ access points (APs) in the plant around where people are likely to congregate. In general, the APs are not deployed ubiquitously, except in markets with the density to make such a strategy sensible.

Wifi™ access points function essentially as Wifi™-to-DOCSIS transducers. Access points that have both Wifi™ and 4G radios also exist. Again, these products are generally available in the market today, so the power consumption associated with them is known. The impact of a Wifi™ AP deployment program can therefore be readily quantified. Of course, unlike the Ethernet switch, which is a module in a node, the AP is instead a separate plant module itself.

For a representative modular 4-port node platform and a generally available Wifi™ AP product on the market, the table in Figure 4 shows that the Wifi™-only AP solution power consumption ranges from 40%-83% of the node platform, where the range is a function of the configuration of the node – from a “1x1” configuration to a “4x4” configuration. In the latter case, the node is configured for four independent serving groups (that may have common broadcast content), and therefore has more optical receivers and transmitters installed to provide this segmentation, thus the relative comparison between the AP and node is friendlier.

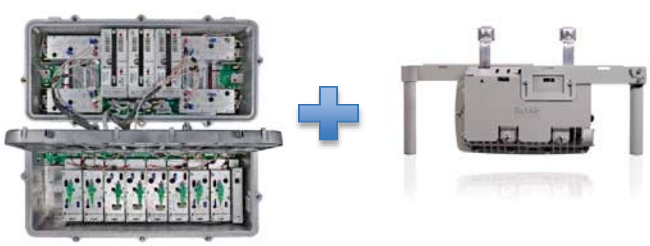

	Wifi	WiFi + 3G Picocell	WiFi + 4G Picocell
1x1	+83%	+90%	+118%
2x2	+56%	+73%	+96%
4x4	+40%	+52%	+69%

Figure 4 – Relative Powering Burden of Plant-Based Wifi™ APs Compared to a Modular Node Platform

Wifi™ APs hang on the plant much like nodes and amplifiers do. Therefore, it is insightful to consider a Wifi™ AP in the context of the other actives and see how it compares. Figure 5 shows this comparison for other representative node and amplifier platforms. As in Figure 4, the cases compared include units that also support 3G and 4G picocell functionality, although all of the analysis going forward for plant calculations will use the Wifi™-only models.

"Bridger"

Wifi	WiFi + 3G Picocell	WiFi + 4G Picocell
+100%	+130%	+170%

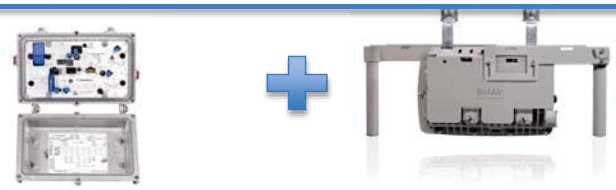

"Line Extender"

Wifi	WiFi + 3G Picocell	WiFi + 4G Picocell
+185%	+240%	+315%

Figure 5 – Relative Powering Burden of Plant-Based Wifi™ APs Compared to Amplifier Platforms

A simple conclusion that can be drawn from Figure 5 is that, from a powering subsystem perspective, a Wifi™ AP looks like a bridger-type amplifier. For a network that has a deep amplifier cascade whereby most of the amplifier count will be line-extender type platforms, this could have significant power consumption implications to extending HSD services with densely penetrated APs. Of course, operators will tend to be placing APs where the demand is most likely, and not ubiquitously. Some high density architectures, however, would lean towards APs more uniformly distributed in the plant. Indeed, some operators with precisely this architecture have quite successfully implemented Wifi™ access in part *because* of the advantages of high-density, since it lends itself to service efficiency.

The Next

There are other potential evolutionary paths on the horizon that may impact power consumption in the plant. The persistent growth of downstream HSD bandwidth has raced along at a compounding rate of about 50% per year (referred to as 50% CAGR). This presents tremendous challenges for operators with existing capacity limitations in the HFC plant and a downstream spectrum that is already 100% utilized. Figure 6 is a sample analysis that quantifies how this aggressive bandwidth growth impinges on various capacity thresholds (horizontal lines), depending on the evolution of video services such as 4k HD (QFHD), IP services, and steps MSOs may take to relieve downstream spectrum congestion, such as Switched Digital Video (SDV) and Digital Terminal Adaptors (DTAs).

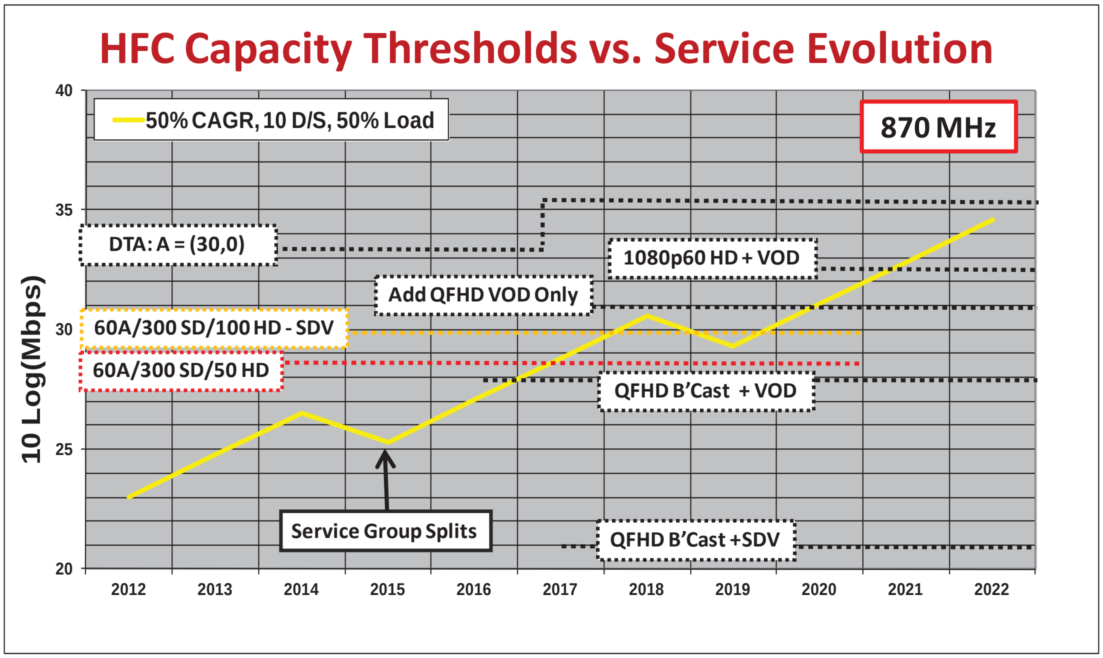

Figure 6 – Persistent Downstream Bandwidth Growth Could Exhaust Available Network Capacity

While new bandwidth efficiencies are being brought to bear on the currently occupied spectrum, to achieve the DOCSIS 3.1 objectives of 10 Gbps in the downstream *and* 1 Gbps in the upstream will require new spectrum beyond 1 GHz. For example, spectrum usage to 1.2 GHz and perhaps 1.7 GHz is being entertained as part of DOCSIS 3.1 and IEEE 802.3bn (EPoC). The impact is significant to the network, so if implemented at all it is likely to occur in phases associated with growth trends observed in the interim and reasonable Capex planning. Figure 7 shows an example of a possible evolution of HFC spectrum over the long term.

Turning on this new spectrum will require new AC power. And, in the coaxial domain, the effect can be disproportionate to the amount of spectrum added because the RF payload is tilted to account for frequency dependent loss. As such, the RF amplification function has to work harder to deliver the identical level at the home for higher frequencies. The good news from an analysis standpoint is that HFC optics and RF amplifiers are known quantities with respect to power consumption, so extrapolating the new power required for wider bandwidths is mathematically straightforward under some assumptions of implementation.

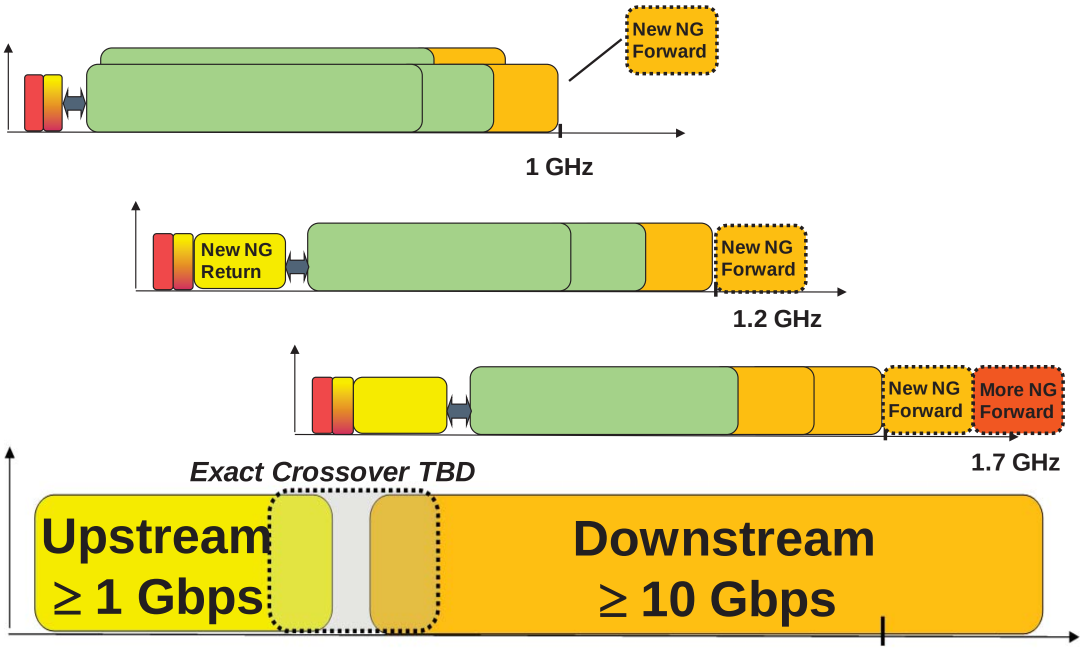

Figure 7 – Possible Long Term Evolution of the Coaxial Spectrum

Another architecture evolution of great interest is the concept of migrating to a digital optical architecture with some form of distributed physical layer, such as the anticipated primary EPoC (IEEE 802.3bn) architecture, or remote CMTS functionality for DOCSIS systems. A very simplified view of this concept is shown in Figure 8.

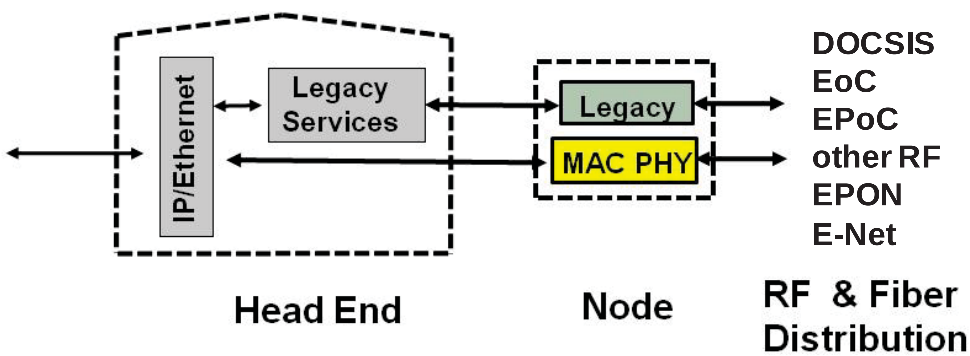

Figure 8 – Simplified Distribute Architecture

The exact split of functionality in the downstream and upstream is the subject of much debate, but a driving force for operators is that moving some processing functionality into the plant has powering and space benefits in the headend where constraints are mounting. There are also performance benefits that result by the removal of linear optics. Estimates of the power consumption effect in the plant can be made here as well by understanding the impact of the key processing chips when pushed out to the node.

With the knowledge and estimates as described above – the Now and the Next – we will consider a hypothetical long term service evolution for a hypothetical cable company, ACME CableCo. The service evolution will be based on the activity and discussions in the industry today, and where that might take operators over the long term (about 15 years). Associated with the service evolution will be the supporting architecture evolution, from which we will calculate the increase in power consumption that will occur over the time period.

It's Not All Bad News!

While the passage of time sees new services and migration paths that cause an increase in power consumption in the plant, those same years create opportunities for saving power as well. We discuss several examples used in the analysis below.

Step on the GaN!

RF technology has been through several semiconductor generations that have increased bandwidth, output levels, and linearity performance. Multiple generations of Silicon technology evolved into multiple generations of Gallium Arsenide technology (GaAs), and we are now into a second generation of Gallium Nitride technology.

RF power efficiency is one of the levers available for future powering savings as broadband hybrid technology is further optimized. The industry is once again considering wider bandwidths, a larger total RF load implied, and without compromising linearity performance as we discuss advanced M-QAM formats such as 4096-QAM. Recent detailed studies, such as the comparison shown in Figure 9 [7], and expectations through 2018 for the future of CATV broadband hybrid technology have been used to project the possibilities for the long term. These enhancements include efficiency gains as "Green" becomes an important objective for operators and the vendors who deploy hybrid amplifiers in their HFC actives.

Some of the efficiency gain that will be made available through hybrid evolution will be spent as bandwidth extension to 1.2 GHz, so the net effect on power consumption works out effectively to a *reduction in the growth* of power consumption. This may sound disappointing, but the decrease is significant, as we shall see.

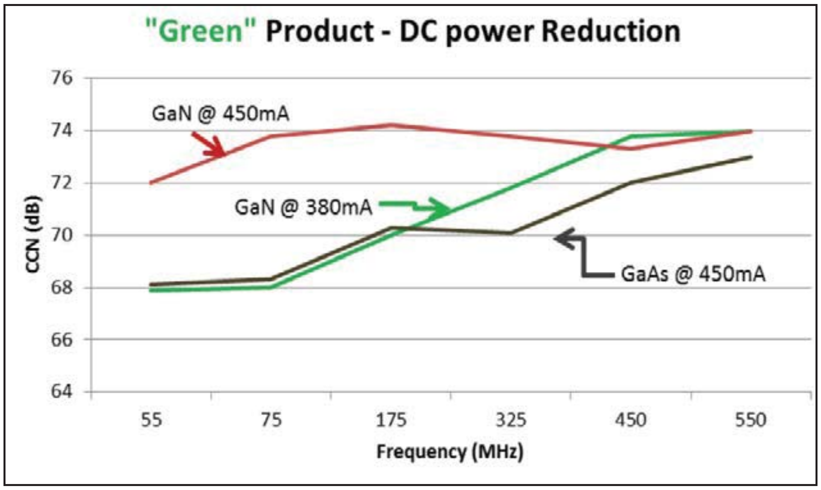

Figure 9 – RF Hybrid Technology Evolutions Will Provide Efficiency Gains [7]

Turn Brown into Green

An architectural efficiency that we will take advantage of is that, armed with the knowledge of a long-term evolution path to FTLA, optimization of locations of the end-state nodes can be built into the evolution plan more effectively. There is a decrease in the total number of nodes in the network end state when optimized compared to an approach that simply turns every current active location into a node location. While this involves some re-plumbing, if planned as part of a multi-year evolution effort, there is an opportunity to spread out the cost impacts.

Typically, HFC brownfields were built with relatively long RF cascades and over large geographical serving areas, not to mention at spacing associated with less forward spectrum. To start from scratch with FTLA in mind, today's technology, as well as an average hhp/FTLA node, the design would be much different. How different and how many fewer actives might exist in the optimized FTLA case can be quantifiably determined on a system-by-system basis using standard HFC design tools and rules. We have done so for an example system with a typical hhp/mile density, and used these results herein.

Off the Grid?

We will take the long view to consider two additional potential efficiency mechanisms. One is a direct attack on plant powering – we boldly consider the idea that in an FTLA architecture there is the potential for some nodes, such as the smallest (single port) ones, to be removed from the grid by powering them with solar panels. This is done with traffic signals in some countries and locations in the Unites States today (with back-up power in both cases, of course). In that application, both municipal power budgets and emergency operations (power outages) are key benefits. Figure 10 shows a deployment on a pole of solar panels used for traffic signal applications.

Note that the consideration of such an approach relies on two premises. First, many FTLA nodes have a strong possibility of being small, single-port device in many common HFC geographies and densities. In our reference example, this ends up at about 40%. Current vendor product portfolios include kits that convert RF amplifiers into nodes. Line extender amplifier conversion kits fit into this category. Being small, the power consumption burden is smaller than a fully segmented, four-port node, for example.

Figure 10 – How Far Can Cable Take "Green" in the Outside Plant?

Second, an FTLA architecture does not have a subsequent amplifier to drive power down to, so solar cells only need to worry about powering their node. Power supplies in place would still be used for co-located Wifi™ Access points. Larger nodes such as those that house Ethernet Switches, for example, would not be as attractive candidates. Of course, as the technology evolves over the years, the opportunity to be even more aggressive may exist.

We also consider the introduction of power factor corrected supplies in the actives themselves to maximize the efficiency of any deployed AC supplies by reducing the reactive losses in powering the loads in the plant. These gains are modest, but contribute to efficiency losses, and are easily determined.

Peaking at the Right Time

Lastly, we will take a long view of something yet even to be deployed – DOCSIS 3.1. In principle, once the entire spectrum is converted to DOCSIS 3.1 OFDM, we will have a single unified transmission system. One of the disadvantages of OFDM signals is that the waveform has an inherently high peak-to-average ratio (PAR) because it is comprised typically of independently modulated QAM carriers summed together. Of course, on the HFC network today are also independently summed QAM carriers – just of wider individual bandwidths and non-overlapping. In both cases, the composite waveform looks noise-like and requires actives with a dynamic range that supports it.

Because this has been a historic issue for OFDM in environments that cannot afford to run with low power efficiency (i.e. battery powered devices, such as cell phones), many schemes have been developed to limit the PAR of the OFDM waveform. Tone reservation, whereby subcarriers are set aside and modulated to reduce the PAR on a symbol-by-symbol basis in real time, is one such technique, and other more elegant approaches also exist. The penalty paid is OFDM overhead, and trade-offs must be made between overhead losses and dB gained.

These peak-to-average power reduction (PAPR) techniques can be applied as well in the cable environment, and this is most straightforward when the system is entirely DOCSIS 3.1 OFDM. A difference in broadband cable is that the accounting for waveform peaking must take into account that the spectrum transmitted onto the cable will be tilted before hitting the wire. The calculation made to adjust the composite waveform to minimize peaks must take this into account, but this works out essentially to scaling factors in the algorithm for the configured tilt used.

Simulations show 3-6 dB may be available in peak reduction, which could translate to power efficiency gains in the RF components of 2-4x. We will use 2x as the long term best case (3 dB of improvement due to PAPR). It is assumed this occurs in a new generation of DOCSIS algorithms timed roughly to when the entire spectrum is OFDM, which is assumed to be >10 years down the road. It is modeled as achieving full 50% improvement over the course of the last seven years of the model (years 8-15). The philosophy here is that as the legacy services wind down and consume a smaller part of the spectrum, PAPR gains begin to be seen and gradually increase as OFDM dominates the composite power, up until the potential gain is fully realized at year 15.

In summary, the following efficiency gains are incrementally factored in over the course of time per the ACME CableCo Master Plan:

1) RF technology evolution
2) FTLA architecture optimization
3) Power Factor Corrected power supplies
4) Solar powering of some FTLA nodes
5) Peak-to-Average Power Reduction (PAPR) techniques

The efficiency benefits that accrue to RF actives as a whole are specific to the active, based on its function (node or amplifier), and port count. Assumptions based on current products are used in concert with the projections in [7] to calculate the powering of newly introduced actives beginning with the next phase of node splitting identified in 2018. This is also the time period at which power factor correction is applied to all newly introduced actives.

ACME CableCo's: The Master Plan

Clearly, there are many moving parts to consider as services, architecture, and technology evolve together. Some decisions have to be made, and, like all smart decisions, they must have a sound and reasoned balance of the concern for the Now, the Next, and even the perceived end state. For example, it is generally difficult to make technology decisions today for what might be 15 years down the road given the pace of technology evolution. However, while recognizing this fact, it may be reasonable for an MSO to view the larger picture as one that ends in an all-fiber solution, and invest in initiatives aligned with that viewpoint.

Our hypothetical cable company, ACME CableCo, will be assumed to be going through a 15-year service + architecture + technology evolution. ACME's migration will consist of an "all of the above" type of assumption for new services and spectrum, so inherently a relatively aggressive case of power consumption growth, which is the intent – reasonable scenarios to bound the high end of the scale.

While ACME CableCo is hypothetical, the analysis for ACME is based on an actual plant footprint and architecture of a major North American MSO. An in-depth study was performed on this network to estimate the cost, bandwidth, and power consumption associated with a fiber deep evolution, as is projected for ACME CableCo, all the way from virtual segmentation to N+0 and optimized N+0, as shown in Figures 1 and 2.

Figure 2 is the fiber-deep only part of the baseline example taken further for this modeling exercise. That detailed network migration analysis of a real HFC system was the starting point, on top of which were added new service components (switches, Wifi™ APs) and upon which extrapolations such as new spectrum and efficiency effects were performed. While the integer granularity and efficiency of use of plant AC supplies and the active equipment's power supplies was factored into the actual reference design for fiber deep above, we did not attempt to do this with the add-on services and architectures. The results are estimates of pure new power consumption demand of the components being added into the plant.

The 15-year, year-by-year, plan is shown in Table 3.

Table 3 – ACME CableCo's 15-year Plan

2013	Node Split ~1500 hhp Footprint (500 hpvn/N+6 to 300 hp/N+4)
2014	Introduce Ethernet Switched Biz Services Over HFC: Growth @ One New Customer/year
2015	Introduce Wifi @ 10% of footprint; Subsequent Growth 25% / year for 6 yrs (40% of footprint)
2016	CCAP & DOCSIS 3.1 penetrating in scale
2017	
2018	Begin Node Split Phased over 3 yrs N+(1-3) / 150 hhp while Introducing 1.2 GHz D/S (2nd Enet Switch for growth pace)
2019	
2020	Completed Final Phase of Node Splitting
2021	Begin Targeted 1.5 GHz BW for High Growth Zones (20%)
2022	Add 20% new FTLA Segments (10%/yr thereafter); PAPR algorithms available, declining legacy spectrum
2023	Activate 1.2-1.5 GHz (Gbps Biz over Coax) (i.e. EPoC or D3.1); CCAP EPON Sizing, 10% Nodes = 100% Biz Oppt'y
2024	Growth in 1.2-1.5 GHz
2025	70% of Footprint @ N+0
2026	
2027	Complete 1.2-1.5 GHz Build Out (Module per four businesses)
2028	Take Rest of Network to N+0 = NETWORK NIRVANA !

The details of how much, how many, and/or how fast for the initiatives identified in Table 3 can be varied in the analysis model. While the year-by-year list looks imposing, this has to do in part with the fact that major initiatives are spread out over several years of implementation, as is apt to take place in the outdoor plant in a pay-as-you-go fashion. The evolution can, in fact, be simplified by breaking it out into the five major elements:

1) Adding Ethernet Switches to serve businesses
2) Adding Wifi™ Access Points for on-the-go HSD services
3) Business-As-Usual Fiber Deep – multiple phases ending in N+0
4) Extended forward HFC bandwidth (aligned with FTLA only)
5) Digital Optical Overlay architecture/module for 1.2-1.5 GHz Gbps services

Let's evaluate how this roll-out of new services impacts plant power consumption. For manageable numerical purposes and an easier understanding, we will confine the example to a single serving area, which is over 1400 hhp in this example. Entering 2013, it has just been virtually segmented, using a 4x4 modular node, four ways.

Let's Run the Numbers!

Figure 11 shows the cumulative power growth associated with the 15-year timeline, with none of the efficiency and power saving steps included. It is today's technology and capability and, along the timeline where identified (spectrum and architecture), extrapolated and extended to implement the migration.

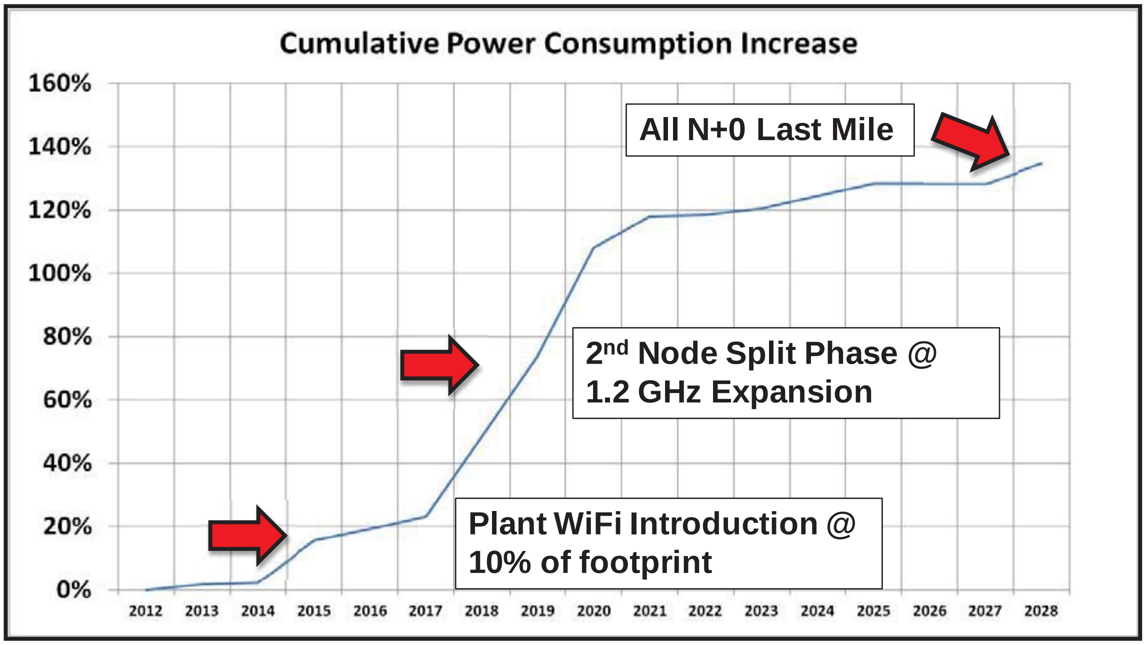

Figure 11 – ACME's Cumulative Power Consumption versus Time

Three notable upticks identified in Figure 11 stand out, and for good reason:

First, introducing Wifi™ APs looks to the powering subsystem as a set of new mid-sized amplifier platforms. As such, introducing Wifi™ deployment will obviously impact the power consumption in the footprint.

Second, as we enter the 2nd node split at the end of the decade, the implication is that we will do so using 1.2 GHz capable nodes, and that this new spectrum will require an increase in DC power to the hybrid devices. The assumption in this case is that the system is not "shortened" to manage the extra loss – the devices must tilt to the desired output power. The effect is shown in Figure 12, where the net result is about a 2.7 dB increase in the total RF load.

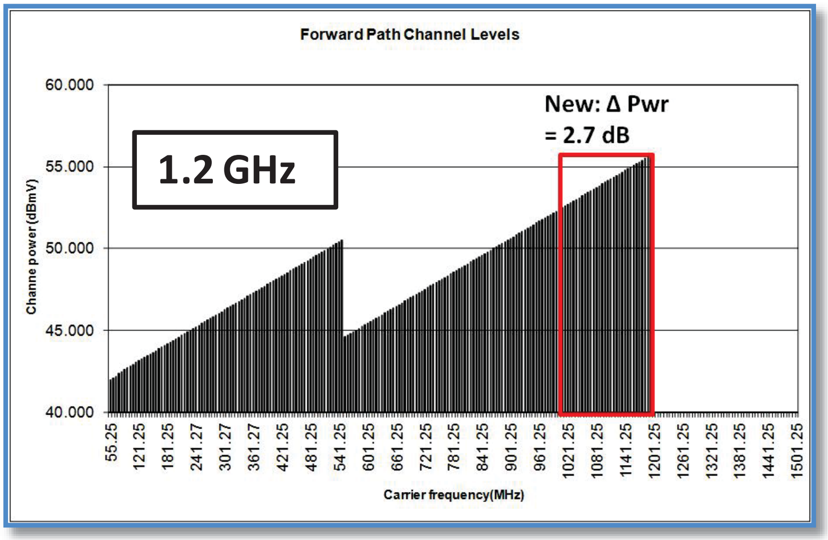

Figure 12 – Adding Spectrum Increases the RF Power Load Disproportionately on the Tilted Coaxial Output

The increase in total load, under the initial assumption, is assumed to be handled by increasing the power consumption of the device to maintain the linearity (ignoring for illustration purposes possible maximum power constraints, thermal design limitations, or platform power supply budgets).

Finally, as we install more and more nodes, the power consumption rises with the exchange of RF amplifiers for nodes, which from a powering perspective can be considered as amplifiers with optics. There is a steady pace of powering growth as every active becomes a node, and a noticeable uptick at the end of the migration cycle in 2028 as the final 30% of the network that has not yet been converted to FTLA is converted.

The result? A roughly 2600 watt (W) system becomes a 6200W system, and, as shown in Figure 13, leads to an erratic growth trajectory that may imply periods of high investment in the powering subsystem. Also, through the end of this decade, the growth in plant powering grows at a compounding rate of nearly 10% before settling during the second half of the migration period.

This reference will be used as a baseline to understand how we might expect to do better with "normal" technology and architecture evolution, and how we might take specific initiatives to ensure that we do "Green" better.

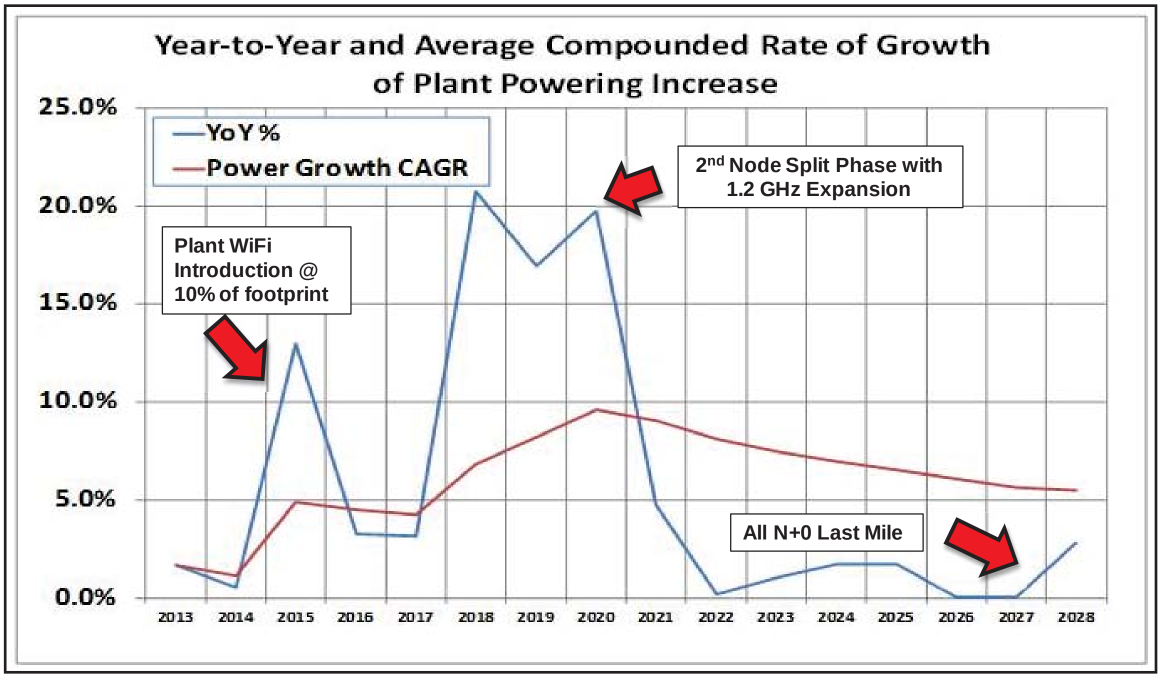

Figure 13 – ACME's Year-on-Year Power Growth and Compounded Average

However, before we proceed with the extrapolation and extensions, we chose a second sample architecture from a different major MSO for comparison. This provides a sense of the sensitivity of the analysis to architecture type and design approaches. It also provides a fundamental sanity check of the inputs and projection analysis approach, since a 15-year evolution term compounds what might otherwise seem to be very small variations when built upon as references for future evolution.

Because it was a different system, there were assumption differences to be made to make as close as possible to apples-to-apples. The most significant difference, however, were "knowns." In the reference case above, the homes passed density was 85 hhp/mile, and there were about 17 plant miles. In the comparison case, a high density 120 hp/mile (over 40% more hhp/mile) over a compact 4 miles of plant miles was analyzed.

The analysis comparison of the reference case and the high density case is shown in Figure 14. The high density case slightly smoothes the growth ramp during the 1.2 GHz upgrade cycle of the second node split. The higher density perhaps allowed more flexibility of HFC design, since such architectures typically involve less coaxial loss and more passive loss, and it is the coaxial lost that suffers from frequency dependency that decreases reach. It stands to reason that a more compact architecture could be less stressed for reach (and therefore number of actives) vs. frequency.

However, at the end of the 15-year term, the cumulative power growth ended up very similar at 135% for the reference case and 146% in the high density case – just an 8% difference.

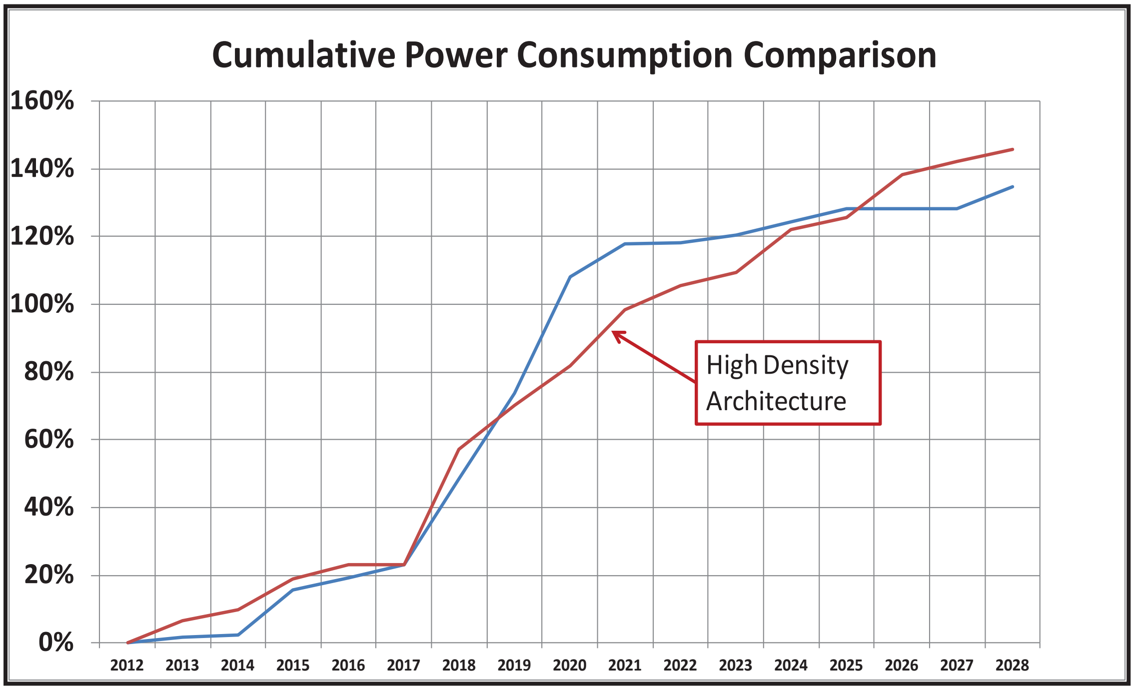

Figure 14 – Cumulative Power Growth Comparison of Reference case and Higher Density Architecture Case

We will carry forth the evolution plan with just the one original reference example. Note that a snapshot of the analysis tool is shown in the Appendix.

<u>Higher Efficiency? Yes We Can!</u>

Using the same 15-year migration plan, Figure 15 does the accounting on the following subset of efficiency assumptions previously described:

1) RF technology evolution
2) FTLA architecture optimization
3) Power Factor Corrected power supplies

Including architecture and efficiency opportunities, as opposed to 15 years of the same technology, has a clear benefit to the growth in power consumption.

The optimization of the long term architecture results in about 25% fewer actives. Also, the continued evolution of RF suggests that only 10-20% of additional power consumption will be needed by the 2018 time frame to enable the extra 200 MHz above 1 GHz over the long term.

In absolute terms, the modified power growth with the above three efficiencies included takes the 2600W system to about 3900W, for what now is about a 50% total increase. This more modest increase translates to $2000 annual savings for this single node footprint. For a 40-node headend, for example, this could mean $80,000 of annual savings, assuming this system design is typical of the overall headend footprint.

The assumption of a new generation of GaN available with higher efficiency in 2018 readily shows itself in Figure 16. There, we can see that the large step associated with extending the bandwidth to 1.2 GHz under current RF technology constraints had caused a noticeable step in year-to-year growth (blue). We can see a much more modest ramp (green) over the three-year period that these next generation actives are deployed. This advantage also manifests in a long term compounding average being reduced by about one-half, from 5% to 2.5%.

In general, a situation that looked very troublesome in Figures 12 and 13 now looks manageable with a properly planned architecture evolution and by taking advantage of likely advances in technology. While a power consumption increase is never a good thing, the extent of additional services and bandwidth managed to be deployed over this less burdensome power growth scenario seems like a very effective tradeoff.

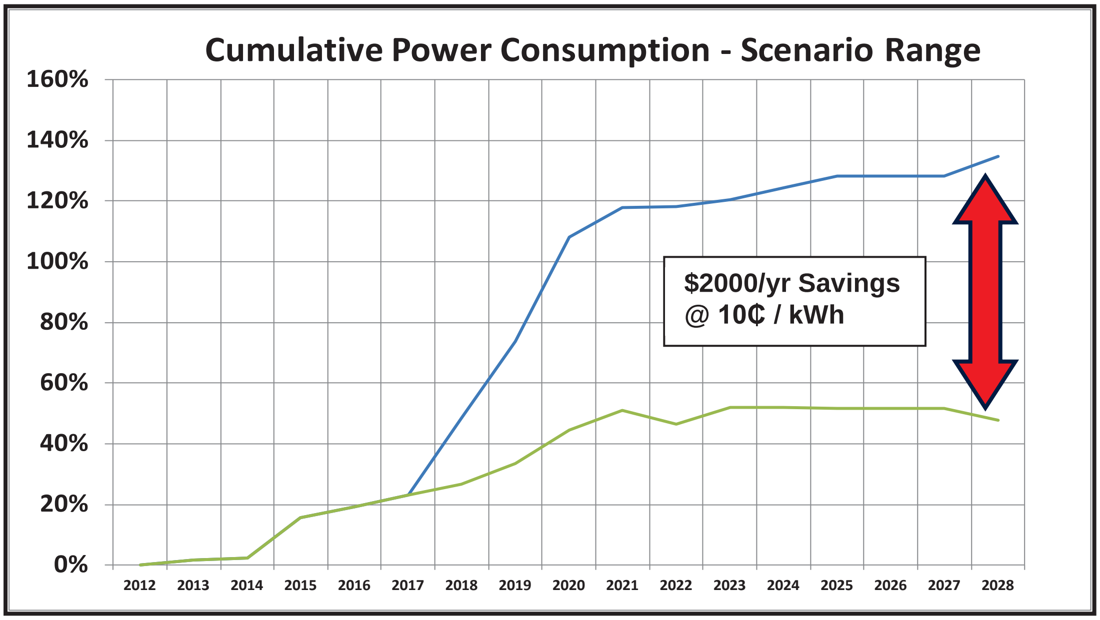

Figure 15 – Accounting for Possible Efficiencies of Architecture and Technology

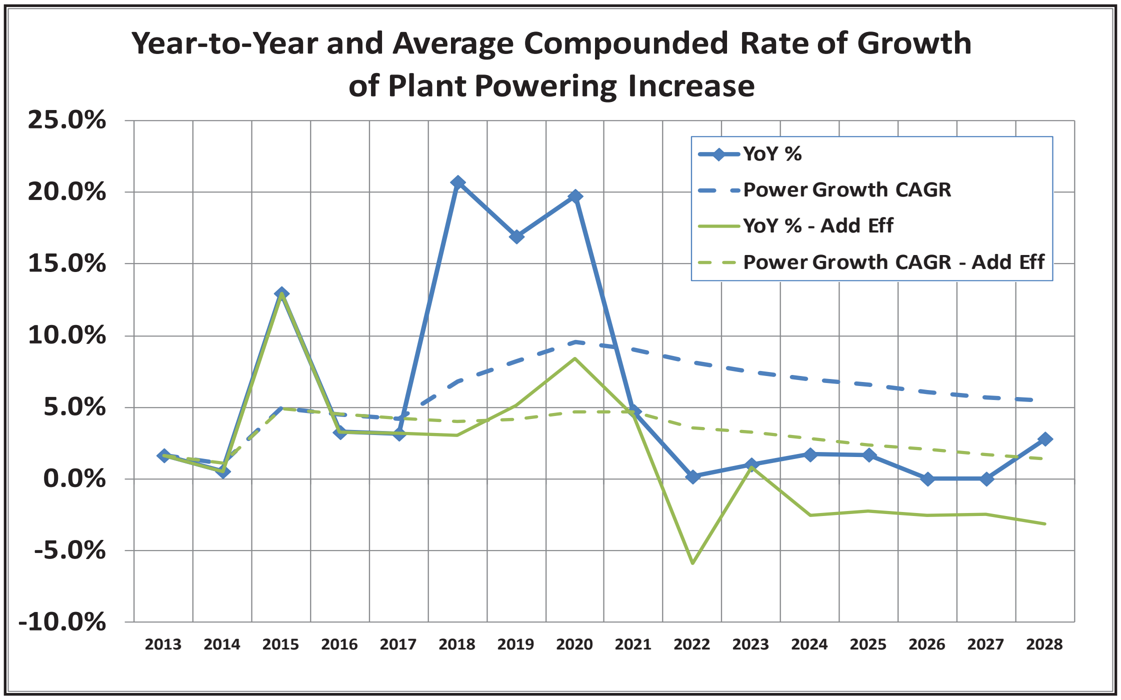

Figure 16 – Comparing Powering Growth Scenarios

<u>Sunsational !</u>

Figures 15 and 16 provided a glimpse of a power consumption growth management with what might be considered standard-fare tools at our disposal over time. Now we embark on some out-of-the-box thinking. Our first example will be the concept of solar powering of the smallest nodes in the fielded end-state, which account for about 40% of the final total. These nodes are anticipated to consume about 40W at that point. By way of comparison, traffic signals based on LED lighting will use lamps that consume up to 25W, which is a major reduction compared to typical incandescent bulbs used in most signals today, and an enabler of solar solutions for municipalities to consider.

The declining cost of solar panels and the range of technologies to trade-off cost, size, and efficiency is an enabler for cable to consider such a long term vision. For an example of modern projections of solar panel technology, a 230W solar panel is expected to be only about 2 ft by 3 ft by 2015, as shown in Figure 17 [9].

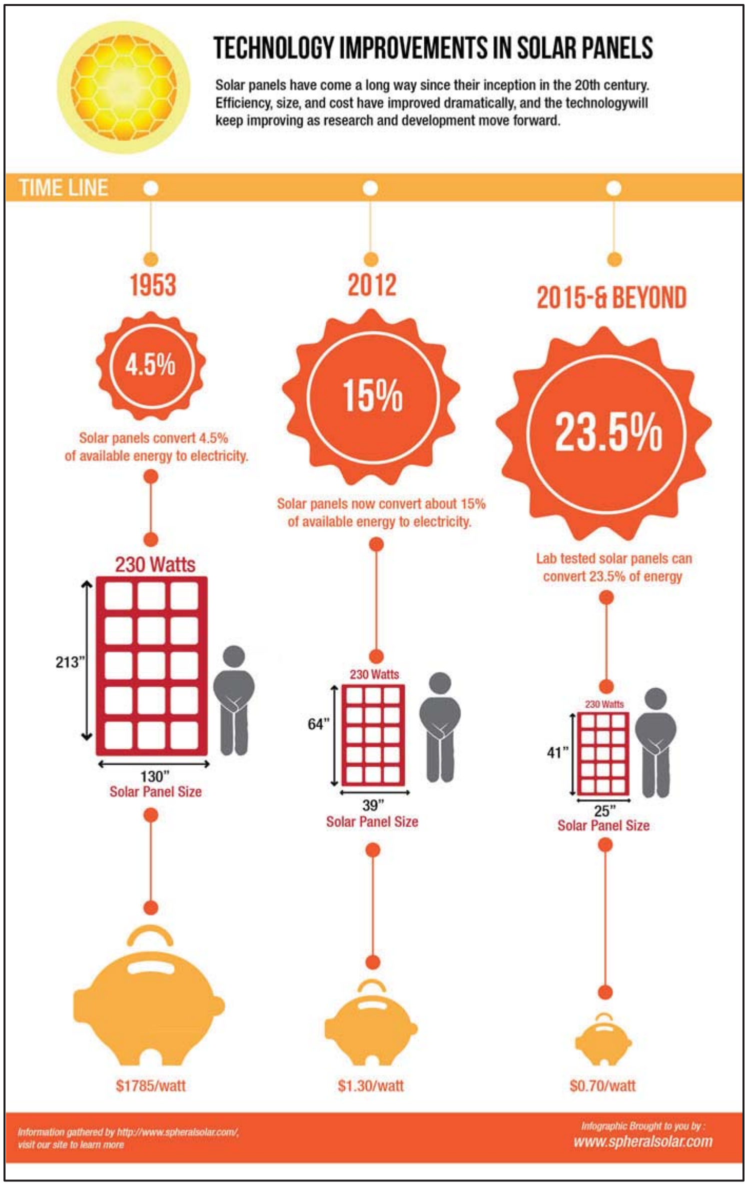

Figure 17 – Solar Panels – Future Costs, Size, and Efficiency

Of course, a complete solar system requires a charge controller, batteries for storage, a power inverter perhaps for AC supplies, and, for 24/7 services, standby power. Actual implementation is more than pointing a panel or two towards the sun and hoping for typical weather. And, there are obviously locations and plant types (aerial or underground) that are better suited to such an adventure. It is often good practice to keep active components away from the sun to the extent possible. However, it is insightful to see where this technology could take us if it continues to mature, HFC architectures (such as FTLA) become compatible, and the significance of “Green” continues to grow.

The modeling problem for the architecture is quite straightforward. We simply select the single port nodes that remain and, during the migration to FTLA, allow these models as they become deployed to be removed in the analysis summation of power consumption – we mathematically take them off of the grid. We are not accounting for the percentage of time they may require back-up powering, etc., just simply using an assumption that panel sizing consistent with solar radiation patterns such as shown in Figure 18 [8] can yield a sensible, robust design beginning in 2021.

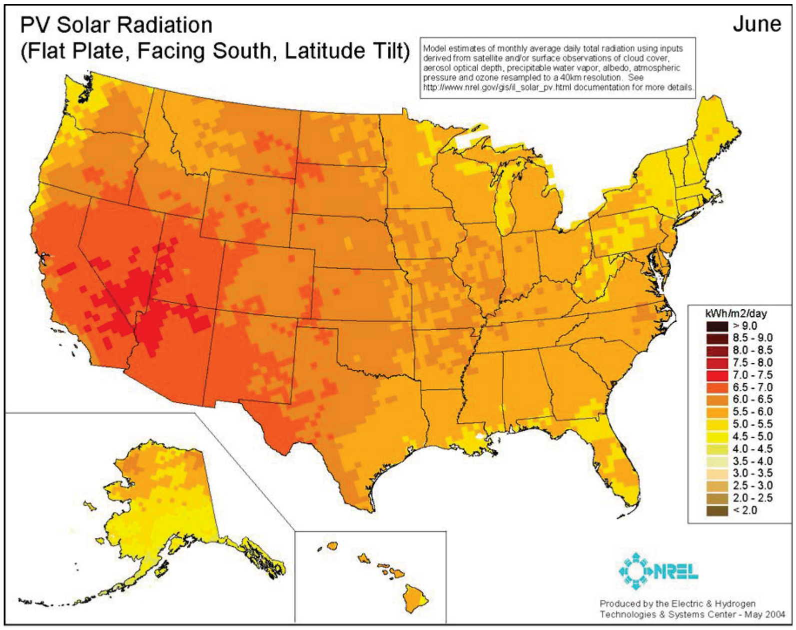

Figure 18 – Seasonal Radiation Models are Key to the System Engineering of Solar Alternatives

Performing the "math of subtraction" for all single port nodes as they are introduced into the architecture beginning in 2021 yields additional reductions on top of the efficiencies shown previously. These results are shown in Figures 19 and 20. We have now reduced the composite increase in power to a growth of about 20%, going from the 2600W start to about 3100W. Compared to the "no new efficiencies" case, the savings in annual cost is about $2700.

And, again, the powering delta and cost also could be looked at as a whole lot of new services deployed and substantial architecture evolution, all for the Opex powering cost of less than $500 extra dollars per year. It would seem that the extent of new services would easily cover this added Opex and provided plenty of new ARPU to spare.

Lastly, on a compounding basis, we have reduced the YOY growth rate to about 1%. While that may sound hardly noticeable in normal operations, the front-loaded growth rate does show a 4-5% compounding growth in consumption for a five year period that runs to the end of this decade.

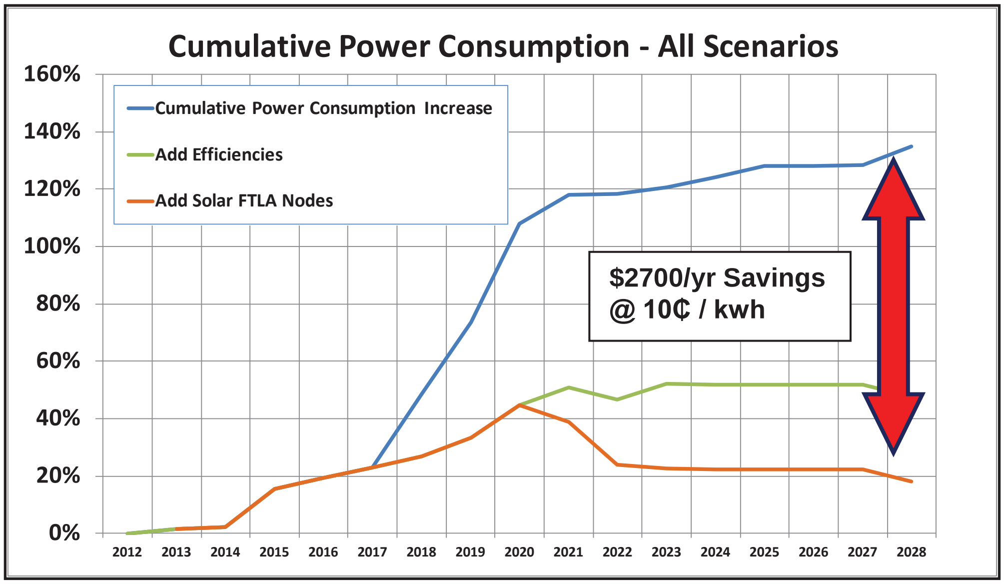

Figure 19 – Off-the-Grid FTLA Implications to Power Growth

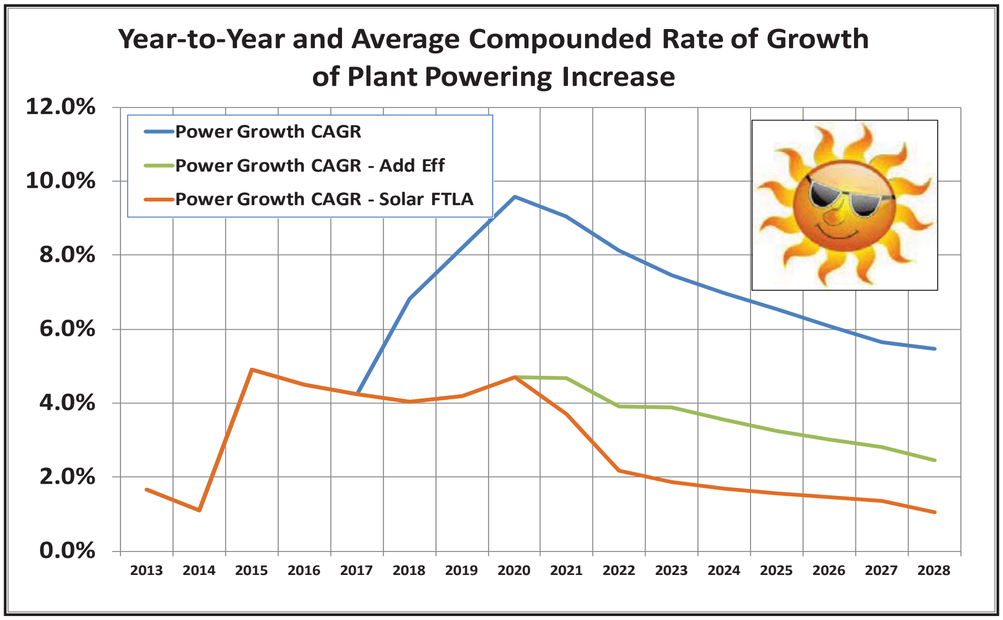

Figure 20 – Enabled by FTLA - No Powering Through the Distribution Network

Operating at Peak Removal Performance

The second outside-the-box consideration for long term evolution is to take advantage of PAPR methods to decrease the peak excursions of the composite noise-like waveform, and thereby allow the RF network to run at a more efficient biasing point for equivalent linearity. Each dB that the peak can be reduced from its noise-like PAR counts towards an equivalent linear ratio of power consumption savings.

As indicated previously, we have assumed that 3 dB will be achievable because of the evolution to an all-DOCSIS 3.1 OFDM system, providing unified control of the aggregate waveform. The modeling assumption is that the full PAPR gain will become gradually available as legacy spectrum becomes a smaller and smaller contributor to the aggregate power load – in other words, legacy services are removed to the point where their effect on waveform statistics is gradually minimized. It is modeled with linear scaling savings per year, although in practice it is likely to be non-uniform.

The savings for each active device can be converted to power consumption savings for each active by knowing what percent of the power consumption is associated with the RF chain, as opposed to the optics, the return path, and to other overhead losses for each 1, 2 and 4-port model used in the architecture.

The (incredibly coincidental) results are shown in Figures 21 and 22. The "incredible coincidence" outcome is that the range of assumptions of evolutions and efficiencies that end with considering the use of PAPR results in nearly exactly a *net zero* increase to long term power consumption.

And, as we see also often see with bandwidth models, there is a power consumption "bubble" or hurdle along the way to clear which also occurs around the year 2020 that may require investment, but the long term outlook looks quite promising. And, as a reminder, this was very much an "all of the above" evolution of services and architectures. In most cases, one or more of the initiatives identified will not be part of an operator's plan. Therefore, it is quite insightful to see that even when a very broad range of possibilities are included in the evolution mix, the situation appears very manageable. And, given that only a subset of the full suite is likely more typical, the situation could be described in the long term as potentially downright comfortable.

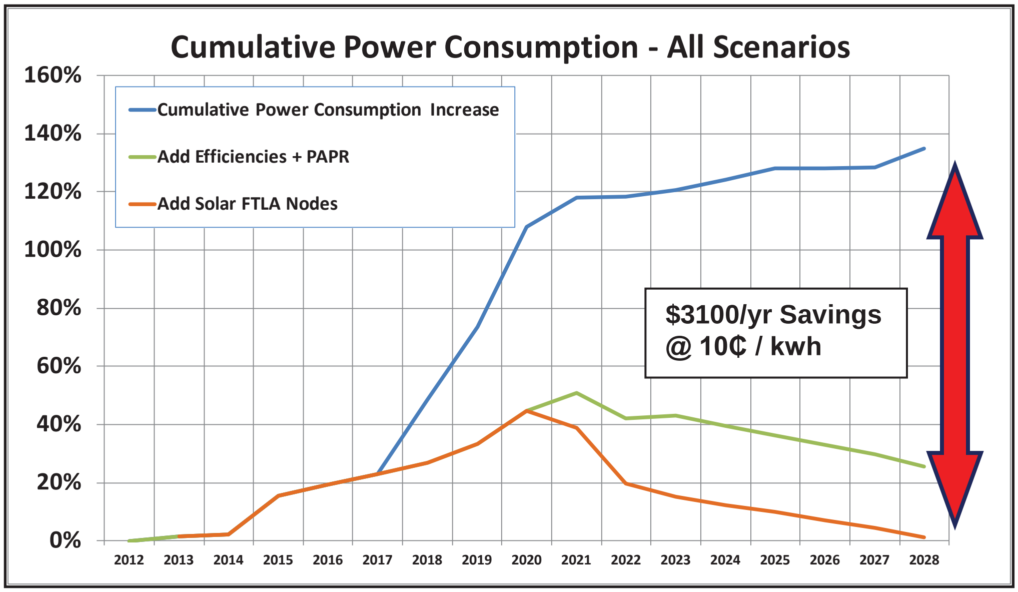

Figure 21 – The Case for Including PAPR – Long Term Power Consumption Growth is Nullified

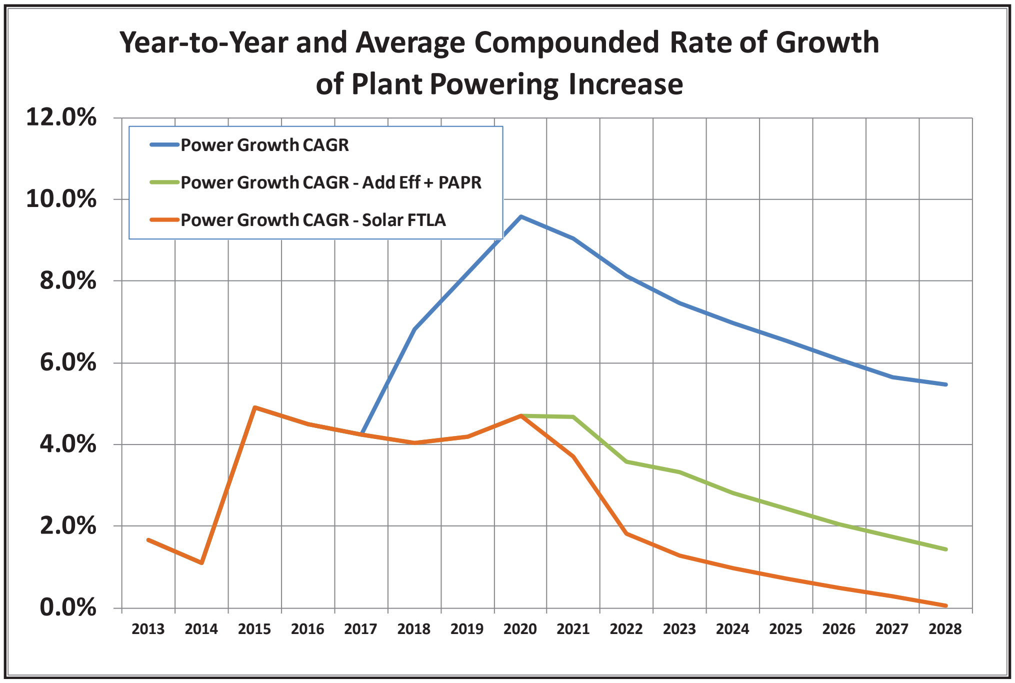

Figure 22 – The Case for Including PAPR – Long Term Power Consumption Growth is Nullified

Conclusion

There is a lot going on in the cable industries these days as operators continue to position themselves as the preferred supplier of choice for consumer video and broadband services. Network migration is expected to deliver the range of new services, data speeds, and video content by providing sufficient capacity at the right time, and by managing the HFC infrastructure optimally for both residential and business customers. Creative solutions have and will continue to be developed that allow operators to maintain their leadership role.

Unfortunately, from the standpoint of energy consumption, these developments introduce more power-hungry components into the plant. So, while an actively powered plant is a key advantage of HFC, operators are keen to find way to save Opex – such as through reduced powering costs –while at the same time making commitments to "Green" initiatives in their headends and premises equipment. The plant is a logical next target for eco-friendly initiatives, and cable operators can get out in front of this by planning their evolution with powering efficiency in mind.

In this paper, we have shown that there is potential for an extreme growth in power consumption given the service and architecture evolutions that could take place, such as those described in this paper. However, we have also showed that this growth can be tamed with effective use of what might be considered "normal" technology evolution and a well-managed migration. The return in the form of ARPU for the modest added powering costs seems favorable when taking into account tools to manage powering efficiency.

Furthermore, we have shown that with some out-of-the-box thinking, we can paint a picture whereby in the long term we may even be able to nullify *completely* the growth in network power consumption associated with very aggressive plans, such as those laid out by ACME CableCo.

It is our hope that this analysis helps operators think about the powering implications in the plant associated with the multitude of exciting initiatives. In so doing, we also hope that it plants the seeds of proactive thinking around what elements of technology and architecture make sense in order to address the potential burden these initiatives may place on the powering subsystem.

Acknowledgements

The authors would like to thank Marcel Schemmann, Fred Slowik, and Brent Patterson of ARRIS for their valuable contributions to this paper.

Bibliography

[1] Howald, Dr. Robert L, "A Balance of Power," SCTE SEMI Forum, April 10, 2013, Atlanta, GA.

[2] Howald, Dr. Robert L, "Powering Consequences of New Services and Architectures," SCTE SEMI Forum, March 15, 2012, Philadelphia, PA.

[3] Howald, Dr. Robert L and Fred Slowik, "Unleashed Decibels," SCTE SEMI Forum, May 24, 2011, Bellingham, WA.

[4] Howald, Dr. Robert L and Floyd Wagoner,"Right-Sizing Your Fiber Deep Architecture," Communications Technology Webinar, August 10, 2010.

[5] Miguelez, Phil, Fred Slowik, Stuart Eastman, "Refueling the Cable Plant – A New Alternative to GaAs," 2010 Cable Show Spring Technical Forum, May 11-13, Los Angelas, CA.

[6] Slowik, Fred, Jon Bordner and Raju Venkataraman, "5-4-3-2-1-0: A Model Approach to HFC Network Migration," 2009 SCTE Expo, October 28-30, Denver, CO.

[7] Young, Conrad, "CATV Hybrid Amplifier Modules: Past, Present, Future," http://www.rfmd.com/cs/documents/CATV%20Hybrid%20Amplifier%20Modules%20Past%20Present%20FutureWP.pdf

[8] http://www.freesunpower.com/radiation.php

[9] www.spheralsolar.com

Abbreviations and Acronyms

AP	Access Point
ARPU	Average Revenue per User
BAU	Business As Usual
CAGR	Compound Annual Growth Rate
CCAP	Converged Cable Access Platform
CMTS	Cable Modem Termination System
DTA	Digital Terminal Adaptor
Eff	Efficiency
EoC	Ethernet Over Coax
EPoC	Ethernet PON Over Coax
EPON	Ethernet PON
FTLA	Fiber-to-the-Last-Active
FTTH	Fiber-to-the-Home
GaAs	Gallium Arsenide
GaN	Gallium Nitride
Gbps	Gigabits per Second
HFC	Hybrid Fiber-Coax
hhp	Households Passed
HSD	High Speed Data
kwh	Kilowatt-Hr
LED	Light Emitting Diode
MEF	Metro Ethernet Forum
OFDM	Orthogonal Frequency Division Multiplexing
PAR	Peak-to-Average Ratio
PAPR	Peak-to-Average Power Ratio
QAM	Quadrature Amplitude Modulation
QFHD	Quad Format HD (aka 4k HD)
RF	Radio Frequency
SDV	Switched Digital Video
WDM	Wavelength Division Multiplexing
YoY	Year-on-Year

Appendix – Analysis Spreadsheet Snapshot

	2012 HFC N+3	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028
Mileage	4.16																
Homes/Mi	120																
Homes in Node Serving Area	500	500	500	500	500	500	500	500	500	500	500	500	500	500	500	500	500
4x4 Segmentation (125 Homes / Node)	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Node +1 (50 Homes/ Node)	0%	0%	0%	0%	0%	0%	35%	70%	100%	80%	60%	50%	40%	30%	20%	10%	0%
% Node +0 (26 Homes / Node)	0%	0%	0%	0%	0%	0%	0%	0%	0%	20%	40%	50%	60%	70%	80%	90%	100%
% Node + 0 Re-plumbed (26 Homes / Node)	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
# Enet Switches	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
% WiFi Coverage	0%	0%	0%	10%	15%	15%	20%	25%	30%	40%	40%	40%	40%	40%	40%	40%	40%
1.2 GHz Upgrade	N	N	N	N	N	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Power Factor Corrected PS	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
Activate 1.2 - 1.5 GHz	0	0	0	0	0	0	0	0	0	0	0	0	1	1	2	2	2
Solar Power BLNs	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
Power Draw in Amps	17.45	18.57	19.16	20.72	21.50	21.50	27.41	29.68	31.74	34.58	35.87	36.51	38.73	39.38	41.60	42.24	42.89
Power Increase	0.0%	6.4%	9.8%	18.7%	23.2%	23.2%	57.1%	70.1%	81.9%	98.2%	105.6%	109.3%	122.0%	125.7%	138.5%	142.1%	145.8%