With Great Power Comes Great Electricity Bills

Reducing grid dependance of the Access Network as it evolves toward 10G and beyond

A Technical Paper prepared for SCTE by

Tobias Peck
Sr. Director of Broadband Product Management
EnerSys
3767 Alpha Way, Bellingham, WA 98226
(360) 392-2247
tobias.peck@enersys.com

Table of Contents

List of Figures

List of Tables

1. Introduction

The coaxial network that is the backbone of the Cable Broadband industry was designed with a simple, ingenious powering scheme that allows multiple "actives" to be powered from a single point of connection to the utility grid. This tried-and-true powering strategy allows operators to minimize their points of interaction with an unpredictable utility grid and create a more reliable HFC grid to power network actives that support customer connectivity.

However, unlike the utility grid, the coaxial network's primary functions are both optimizing the flow of data to customers and transporting energy. As we continue to move toward 10G and beyond, exponential progress has been made in speed and volume of data that can be delivered over HFC, often by developing new architectures that increase dataflow with only minimal increases to power. But, as our industry seeks to be more sustainable and reduce operational energy spending, what should we be doing to improve power usage efficiency? Are there network efficiency gains that can be applied to all architectures to reduce energy usage and further improve the amount of data per dollar that the network can deliver?

This paper will first give a brief overview of network powering as we move toward 10G and beyond to understand the urgency of developing more efficient network powering. From there, the goal will be to map the power consumed by the network from the point of connection to the utility, through every transition point where the energy is transformed, transported, split, or ultimately consumed, to identify where energy is being used without creating value for customers. Finally, with our map in hand, we will discuss where efforts can most effectively be applied to reduce energy consumption and associated cost from the Access Network.

2. Powering Challenges into the Future

One operator recently commented that their outside plant (OSP) coax installed in the 1980s will meet DOCSIS 4.0 performance requirements with only passive component upgrades (splitters and taps). Similarly, OSP powering elements such as enclosures, power inserters and ferro-resonant transformers can function for decades. And while OSP power electronics and batteries will require expected periodic replacement resulting from normal use, the components that currently comprise coaxial network can support network demand into the foreseeable future. The question that needs to be addressed is how the OSP network can deliver power to the network more efficiently and sustainably. To fully understand the impact of the OSP on overall industry power usage, the SCTE Energy 2020 subcommittee recently updated the "Power Pyramid" showing that OSP power usage makes up nearly half of the overall consumption in the Cable Broadband industry.

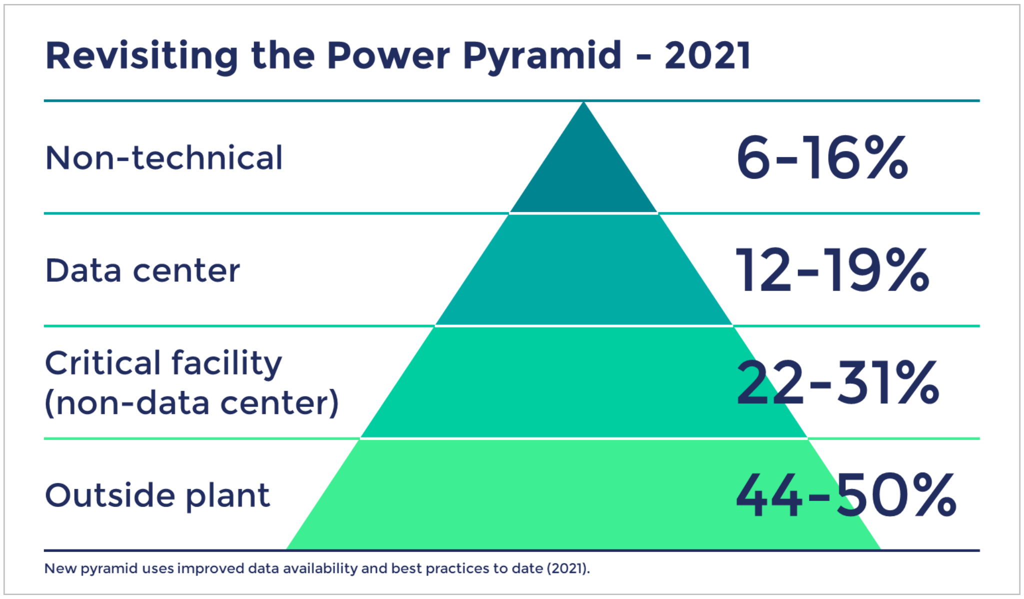

Figure 1 - Cable Operator Power Consumption Pyramid[i]

The urgent need to increase power efficiency is inherently connected to not only an understanding of the past and present of the HFC network, but also the future. The trek to 10G requires us to understand what the network of tomorrow will look like in order to plan for optimum power efficiency. Access network evolution has created unique powering challenges that must be addressed to assure that network powering upgrades will endure for decades to come. Let's review some of our powering challenges from the lens of the evolving access network with some help from CableLabs.

2.1. Powering 10G Vision

The CableLabs 10G initiative is the catalyst behind several technology innovations designed to deliver future proof internet speeds up to 100 times faster than most consumers are experiencing today. 10G aims to provide 10Gbps symmetrical, secure, low latency data services. 10G innovations will affect every aspect of the broadband network including headends, the access network and the customer's premises. The access network specifically must undergo enhancements to support new performance levels. Underlying technologies used to move network performance towards and beyond the 10G vision requires power. Assuring the availability of additional, reliable, and intelligent power for the 10G capable network is both essential and challenging since network architectures are evolving and much of the 10G enabling technology is still being developed. To approach this dilemma, let's review the 10G architecture vision as it exists today to ensure we're planning the appropriate powering infrastructure to support this near-future vision. For our 10G powering discussion, refer to this CableLabs 10G network architecture shown in Figure 2.

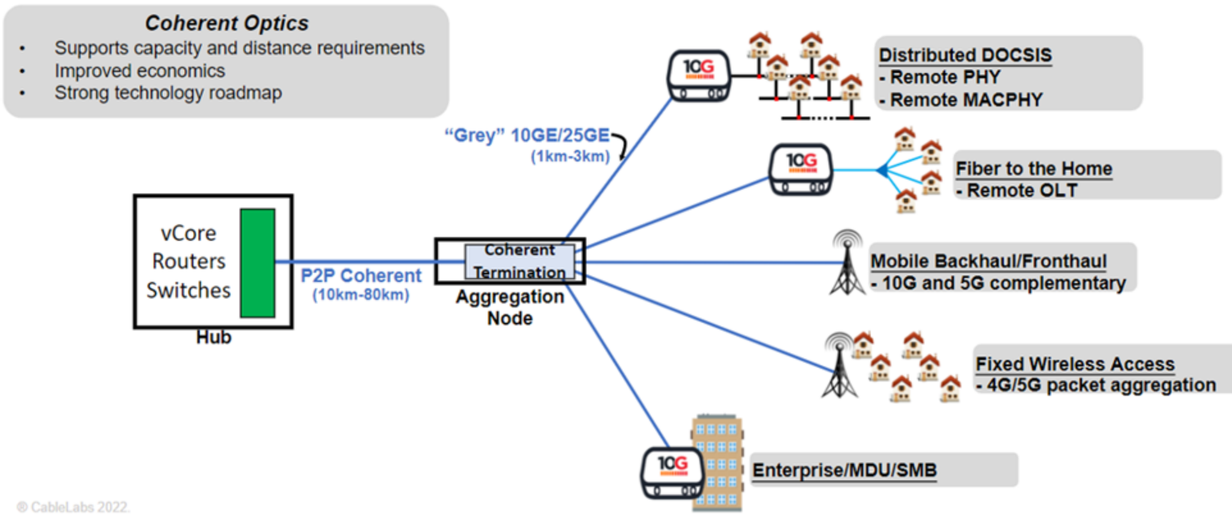

Figure 2 - Diagram of the Access Network in the near future

This near-future network contains several new concepts all combined into a single diagram. Let's review these network elements through the lens of power.

Aggregation Node: Fiber leaving the headend is routed to a new network component, the aggregation node (AN). In this architecture, new high speed fiber communications are pushed deeper into the network by use of high-speed coherent optics between the hub and the new AN. This link could be 50Gb, 100Gb or faster. The AN will act as an optical Ethernet switch providing 10Gb (or higher speed) optical links to downline elements such as DAA nodes and OLTs. The result is higher capacity services deeper into the network and closer to subscribers. From a powering vantage point, we've simply added another remote active element to the mix. Although the AN concept is new, one likely implementation would be a clamshell enclosure with fiber ports for upstream and downstream links, with a traditional node power port connecting to a standard 90VAC HFC power source. Once AN units begin to be deployed, power requirements will need to be analyzed. Given today's state of coherent pluggable modules combined with known high-speed switching elements, we anticipate that the AN may require approximately 200W in a fully configured state. This estimate remains highly speculative until actual architectures using AN prototypes are deployed.

DAA Nodes: Today's DAA nodes are fed from 10G Ethernet backhaul. Future DAA nodes may utilize 50G, 100G or faster backhaul links from an AN. These higher capacity links would be implemented using pluggable coherent optic modules. DAA node outputs support four DOCSIS 4.0 RF QAM channels over coax. Future looking, fully loaded DAA nodes may approach 180W power requirements.

R-OLTs: PON networks are deploying at increasing rates. The RDOF initiative is contributing this increase. Today's strand-mount R-OLTs support up to 512 subscribers using either 10G-EPON or XGS-PON protocol. Higher speed PON architectures are being discussed. 50Gb coherent optics support is planned for next generation PON implementations. Current R-OLT units can consume up to 140W. Future 50Gb PON R-OLTs will likely require up to 165W.

Wireless Backhaul: Wireless backhaul for 5G using the existing HFC infrastructure is expected to grow significantly in the near term as 5G scaled rollouts continue to increase. Coax powered radio units (RU), access points (AP) represent a significant incremental power draw to the coax plant.

Enterprise/MDU: high-capacity enterprise applications as well as large MDUs have historically used P2P dedicated optics for backhaul. DOCSIS 4.0 may be used for many new commercial applications when coax already exists in the region. New ANs providing P2P services or utilizing 10G or 50G PON to commercial and MDU customers are anticipated in the future.

Having this vision in mind of the near-future HFC network and its ability to support the data and backhaul demands for 10G and beyond, let's now discuss how to make the future more energy efficient.

3. Power Mapping the HFC Grid.

Reducing energy use and OpEx spend from the access network starts by understanding where power is used in the plant. More specifically, identifying and reducing or eliminating energy waste is the key to optimizing the effectiveness of the network. By mapping the flow of power from the point of connection to the grid and cataloging how each Watt of energy is consumed or wasted, we can develop data-driven strategies to optimize network sustainability and significantly reduce utility spend.

As we look to understand how to save energy, it is important to keep in mind these simple conversions based on the US average utility cost and CO2 impact per Watt:

1 Watt of power savings = 8.76 kWh annually

8.76 kWh @ $0.119 per kWh[ii] = $1.04 of savings = 7.45 Lbs of CO_2[iii]

So, every Watt of savings yields a dollar of annual savings on energy bills, with the potential for 2 to 3 times that in regions with higher utility rates. This combined with the fact that there are more than 750,000 power supplies in North America alone, help to emphasize the value that can be gained by reducing wasted power in the OSP.

3.1. Where does the power go?

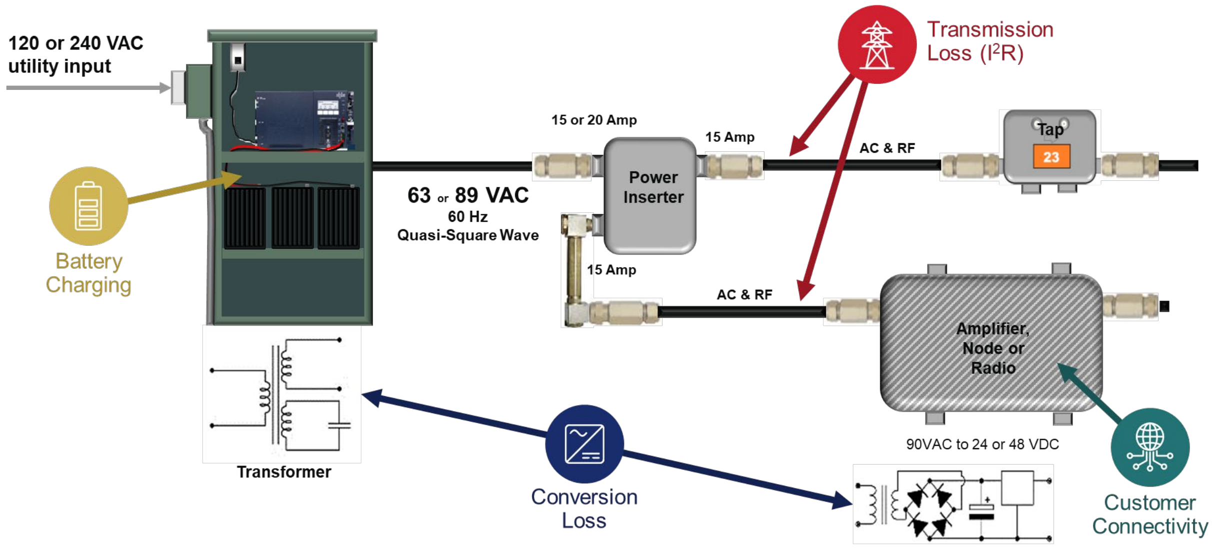

Figure 3 - Areas of power consumption in the HFC Grid

From the point of connection to the utility, power in the HFC network is essentially consumed in 4 ways:

1. **Customer Connectivity** – Powering equipment necessary to transport data through the network and connect customers to the digital world requires energy
2. **Conversion Loss** – Converting electrical power to a different voltage or from AC to DC and back
3. **Transmission Loss** – Power loss inherent from moving power across the network (I^2R loss)
4. **Battery charging and management overhead** – power used for operation of the network power supplies including the cable modem (transponder) and battery charger

Understanding the reasons behind each of these areas of energy consumption and the locations where they happen will allow us to develop strategies to greatly reduce them. Because the plant is a complex electrical circuit, changes to any of the four elements of power consumption above can often have an additive effect. For example, conversion loss in actives can have an impact on transmission loss throughout the plant. We will discuss this in more detail later.

It is also necessary to consider power usage in the network wholistically, as in some cases small power sacrifices in the Access Network are made to make the network more robust or intelligent, and often are offset by significant power savings in other areas of plant operations. In order to better quantify the impact of each of the four types of power consumption for comparative analysis against operational impacts we will build a model network and observe the impact of changes to each of them. But first, let's dive a little deeper into each of these four types of power consumption in the network.

3.2. Customer Connectivity

The HFC network is the central nervous system of the connected world. The actives that extend the reach of the network or increase the speed and bandwidth of data that flows through it are the reason that the HFC grid exists. Significant effort is consistently being invested into network actives to enable them to do more with less power. While these efforts have significant impacts in reducing network power consumption, for the purposes of this paper we will view their power draw as a mathematical given within our power calculations, understanding they are a non-negotiable within the network. We will discuss potential savings from maximized conversion efficiency within actives as a piece of Conversion Loss.

3.3. Conversion Loss

When converting electrical power from one form to another there will always be some power loss. The method by which that power is transformed can have a significant impact on power consumed in the HFC network. However, any conversation around transformer efficiency needs to be had with a more wholistic view considering effects on plant resiliency and operating costs. Currently in the HFC network there are 3 main points of conversion where loss needs to be considered.

3.3.1. Converting utility voltage to plant voltage

The first point of conversion is at the utility input to the power supply that feeds power to the plant. At this point utility power at either 120VAC or 240VAC is converted to plant voltage of 60VAC or 90VAC. In the overwhelming majority of powered coaxial plant in the world this conversion is done by a power supply with a **ferroresonant transformer.**

3.3.1.1. A Brief description of ferroresonant transformers

“Ferroresonance” is a phenomenon associated with the behavior of iron transformer cores, which in the instance of broadband power supplies, have two separate sections. The input core section is designed to prevent saturation at the maximum utility input so it cannot trip the input circuit breaker. The output section is designed to saturate in a controlled manner to provide tight output voltage regulation. The core is designed to optimize the efficiency at full load and allow the saturation loss only at lower load for output voltage regulation.

Normally, a ferroresonant transformer causes distortion of the output sine wave shape at lower load to regulate the output voltage. Ferroresonant transformers also have a magnetic shunt between the input winding and the auxiliary secondary winding, which is paralleled with one or more capacitors, forming a resonant circuit tuned to the power supply frequency. This resonant LC or “tank” circuit amplifies the input voltage and drives the output core to saturation to regulate the output voltage. In addition to providing isolation from input to output windings, this tank circuit provides additional filtering of input noise and resilience against high utility energy surges, such as lightning and industrial surges. The tank circuit also limits the output current during plant short circuit conditions to prevent tripping the input circuit breaker.

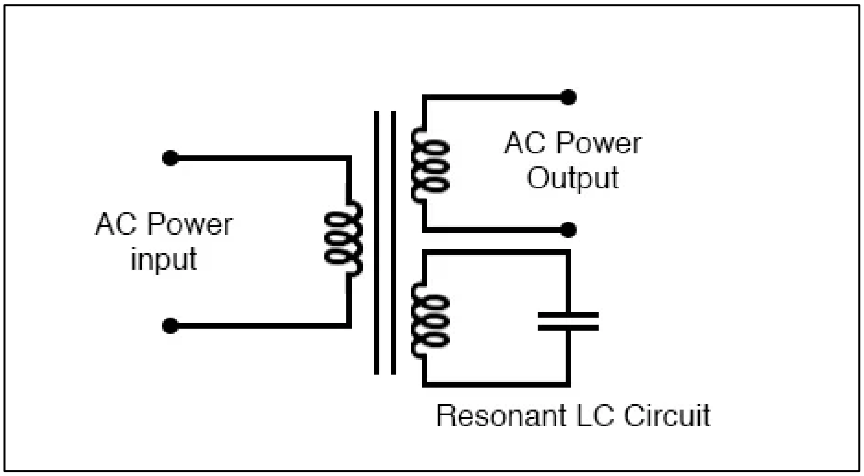

Figure 4 - Simple circuit diagram of a ferroresonant transformer

Because of their unique design, ferroresonant transformers provide several key benefits in powering the OSP network. First, they provide network equipment protection for surges and other transients as well as filtering out input noise that might otherwise be passed on to the plant. Additionally, they provide significant protection from shorts on the plant by current limiting to protect other equipment, and no-touch recovery once the short condition has been repaired. Finally, their relatively simple design, with minimal component count makes them extremely reliable and resilient to the uncontrolled environmental conditions in the majority of OSP locations.

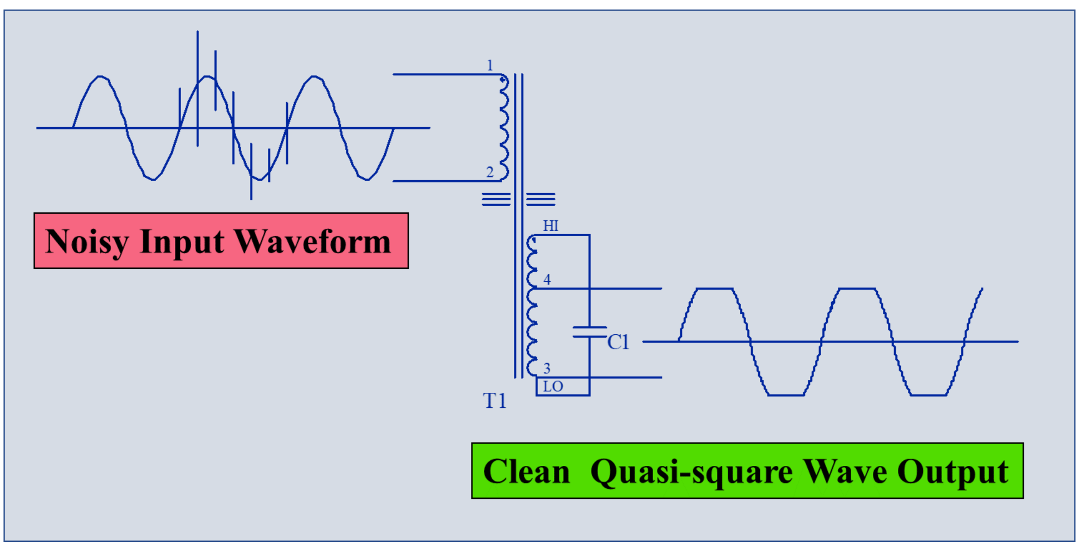

Figure 5 - Noise and transient reduction capability of a ferroresonant transformer

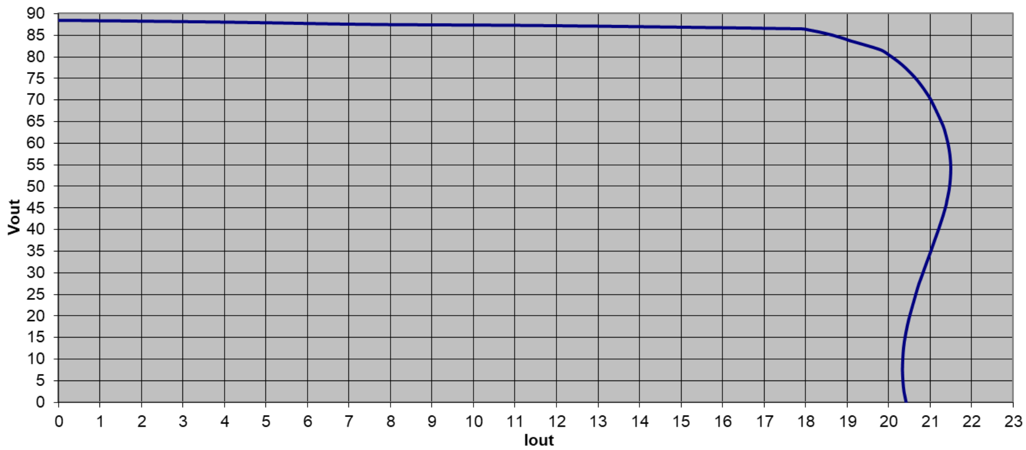

Figure 6 - Current limiting of a ferroresonant transformer during a short-circuit condition

However, with the additional power required for both the tank circuit, and the iron core saturation, the benefits mentioned above come at the cost of some efficiency loss. Since these elements of a ferroresonant transformer are optimized to handle a peak load, there is an inherent overhead to running these transformers that generally makes them significantly less efficient at lower loads. For example, in a power supply that is rated for up to 18A of output, there may be an inherent load of 100W or more to ensure resiliency up to peak load. At 3A of output at 90VAC or 270W of load, an additional 100W of inherent load means the transformer is operating at just under 73% efficiency.

Fortunately, much has been done over the past 30 years to gradually decrease the inherent load required for a ferroresonant transformer and make these devices more efficient. For example, a standard 15A power supply outputting 700W before 1999 would have had an efficiency of around 82.5 %. In the 2000s the standard 15A power supply was up to around 86% efficient at that load and today it is an additional 2% higher. Additionally, as the peak rating of a ferroresonant transformer is lowered so is the inherent load required to support that peak rating. Therefor a transformer that was rated for 5A, for example, would have an inherent load significantly lower than the 18A transformer.

3.3.1.1. Load Matching

It is important to note that when loaded to 50% or more of max rating, ferroresonant transformers can reach efficiencies of 85% to 94%. Knowing this, one strategy that could be deployed to reduce energy usage without losing the operational benefits of a ferroresonant transformer-based power supply is **Load Matching**. This strategy uses power supplies designed to tiered peak capacities in order to lower the inherent load on power supplies at lower amp draws. For example, if power supplies designed to peak loads of 5A, 10A, and 18A are deployed in an operator’s network, power supplies can be targeted to unique plant loads between 2.5A and 18A so that all power supplies operate at efficiencies above 85%, with many above 90%. While this strategy comes with some added management operationally, it could be an effective strategy to deploy during regular replacement cycles that shows moderate improvement for little added cost. We will model the potential impact of this strategy later on.

3.3.1.2. Linear and Switch-mode transformer efficiencies

While other transformer architectures, such as linear or switch mode, may be able to achieve conversion efficiencies of 95% or greater, these efficiency gains need to be weighed against either lost operational benefits, or the energy cost to add intelligence or redundancy into their architecture to achieve desired resiliency targets. For the purposes of this paper, we will model potential energy efficiency gains from increased efficiency of moving to these architectures, but thorough analysis should be done on real-world impacts of switching to a non-ferroresonant transformer.

3.3.2. AC to DC battery charging or DC to AC battery backup

The next conversion to be discussed briefly is the conversion in the standby power supply from AC to DC for battery charging, or DC to AC to power the plant during utility outages. These conversions have very little impact on overall power consumption since they happen very infrequently. And, while there is a constant float charge generally being applied to lead acid batteries in a standby power system to keep them healthy and ready for discharge, even in the worst-case scenario a 20% change in the efficiency of this conversion would yield less than a Watt of savings.

3.3.3. Conversion of plant power to DC power for actives

Another location in the plant where conversion loss is notable is at each active. Plant actives are generally powered internally by 48VDC or 24VDC which means that conversion from 90VAC plant power to the required DC voltage is necessary. Most actives do this conversion through small voltage conversion devices which, generally have low losses with efficiencies generally ranging from 95% to 98%. The presence of the ferroresonant transformer providing power to the network allows these higher efficiency transformers with less resiliency against transients and surges to be safely used in network actives. The efficiency of the these, and all transformers, does tend to decrease with age as transformer insulation wears down and more energy is lost to heat. Reducing network impact by ensuring maximized transformer efficiency of actives is key to reducing wasted energy in the network, as inefficiency of actives can also create additional transmission loss in the plant.

3.4. Transmission Loss

Transmission loss refers to the power lost from pushing current down the coax to feed actives, stated mathematically by Ohm's law, and it is one of the most important factors to understand in order to reduce the power consumption of the OSP. Depending on the particular section of plant, transmission loss can account for up to a 25% of the power consumed in the plant.

Since the plant is, from a powering standpoint, an electrical circuit with a series of loads drawing current and a series of resistors between them, in a single span of coax feeding an active we can express the transmission loss by Ohm's law as:

$$P(loss) = I^2R$$

Where:

P(loss) = power lost from coax line resistance, measured in Watts.

I = current through the cable required to feed an active, measured in amps

R = resistance of the length of cable, measured in Ohms

To understand the total plant load and how and where transmission loss occurs, it is also important to remember that active loads on the plant are constant power devices, or again from Ohm's Law

$$P = IV$$

Where:

P = power required for the active to function

I = current through the cable required to feed that active, measured in amps

V = voltage feeding the active

While these concepts are fairly simple mathematically, the interaction between them and a number of properties of the plant make them more complex to model in the plant.

3.4.1. Modelling power transmission loss

The layout and makeup of the HFC plant varies greatly from section to section. And while the number and type of actives in the plant are the most significant factors in the power draw of the plant, the coax and the passives have an impact on transmission loss.

3.4.1.1. Coax Loss

Everywhere that the plant is not drawing power it is resisting its flow. The resistance of any span between two actives is generally determined by three factors which, again, vary greatly from plant to plant – the length of the coax, the diameter of the center conductor of the coax and the passives in the power path. The coaxial resistance is fairly predictable as long as the cable size is known, which is not always a given especially when reaching toward last-leg actives. For modelling purposes, assuming that we know the length and diameter of coax, each span from active to active can be treated as a resistor with a certain ohmic value calculated from the table below.

Table 1 - Quick reference guide for coaxial cable resistance

P3 Cable Resistance	
Cable Dia	**Ohm/Ft**
0.5"	0.00172
0.625"	0.0011
0.75"	0.00076
0.875"	0.00055
625 PF	0.0003

3.4.1.2. Passive impact to transmission loss

Passives such as taps, splitters, directional couplers and power inserters are less impactful to plant resistance, but also less understood. Their per length resistance is inherently greater than that of the coax that surrounds them in the plant, but the value of this addition resistance is not well-documented and therefor hard to quantify. A cursory check of resistances for several random passives recently showed resistances from input to output ports of between 10 and 35 milliohms, which based on the table above is the equivalent of an additional 10 to 50 ft of coax depending on diameter. Assuming that approximately one passive is in each 200 ft. span of coax, this means that passives could add up to somewhere between 5% and 10% to the overall resistance of the plant.

To be clear, these results do not come from a well-organized, systematic study of a large sample of various types and manufacturers of passive equipment. They are an initial glimpse at a potential opportunity for plant savings that should be further understood. With the need to update passives to handle higher frequencies required by DOCSIS 4.0, this could be an opportune time to determine a resistance specification for passives that is not cost prohibitive.

3.4.1.3. Constant power actives

The vast majority of active equipment in the access network today is of a constant power nature, or simply the input current and voltage can vary as long as their mathematical product provides the requisite power for the device to operate. This is not to say the power draw of actives cannot vary moment to moment due to changes in throughput required. The principles of 'constant power devices' dictate that any change in voltage is offset by an inversely proportional change in current. Thus, if a 90W load initially is receiving 2A of current at 45V, and the voltage on the plant is doubled, the current through the lines will be halved. For the HFC network this has the added benefit of reducing the voltage drop through the coax, thereby reducing power loss. As you can see, this process would continue to occur until the changes in voltage and current were immeasurable and the plant voltage "settled." In the plant this settling of voltage happens almost instantaneously. For a mathematical model of the plant, it can be more time consuming to calculate without the use of some additional calculating power.

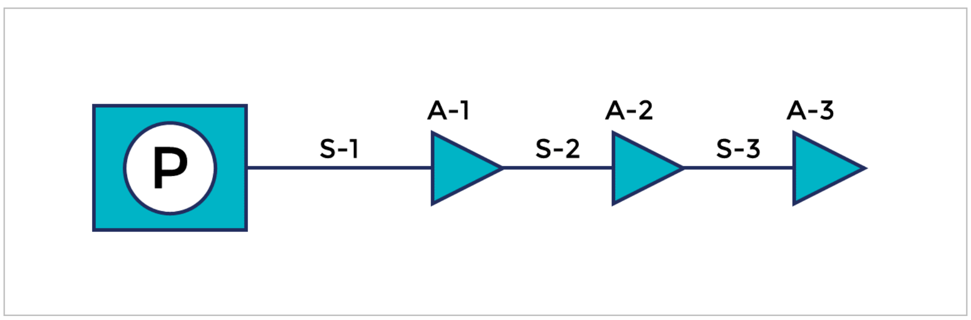

Figure 7 - Sample network for power loss calculation

Consider the simple section of plant above. To calculate the total power draw from this section of plant you would sum the power drawn from each active (A-1,2,3) and the total transmission loss from each span (S-1,2,3.) Before calculating the loss in span S-2, you need to calculate the loss in S-3 as the power loss in that span is a factor in determining the amount of current that is drawn through S2. The same goes for S-1 where you need to calculate the loss in S-2 and S-3 before you can calculate that span.

So first:

$$P_{(\text{loss in S-3})} = (P_{(A\text{-}3)}/V_{(A\text{-}3)})^2 \text{ x } R_{(S\text{-}3)}$$

Then:

$$P_{(\text{loss in S-2})} = ((P_{(A\text{-}3)} + P_{(A\text{-}2)} + P_{(S\text{-}3)}) / V_{(A\text{-}2)})^2 \text{ x } R_{(S\text{-}2)}$$

Finally:

$$P_{(\text{loss in S-1})} = ((P_{(A\text{-}3)} + P_{(A\text{-}2)} + P_{(A\text{-}1)} + P_{(S\text{-}3)} + P_{(S\text{-}2)}) / V_{(A\text{-}1)})^2 \text{ x } R_{(S\text{-}1)}$$

For purposes of power modelling the plant later in this paper a computer model was built that sets the end-of-line (EOL) voltage at the last active to a relatively low value then calculates loss back to the point of power. The model then checks the voltage against the known output of the power supply then repeats the process iteratively by raising the EOL voltage incrementally until the calculated voltage at the power supply matches the known output voltage. From there all transmission loss in the plant can be summed and analysis performed on the impact of any power saving strategy.

3.4.2. Concepts for voltage increase in the HFC

The proportional relationship between voltage and current with regard to power opens up interesting opportunities for reducing transmission loss by raising the voltage of the plant to lower the current, and therefor transmission loss, through the coax. First, there is still roughly 25% of plant in North America that is 60V. We will show later in our model that raising the voltage of plant from 60VAC to 90VAC can yield a significant reduction in transmission loss. Second, the National Electric Safety Code (NESC), Section 2 (excerpt in Figure X below) allows up to 150VDC in the communication space and while there are details that need to be better understood with DC powering of the HFC plant, there may be some relatively simple solutions that could allow the voltage to increase to that level. This increase would not only reduce power draw on the plant, but would also increase the reach of the plant and potentially allow for the reduction of operational energy usage by reducing the total number of sites to be maintained. It could also enable easier deployment of wireless technologies powered from the HFC grid. We will model all these scenarios below to better understand their impact on reducing utility power draw.

1. communication lines.

a. **located in the communication space.** The conductors and their supporting or containing structures, equipment, and apparatus that are used for public or private signal or communications service, and which operate at potentials not exceeding 400 V to ground or 750 V between any two points of the circuit, and the transmitted power of which does not exceed 150 W. When operating at not more than 90 V ac or 150 V dc, no limit is placed on the transmitted power of the system. Under specified conditions, communication cables may include communication circuits exceeding the preceding limitation where such circuits are also used to supply power solely to communications equipment. Fiber-optic cables are considered as communication lines, regardless of whether they are installed in the communication space or supply space in accordance with applicable rules,

NOTE: Public and private telephone, telegraph, railroad-signal, data, clock, fire, police-alarm, cable-television, and other systems conforming with the above are included. Lines used for signaling purposes, but not included under the above definition, are considered as supply lines of the same voltage and are to be so installed. Traffic signal light lines are considered as supply lines, not communication lines.

Figure 8 - NESC section 2, excerpt defining allowable voltage in the communication space

3.5. Battery Charging and Management Overhead

Standby powering has made the HFC grid an extremely resilient architecture. There are two key elements required to make standby powering such a robust architecture that draw additional power.

3.5.1. Power use from battery charging

Standby power requires batteries, battery charging can use unnecessary power if not properly managed. Typical gel cell VRLA batteries require a trickle/float charge of up to 3-4W constantly to stay healthy.

This is necessary to negate the effects of self-discharge and ensure batteries are topped off and ready for an outage. As batteries age and internal resistance increases the power used for float-charging these batteries will continue to increase. Therefore, it is important to ensure that batteries are properly maintained so that aging batteries with higher internal resistances, which are often no longer able to provide adequate back up time, are replaced regularly.

Some newer advanced AGM lead-acid batteries present a small opportunity for power savings. With lower internal resistances and self-discharge rates, these batteries have the ability to "rest" without being charged for up to 75% of their life, reducing the power needed to float charge them. Lithium-Ion batteries have extremely low self-discharge rates and, therefore, do not require float charge at all. There are small power gains of around 4 Watts to be had here, but as a part of a broader strategy it could help add up to larger total reduction of power.

3.5.2. Power use from DOCSIS status monitoring

Remote DOSCIS status monitoring of network power supplies has helped to greatly improve power reliability in the OSP by providing insight into potential issues before they become customer impacting. As our ability to extract additional data and perform more advanced analysis on these devices continues to improve, plant operations practices will become increasingly more streamlined. While these devices have become mission-critical to plant operations, they do have some draw that should be noted. Like other cable modems, the majority of devices currently in the OSP have an average draw of 5 to 7 Watts which will increase as networks advance and require higher-powered DOCSIS devices.

4. Modelling the impact of loss factors in the HFC grid

To understand how power is used in the HFC plant and more importantly where it is being consumed without adding value, we can look at an average section of plant and map the power flow. Then we can run mathematical scenarios based on targeting each type of loss and measure the impact of each strategy.

As mentioned earlier, the layout and makeup of the HFC plant varies greatly from section to section. This exercise is designed to show potential impact from loss factors, but the weighting of these factors may vary slightly depending on architectures or location of the plant– rural, suburban, urban –being viewed. The average power supply in North America has an output load of around 650W-750W so for the model it is ideal to find a section of plant with an output load within that range. Fortunately, in order to find this model section of plant all I needed to do was take two aspiring young cable engineers on a walk around my neighborhood to assist me in mapping the power flow from my local power supply. They were very engaged and helpful, until we reached the node which had an ice cream truck parked in front of it.

Figure 9 - Pictures of the model plant and the ice cream truck where I lost my assistants

For the sake of simplified mathematical modeling I used some generalized estimations of power draw for various actives and calculated the output power from my local power supply to be approximately 720W, which just happens to fit nicely into the average load range. I've mapped our model network in **figure 10**, but for the purpose of this exercise I have arranged the elements to more effectively show power flow instead of a more standard optical and RF-path layout. It is also important to note that, while loads were based on data from a cursory review of real actives of varying models and capabilities, the accuracy of the plant active power estimations used for modelling power draw is somewhat irrelevant due to the broad variation of active loads that the plant can support today and will support in the future.

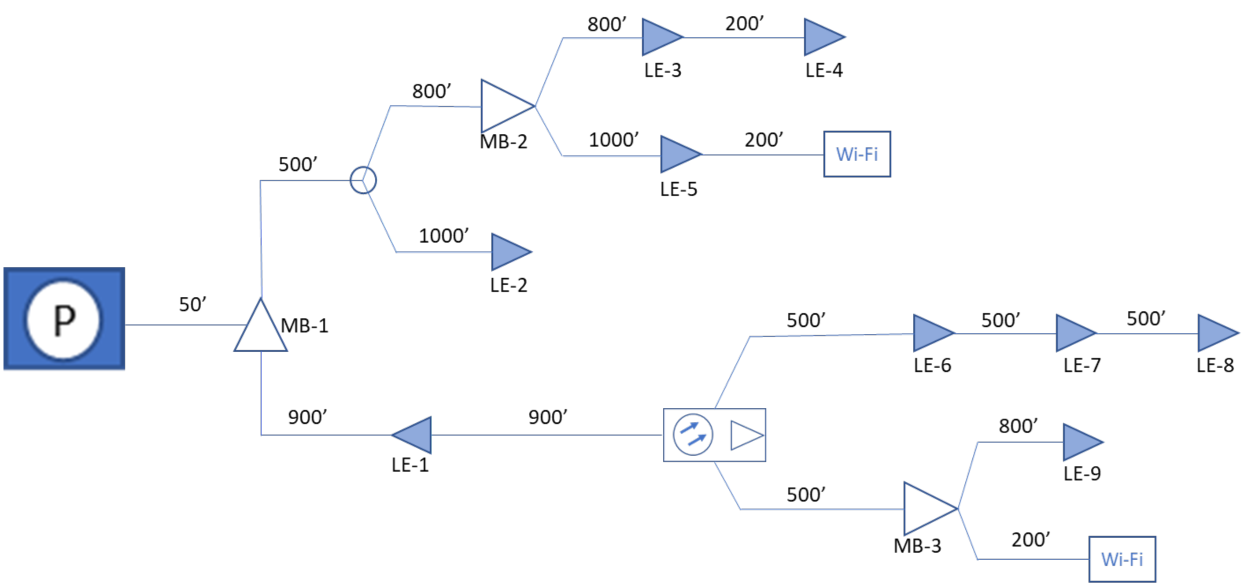

Figure 10 - Network power diagram for scenario modelling

4.1. Baseline assumptions for modelling

Below are the baseline assumptions that were used to build our initial power map.

1) Estimates for power draw of plant actives

Table 2 - Power draw for modeled plant actives

Active type	Power Consumed (W)	Quantity	Total Draw (W)
Legacy Node	90	1	90
WiFi Access	55	2	110
Amplifier (LE)	35	9	315
Mini bridger	50	3	150
Total Customer Connectivity Load			**665**

2) Per the chart above the total Customer Connectivity Load is assumed to be 665W.
3) All actives are assumed to have a conversion efficiency of 95% which is included in power consumption estimates.
4) Based on 2&3 above, the theoretical minimum power to run the network actives at 100% plant efficiency would be 631.75W.
5) Power supply conversion efficiency is assumed to be 86% based on a 15A power supply deployed prior to 2010 running at this load. This represents more than 55% of the deployed population of broadband power supplies in North America today.
6) Plant passives are assumed to add approximately 10% to the total plant resistance
7) Coax carrying power in the plant is assumed to be a 75/25 mix of .75 and .625 P3 cable to represent real-world variation in cable types.
8) Battery charging and management overhead is assumed to be a fixed 10W.
9) There is assumed to be 1 passive for every span or approximately 200 ft.

Inputting the baseline assumptions into the computer model, here are the baseline results:

Table 3 - Baseline power breakdown of modeled network

	Original Baseline
Required power for connectivity	631.75 Watts
Conversion Loss (Actives)	33.25 Watts
Transmission Loss	59.48 Watts
Battery Management Overhead	10 Watts
Conversion Loss (Power Supply)	119.57 Watts
Total Utility Power Draw	**854.05 Watts**

So, based on the above assumptions, my local network has a customer connectivity load 631.75W. This is the theoretical minimal power that is needed to keep the actives in this section of plant functioning if no other power was consumed in the plant. This is impossible to attain, but the goal is to get as close as possible with reasonable investment. The annual utility usage cost of this power supply based on the US average utility price of $0.16 per KWh is about $892.

4.2. Conversion Loss impact scenarios

4.2.1. Power supply conversion loss

Based on our baseline above, this is the most significant loss factor so let's look at the real-world impact of improving power supply efficiencies on conversion loss. We will start by looking at the impact of upgrading to a newer 15A power supply as a part of regular planned replacements. From there we will compare the impact of load-matching a more appropriate 10A power supply to this section of plant. Finally, we will look at the impact of a theoretical power supply that can maintain a maximum efficiency of 94% across a wide range of loads. See the results in the table below.

Table 4 - Model of power consumption varying power supply efficiency

	Original Baseline	Newer 15A PowerSupply	Load-Matched 10A PS	Theoretical High-Efficiency PS
Required power for connectivity	631.75 Watts	631.75 Watts	631.75 Watts	631.75 Watts
Conversion Loss (Actives)	33.25 Watts	33.25 Watts	33.25 Watts	33.25 Watts
Transmission Loss	59.48 Watts	59.48 Watts	59.48 Watts	59.48 Watts
Battery Management Overhead	10 Watts	10 Watts	10 Watts	10 Watts
Conversion Loss (Power Supply)	119.57 Watts	100.16 Watts	80.7 Watts	46.88 Watts
Total Utility Power Draw	**854.05 Watts**	**834.64 Watts**	**815.18 Watts**	**781.36 Watts**

Overall, strategies to increase power supply efficiency seem to have good results in our model network. Moving from an aging power supply to a Load Matched 10A power supply yields a 4.5% reduction in overall power consumption or an annual savings of around $40 annually. Any additional economically viable efficiency that can be gained toward the maximum theoretical efficiency above should obviously continue to be pursued as well.

4.2.2. Conversion loss in active devices

Maximizing the efficiency of power conversion in plant actives is not an easily deployable strategy as it requires multiple touch points for every section of plant. For example, in our model network there were 15 actives versus just one power supply. That being said, as we transition networks from current architectures to the near-future network discussed earlier, we must remember to be vigilant about ensuring maximum efficiency in plant actives as gains or losses out the plant have a cascading effect that you can see in the table below. This scenario models hypothetical increase in efficiency of all actives from 95% to 98%.

Table 5 - Model of improved power conversion efficiency of actives

	Original Baseline	Plant actives at 98% efficiency
Required power for connectivity	631.75 Watts	631.75 Watts
Conversion Loss (Actives)	33.25 Watts	12.89 Watts
Transmission Loss	59.48 Watts	55.52 Watts
Battery Management Overhead	10 Watts	10 Watts
Conversion Loss (Power Supply)	119.57 Watts	115.61 Watts
Total Utility Power Draw	**854.05 Watts**	**825.77 Watts**

This scenario is interesting as we get to see just how complex plant power is through the ripple effect that happens by changing the load on the plant. Gains or losses in conversion efficiency at the actives can create carry an additive change of up to 30% by impacting transmission loss and power supply conversion loss. Notice here that by reducing active load by 20W the rest of the plant load was reduced by an additional 7W, or almost an additional percentage point.

4.3. Transmission loss impact scenarios

In the prior scenario we began to see a small impact on transmission loss when adjusting slightly the amount of power drawn by actives. However, the best strategy to reduce transmission losses, as previously mentioned, is to increase the voltage at which power is delivered to actives and thereby reduce current flowing through the plant. In this scenario we will take a small step backwards to show what our model network would look like if it were running at 60VAC and then take a big step forward to the potential future scenario of powering at 150VDC.

Table 6 - Model of the impact of increased voltage on transmission loss

	Plant modeled at 60V	Original Baseline	Theoretical 150VDC
Required power for connectivity	631.75 Watts	631.75 Watts	631.75 Watts
Conversion Loss (Actives)	33.25 Watts	33.25 Watts	33.25 Watts
Transmission Loss	154.67 Watts	59.48 Watts	18.46 Watts
Battery Management Overhead	10 Watts	10 Watts	10 Watts
Conversion Loss (Power Supply)	135.06 Watts	119.57 Watts	112.89 Watts
Total Utility Power Draw	**964.73 Watts**	**854.05 Watts**	**806.35 Watts**
Voltage at last active	**47.2 Volts**	**79.8 Volts**	**145.3 Volts**

The biggest thing that jumps out in these scenarios is the power savings going from 60VAC to 90VAC, where we see the impact of transmission loss in older sections of plant and again the additive effect that loss can have on power supply conversion loss. Fortunately, adjusting the output voltage of the plant in most locations can be done in a few minutes without an upgrade to the power supply. Unfortunately, as many operators are already aware, in some of these sections of plant actives would need to be swapped out to handle higher voltages.

The theoretical 150VDC power supply does show significant power savings due to a big reduction in transmission loss. Another key detail to note in the model is the voltage of the last active in each scenario. As mentioned earlier, potentially the biggest gain from going to an increased output voltage like this is the ability to combine sections of plant and reduce power supply locations that need to be managed by plant operations. In 60V plant not only does current need to increase to deliver adequate power, but voltage decreases much more quickly and only needs to drop 18V (from 63VAC to 45 VAC) before it can no longer power actives. In 90V plant, voltage drop toward the minimum voltage for actives is much less pronounced, but would still be impacted significantly if the powering needs of two sections of plant were combined. Additionally, as the next generation of actives with higher power draw are deployed, the additional current required to power these actives will accelerate voltage drop and power loss beyond what we see in the scenarios above. In 150V plant, voltage drops so insignificantly, and the gap between power supply output voltage and minimum voltage for actives is so great, that two sections of plant could easily be combined without adding significant additional losses or risk of dropping actives due to low voltage.

4.4. Battery Charging and monitoring overhead

Assuming that the load of the cable modem used to monitor the health of network power is a given and that these devices are at maximum efficiency, reductions in this can be modeled based on a simple binary scenario. If we can, as a part of near-future network upgrades, leverage an energy storage technology that reduces or eliminates the need to float charge, we can eliminate 4 Watts of draw on the HFC network or the equivalent of 35 KWh annually. No need to create an elaborate model for this, however, there are other strategies to leverage the existing energy storage to create additional energy or utility cost savings that we be discussed in a moment.

4.5. Cumulative effect of power savings strategies

One last scenario to look at is the potential cumulative effect of implementing all of these strategies to have an idea of how much wasted energy we could theoretically remove from the plant. In this scenario, we will optimize all the variables to the limit of their realistic values.

- Increase conversion efficiency in actives to 98%
- Increase power supply conversion efficiency to 94%
- Increase line voltage to 150 VDC
- Reduce the additive resistive load from passives to 5%
- Remove battery float charging of 4W from the battery overhead

Table 7 - Cumulative effect of power savings strategies

	Original Baseline	Theoretical Best Case
Required power for connectivity	631.75 Watts	631.75 Watts
Conversion Loss (Actives)	33.25 Watts	12.89 Watts
Transmission Loss	59.48 Watts	16.48 Watts
Battery Management Overhead	10 Watts	6 Watts
Conversion Loss (Power Supply)	119.57 Watts	42.84 Watts
Total Utility Power Draw	**854.05 Watts**	**709.96 Watts**

While this model showing a 17% reduction in power draw is only theoretical, it is a realistic view of what is possible. It emphasizes the point that small cumulative gains will add up to significant progress in reducing our industry's footprint and drastically reducing operational costs. Using this average example as a baseline, if we could realize similar savings across all 750,000 power supplies in North America, our industry would be able to save nearly a terawatt hour of energy, 800,000,000 lbs of CO_2, and more than $110M of utility operational cost. In addition to that, raising plant voltage and increasing reach could save millions of additional dollars in maintenance costs.

The biggest question for many of these opportunities is how to make the solutions economically viable. To achieve many of the gains we've identified, investment is required and the payback period for any of these initiatives individually often makes them difficult to justify. One solution may be to pair these concepts for power savings into network upgrades already planned for the access network of the future, specifying key power saving elements into next-gen actives and passives. Additionally, prioritization of plant upgrades for next gen architectures could factor in opportunity for power savings at each sites so that gains are accelerated.

5. Other power reduction strategies

While significant gains can be made by targeting wasted energy in the plant, that strategy is limited. In our model example our maximized efficiency scenario removes 144W of grid energy and cost but still leaves 710W of dependency. To accelerate sustainability efforts it will be necessary to find additional ways to reduce energy draw from the grid. Here is a brief overview of two other concepts that should be considered to reduce grid dependence and operational spend.

5.1. Solar power augmentation

Renewable sources of energy bring the promise of sustainability and reduction of our carbon footprint. While many sources of renewable energy don't scale well to an OSP powering installation, solar energy not only scales effectively, but also could have the additional benefit of shading installations from excess solar load.

Over the past decade, the cost of an installed renewable energy system has gone down between 60% and 70%[iv],[v],[vi], to the point where the cost of solar energy is around $2.50 per Watt fully installed. This means that two common 370W panels would have an installed cost of less than $1500, including permitting. In most locations in the US these panels would generate around 3.2 KWh of power per day. With the average site drawing 700W, solar production could reduce overall energy usage by an additional 20% or around $140 of annual savings per site, with potential upside from more intelligent implementations. This could be scaled up to four or more panels where more space is available to cover larger loads.

Figure 11 - Example of a solar augmented cable power supply

5.1.1.– Options for solar implementation in the HFC plant

There are a couple of options for installation of solar in the OSP. The first is to simply leverage the real-estate of the plant to install an AC based grid-tied solar system and sell back energy produced by the panels to the utility company. This option is generally considered the simplest to implement, but it has several drawbacks. First, most utilities buy back at wholesale rates and not the retail rate that users generally pay so the payback model is reduced. Second, the power cannot be used for any other plant benefits like increasing runtime in an outage or charging batteries with free power after an outage. Finally, this type of "unintelligent" solar production has flooded the grid and created uneven demand for utilities. Thus, utilities have begun to discourage this scenario by removing incentives and increasing costs for permitting and metering of these sites.

An alternative, patented method exists to harvest solar energy locally at the point of load. This approach, which is shown below in **Figure 12**, creatively leverages the typical powering elements already existing in the OSP and does not involve any grid interactive utility permitting. By coupling the solar power directly onto the battery (DC) bus, the harvested energy can be utilized automatically for battery charging or extended runtime in an outage. By adding supplemental battery capacity, the system can also be used intelligently for offsetting utility power consumption when it makes the most sense (e.g. higher billing hours). The key element of the system is the master controller which can tailor solar consumption to various system sizes, standby runtime requirements, and utility billing scenarios to optimize the available solar resource and minimize utility billing while still maintaining critical backup capacity.

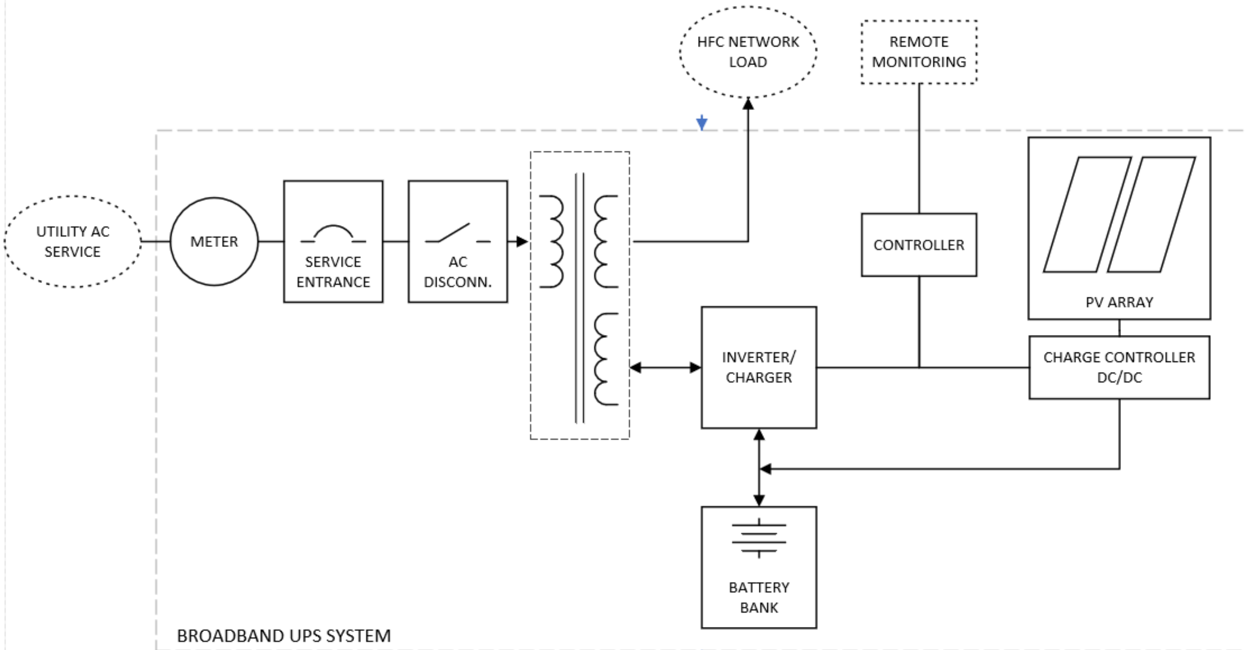

Figure 12 - Basic diagram of a cable power supply augmented by solar

5.2. Time-of-use mitigation

Large scale generation of energy is not inherently scalable, especially with cleaner sources of energy like solar and wind. When customer demand increases past a certain point, utility companies pay premiums to turn up secondary “Peaker Plants” to meet demand. Many large utility companies have rate structures that increase the price of energy during defined peak hours to encourage customers to reduce energy use during these times. This creates another interesting opportunity for reduction of operational energy cost as most OSP installations already have energy storage on site for emergency backup.

Business

Business Time Periods and Delivery Rates

	Peak	Off-Peak
Hours	8 a.m. to 10 p.m.	10 p.m. to 8 a.m. and all day on weekends
TIME-OF-USE DELIVERY RATES		
June 1 to Sept 30	29.38 cents/kWh	1.08 cents/kWh
All other months	14.47 cents/kWh	1.08 cents/kWh

Figure 13 - Example utility time-of-use rate structure

By slightly oversizing the battery array and deploying a battery with the capability for frequent cycling, batteries can be used to power the plant during hours with more expensive rates and recharged later with lower cost energy. For example, using the rates above a power supply drawing 700W from the utility would about $3 a day to power in the summer. By using batteries to remove it from the grid during peak rates for 4 hours and recharging later in the day when energy is inexpensive you could reduce your energy bill by $1 each day in the summer.

With major utility companies deploying time of use rate structures with peak rates as high as $0.57 per KWh this strategy should be explored further. Although, this strategy does not actually reduce total energy use, peak energy usage historically tends to be more carbon-intensive. Thus, this strategy does reduce carbon footprint, if not overall energy consumption. However, by combining time-of-use mitigation with solar augmentation and building intelligence into network power supplies to use the energy source that is most economically viable at any point of time - grid, solar, or batteries – This strategy can yield significant savings in power usage and operational energy spend.

6. Conclusion

As we progress along the path to 10G and beyond it is imperative that we continually seek out new solutions to reduce energy consumption of the network. There is significant opportunity to improve the way we power the access network and specifically reduce the inefficiencies in the plant. The possibility of leveraging renewable energy to reduce grid dependency of the OSP is becoming more economically viable. There is opportunity in front of us, however there is much work to be done.

The power map that we have created in this paper and the mathematical model used to analyze it are an early step to realizing energy savings. The opportunities identified above must each be investigated in depth and economically viable solutions developed to address them. More efficient power conversion, reduced transmission losses, and renewable energy augmentations all can have huge impact on energy usage. Even smaller changes like ensuring minimal resistive impact from network passives and reducing the need to float charge backup batteries can have a cumulative impact that can save millions of dollars and megawatt hours of energy. With significant network upgrades on the horizon to facilitate the path to 10G and beyond, the time is now to understand everything that can be done to reduce power consumption in the access network and ensure that near-future network is b energy efficient as it is fast.

Abbreviations

AC	alternating current
AGM	absorbent glass mat
AN	aggregation node
DAA	distributed access architecture
DC	direct current
HFC	hybrid fiber coax
MDU	multi-dwelling Unit
MSO	multiple-system operator
OLT	optical line terminal
OSP	outside plant
P2P	peer-to-peer
PON	passive optical network

RDOF	Rural Digital Opportunity Fund
UPS	uninterruptible power supply
VRLA	valve regulated lead acid
kWh Google	kilowatt hour

Bibliography & References

R. Anderson, "Network Power Considerations for 10G Enhancements", SCTE- Technical Journal, vol. 1, no. 1, pp.23-37, June 2021.Institute of Electrical and Electronics Engineers, and American National Standards Institute. National Electrical Safety Code C2-2017. 2017th ed., Institute of Electrical and Electronics Engineers, 2017.

---

[i] https://www.scte.org/standards/energy-2020/energy-2020-powering-cables-success/

[ii]https://www.eia.gov/tools/faqs/faq.php?id=74&t=11#:~:text=In%202020%2C%20total%20U.S.%20electricity,CO2%20emissions%20per%20kWh.

[iii] https://www.eia.gov/electricity/data.php
[iv] https://www.marketwatch.com/picks/guides/home-improvement/solar-panel-costs/

[v] https://www.solarreviews.com/blog/how-has-the-price-and-efficiency-of-solar-panels-changed-over-time

[vi] https://www.nrel.gov/news/program/2021/documenting-a-decade-of-cost-declines-for-pv-systems.html