Adaptive Power Management for Node Clusters

A Technical Paper prepared for SCTE•ISBE by

Fernando X. Villarruel
Architect, Office of the CTO
Cisco Systems, Inc.
5030 Sugarloaf Pkwy, Lawrenceville GA, 30444
villarf@cisco.com

Michael Mobley
Technical Leader
Cisco Systems, Inc.
5030 Sugarloaf Pkwy, Lawrenceville GA, 30444
mobleym@cisco.com

Curt Dalton
Software Development Mngr.
Cisco Systems, Inc.
5030 Sugarloaf Pkwy, Lawrenceville GA, 30444
daltonc@cisco.com

Lamar West, Ph.D.
President
LEW Consulting, LLC.
7840 Holly Springs Road
Maysville, GA 30558
lamar@lamarwest.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

Introduction

Arguably the biggest challenge for the HFC plant moving forward is implementing next generation technologies in nodes that operate in a limited power consumption environment. This, for example, is the backdrop for the evolution to distributed access architecture (DAA,) full duplex (FDX), and networking nodes which are expected to require more average, and larger ranges of power consumption. There is also a discrepancy between the higher power/thermal dissipation capability of new nodes and the maximal power that is allotted to them per multiple system operator (MSO) product specifications. This creates an opportunity for an analytics-based application of power management that maximizes the performance capabilities of nodes but at the same time keeps in check or reduces the consumption of their system power footprint.

Because of MSO desires to avoid adding new power supplies when deploying additional nodes as part of, for example, fiber-deep deployment, nodes are now part of a power consumption cluster, which is effectively a collection of nodes serviced by a power supply where ultimately the power envelope that matters is that for the collection of nodes in the cluster, and not of any single node in particular. In effect, nodes are reverting back to a more centralized powering schemes from their previous distributed powering architectures, where the power supply placement is often no longer optimal due to new nodes being added downstream of the original node. Significant additional losses in power due to the Joule heating or I^2R losses in the coax used to transport power to the new node locations are now added to the powering requirements of the new nodes themselves. The new power consumption means that many power supplies may become challenged to supply sufficient power as new devices such as wireless strand-mounted devices are added to the HFC plant.

Therefore, power sensory information of node function is now necessary. Sensors that are hardware parts can be added to nodes to make the reading and reporting of power state information possible. This sensory information can be collected and maintained centrally, within a general cloud infrastructure. Making such energy consumption information available from sensors allows for an analytic comparison and optimization of power and/or performance settings of elements in the cluster to optimize performance while improving energy efficiency. And finally, in keeping with the current trend for increasing intelligence in network operations, these sensors and associated data can be the inputs to machine learning algorithms that provide predictions and necessary decisions to optimize or evolve a system and thereby facilitate introduction of new or different elements into the cluster.

This paper is a novel proposition for an analytics-based application that manages the problem of wide deployment of new technology in nodes. In this paper we describe the hardware, sensory capability expected, the cloud-based architecture needed for data collection, storage and analysis, the logic applied to analytical engines, and the process for execution of optimal states in the presence of a broader policy mechanism, and finally the inclusion of a machine learning principles for integration of new technologies as nodes evolve.

Background

The HFC plant finds itself in a precarious position when it comes to the topic of power consumption. On one hand there is the recognized desire to minimize as much as possible the power consumption of the plant. The effort of the SCTE/ISBE Energy Management Subcommittee has documented the tremendous cost burden that power consumption will have for MSO's in the near years to come. With this in mind there has been a rally to gather minds and technologies around energy conservation measures. Simultaneously however, the ever-increasing numbers of ports required for more granular service groups, in DOCSIS and video, have led to the favoring of distributed architectures, where the PHY layer of the converged cable access platform (CCAP) is migrated to the outside plant. This evolution significantly increases the capacity of the fiber feeding the nodes by making the link digital, uses more efficient gallium nitride (GaN) technology that enables higher RF spectrum use, and can add to the intelligence of the node, but it can also dramatically add to the power consumption of these field-distributed endpoints. These trends are just at the beginning, where energy conservation measures will be increasingly important in the future as added capability of intelligent nodes is expected to be more prevalent, along with their need for more power to accomplish their functions.

Another dimension is the operational challenges with rearranging or upgrading the powering infrastructure of the plant. A wholesale revision would be prohibitively costly, and even impractical due to the nature of upgrading utilities. This dynamic leads operators to prescribe strict specifications for the power consumption of nodes because in the absence of any specific knowledge of the power state in a specific node situation or geography, it is the only guidance that can be given. For example, MSO's typically give a power consumption number for all nodes independent of function or location. As mentioned in the introduction, the deployment of fiber-deep architectures in particular with the constraint of not adding or moving power supplies often leads to new I^2R losses in the HFC plant that make power consumption per node a more variable quantity.

In this paper we propose the question what if that were not the case? What if a system was not blind to the power state of it parts? What if guidance about power consumption could be given in broader and more granular terms, and not just for isolated nodes? If this were the case then optimization of power consumption could be done in the context of performance. It could vary from place to place, and it could be updated over time. We develop this proposition in the next sections of this paper.

Node Cluster

We define a node cluster as a group of functional nodes that share one power supply. This distribution of parts is generally referred to as the centralized power supply model. (Note, the principles we present would also work with alternate power supplies and node relationships, but we describe the solution here with the most likely scenario.) Figure 1 is visual representation of a node cluster. We note fundamentally that the possibilities for node parts can include various technologies. We list some of the ones we know now, with the understanding that in the future there could be others. We include legacy analog optic/RF nodes as DOCSIS3.1 or D3.1 in the figure. We include DAA nodes as remote PHY and FDX. We also include packet processing nodes such as field aggregation routers, optical transport nodes (OTN) multiplexing transponders or Muxponders and optical line termination units (OLTs).

These nodes within a cluster can require a range of power consumption profiles. Some could be less than the guidance typically given for power consumption and some could be more. The commonality for all nodes is their shared line power supply. For this reason, the best guidance with regards to power capabilities and limitations is given by the line power supply. This is really the only power envelope that matters.

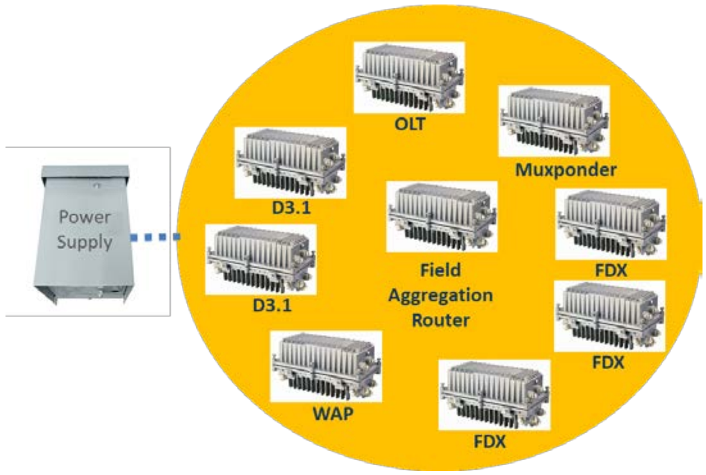

Figure 1 - Node cluster definition. Power supply and its node parts.

While the statement of the power supply envelope is true, at the moment the relationship of the power supply to its subtending nodes is not known. The reason is that nodes at the moment do not measure (or are not known to measure) or have a mechanism to report their power consumption. Third generation power supplies on the other hand have a way to accurately report power state, but have no way to compare their capabilities to the context of the usage of their power-consuming devices in the node. Thus along with the concept of a cluster we propose the basic principle of nodes that are able to report their power state information.

Sensors and Measurement

The inclusion of power management integrated circuits (ICs) to a node design allow for measurement and reporting of electrical current and power consumption. These ICs are typically small and are cost and power-consumption effective. They have the task of measuring current drawn at a range of voltages and also of reporting current or power. Depending on the model, one can measure several lines by toggling, of have dedicated measurement. The reporting structure is typically facilitated by an integrated micro that allows for two way communication via a rudimentary form such as Intra-Integrated Circuit (I2C). The data set that includes power information is not meant to live in the node, so it would have to be stitched into the data plane for signaling that is already being transported to service packet cores. We expand on this in later sections.

Figure 2 below shows an example for power measurement and reporting of a remote PHY node via a power monitor chip. Note that this example tracks the usage throughout the power tree. The node power supply is monitored at entry and exit, and the granular parts of the node are measured as well. From a data collection perspective the input into the node power supply is the minimum required power measured, but the granularity added to measuring various points of the node allow for a broader application of both power consumption and performance tuning. In the next sections we will show that understanding the usage through out the power tree within the node is quite useful, in particular for balancing node component setting for energy consumption and performance targets. Note that Figure 2 is an example of an HFC Remote-PHY node, but a similar approach can be taken for analog, or packet processing nodes.

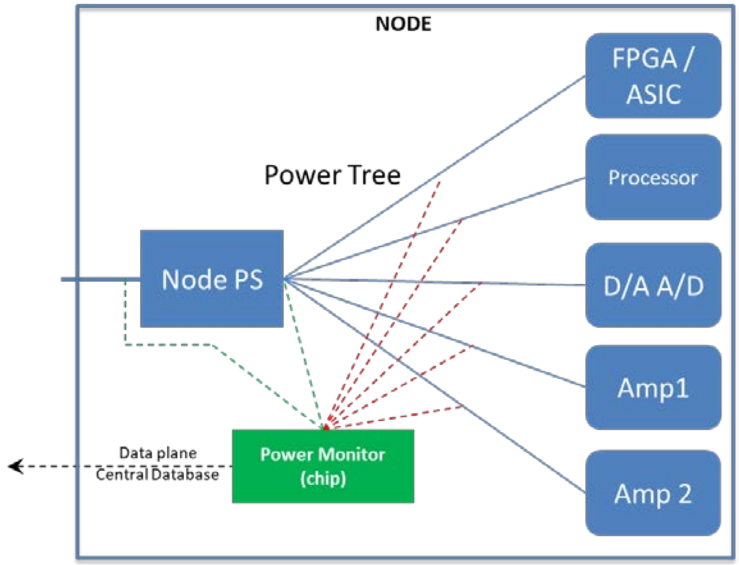

Figure 2 - Power measurement and reporting of elements within a DAA node.

Power Knobs for HFC RF Nodes

The power consumption of an HFC RF node can be influenced by various settings for the parts within the node. The settings are related to throughput and performance expectations of the node. However, it is important to understand that there can be multiple settings that achieve the same customer end-line performance or throughput from a node. Figure 3 shows the various levers or knobs we have available, for which through varying settings can affect the power consumption of the node. Unlike Figure 2 where we call out the node parts that can vary, in Figure 3 we focus on the signal settings and power supply settings of these internal devices that have power consumption effects. In this case it is not just a matter of power consumption but changes that can be done with the perspective of performance. We point out the modulation of the signaling, which with current and next generations of DOCSIS can vary greatly depending on network and end of line conditions. Varying the bandwidth is related to the settings the packet processor of a remote PHY device (RPD) for instance, as we see later on. Varying the RF bandwidth is related to the general settings for the plant of licensed spectrum being used. This of course can impact the necessary settings on both the RF management of the node and the packet processor. RF power refers to the absolute value of RF power being put out by amplifiers in the node design. Depending on the quality of signal needed at the end of line, (through RF drops, or related to the modulation setting, or the bandwidth allocation,) this value can change and can sometimes be maximized, while at other times not so much.

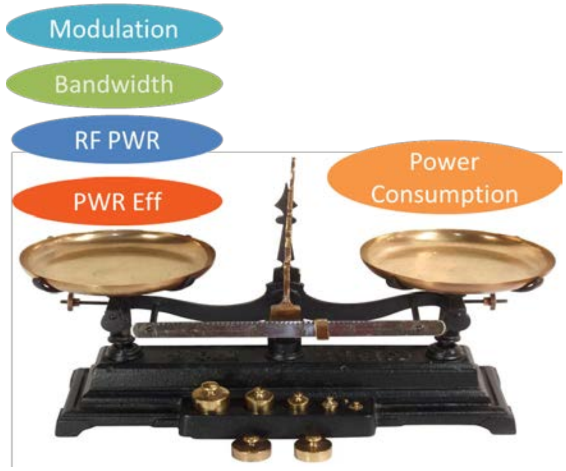

Figure 3 - Power consumption knobs for an HFC RF node

1. Example: RF FWD Amplifier

The forward path amplification structure of HFC nodes allows a good example to understand the dynamic of power consumption and performance. Figure 4 shows the typical forward (FWD) path amplification structure for a four port HFC node. There is a common pre-amplifier and four independent post amplifiers that provide RF output power for the four output ports of the node. The structure itself gives us the first type of power efficiency tactic. In the case where there is an inactive leg of the node, then the mere savings of turning off one post amplifier can be up to 20%. Similarly, if the accompanying RF structure allows for turning off of an amplifier during low traffic conditions (e.g. at night or in summer for university student residences) then those energy savings are also available for selected times of day/year.

The RF amplifiers are effectively the signal (and its noise) interface from the transmission network to customer premises equipment (CPEs), directly or via another set of amplifiers. The RF power value is tied to the needed carrier-to-noise ratio (CNR) in the RF domain and the related signal-to-noise ratio (SNR) after demodulation. Note that in this context, "noise" refers to the sum of thermal noise, interference, and composite intermodulation noise (CIN). Consequently, there is a limited RF output power dynamic range that provides an adequate carrier-to-noise ratio and the related signal-to-noise ratio. And as it turns out, signal-to-noise ratio or equivalently modulation error ratio (MER) for digital (QAM and OFDM) signals, which are used exclusively now, are tied to very particular expectations, where a signal-to-noise ratio that is too low is not usable, and if it is too high, it gives no extra benefit. This allows for the dialing in on a power range that is tied to a particular channel modulation, and by extension the number of channels that are used within a spectrum. As an example, we note that if an amplifier structure that is capable of signaling over the full RF range of DOCSIS 3.1 is operating at half capacity, either by reduction of RF spectrum or reduction of signal order of modulation (and thus lower required MER) it can then accommodate approximately 20% worth of power consumption savings. Note we do not specify any particular numbers as they would be product-specific, but the dynamic should be similar throughout the industry.

Another form of power reduction, a somewhat hot topic now is digital pre-distortion, (DPD), (Chong, 2018). The capability of DPD to help correct noise components from the driven profile of amplifiers is simplified in nodes with remote PHY devices which lend themselves to the digital signal processing (DSP) needed. Overall, accounting for the added DSP power, the saving of DPD can be up to 20% in very straight forward applications. There is potentially more added bonus, but this requires special calibrations and attention to details of loadings, etc.

Overall, it seems that 20% saving from the amplifier usage under reasonable circumstances is a very achievable target.

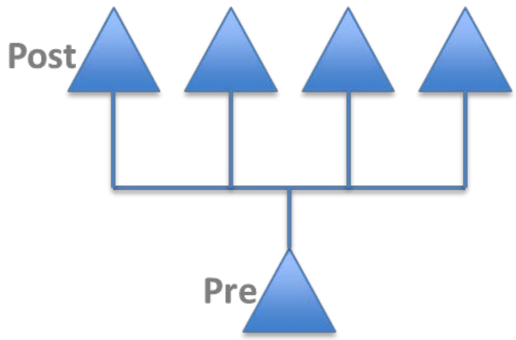

Figure 4 - FWD amplification structure for HFC RF Node

2. DAA Devices

A critical part to power consumption of new nodes is the silicon that is used for remote PHY or similar devices. These devices effectively take Ethernet / IP signals that were generated from a packet core and convert them to RF signals to be fed into the RF plant. At their core these devices have application specific integrated circuits (ASICs) or field programmable gate array (FPGA) technologies, or both which determine their functionality. Unfortunately, this vast amount of intelligence added to the node also comes with a power consumption penalty. This tradeoff is made with performance in mind as well, where ASICs are a naturally lower power solution but with little flexibility once finalized, and FPGAs have a higher power penalty with greater flexibility for evolution in the field. Because these devices are so power-hungry it is important to be able to understand and leverage any power saving capabilities. Without getting too into product detail, a mix of technologies, older and new versions, along with bandwidths, modulations and other ancillary functions dealing with versions of DOCSIS, the silicon within DAA devices can vary power by up to 50%. For this reason, the numerical understanding of power consumption (measured variations) for DAA devices is quite useful, if not necessary.

Load Variations and Cluster Power Savings

With the availability of varied power states in cluster nodes we turn our attention to how these variations materialize into savings to the overall cluster's power consumption. Connecting the cluster is a sequence of unique coaxial segments. Let's look at how the DC loop resistance of the coaxial cable affects the total power output supplied by the line power supply. This is a highly simplified analysis and is intended to illuminate the issue rather than give a rigorous final result. A rigorous analysis, such as that described in Mitchinson (Mitchson G., 2016) would be extremely complicated and is beyond the scope of what we wish to accomplish here. In order to make the current analysis practical we make the following simplifying assumptions:

1. The switching power supply in a node draws a constant power regardless of the input voltage. Actual switching power supplies vary slightly in efficiency as the input voltage varies. This results in a slight variation in power draw over the input voltage range. However, this variation is relatively small.

2. A group of four optical nodes are being fed from a single line power supply. These optical nodes are located relatively close to each other, yet all are a significant distance from the line power supply. For the purposes of this analysis we will simplify the things by ignoring the powering losses associated with the coax that interconnects the nodes and concentrate on the main coaxial span that connects the group of nodes to the line power supply. The effect of the interconnecting coax might slightly change the final value of the results. However, rigorous analysis of the node interconnections will complicate the mathematics well beyond what is required to demonstrate the nature of the relationships given below. This complexity of this analysis further illustrates the need for multiple individual power sensors located throughout the network in order to obtain empirical data.
3. We will assume that the rms voltage provided by the line power supply is the same as the peak voltage. This assumption is a reasonable one when the line power supply is under no load and the trapezoidal voltage output approaches a square wave. However, as the current from the line power supply increases its voltage waveform, it becomes more rounded, resulting in a decreased rms value. Were this change included in the analysis it would simply tend to reinforce the results that we will provide
4. We will not consider changes in the real power and apparent power as a result of the power factor of the switching power supplies in the node.

Consider the situation where the four optical nodes are connected to a 90-volt line power supply by a coaxial span with a DC loop resistance of 3 ohms, as shown in Figure 5.

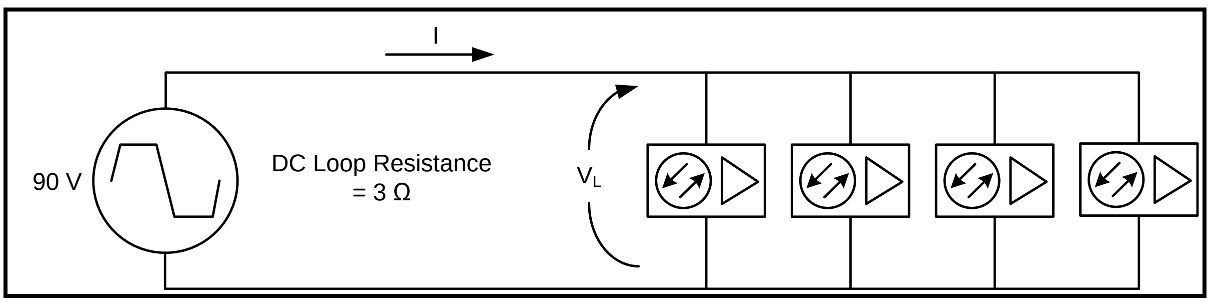

Figure 5 - Four Optical Nodes Connected to a 90 Volt Line Power Supply

For the purposes of this analysis we will assume that each optical node draws 75 watts. The four optical nodes will draw a total of 300 watts. We can calculate the current draw, I from the line power supply and the voltage, V_L, across the four optical nodes.

$$300\ watts = (V_L)(I) \tag{1}$$

Rearranging

$$V_L = \frac{300\ watts}{I} \tag{2}$$

Considering the voltage drop across the DC loop resistance of the coaxial cable we can calculate the voltage across the optical nodes, V_L, by

$$90\ volts - (I)(3\ \Omega) = V_L \tag{3}$$

Rearranging

$$(I)(3\ \Omega) - 90\ volts + V_L = 0 \tag{4}$$

Substituting V_L using equation (X.2)

$$(I)(3\ \Omega) - 90\ volts + \frac{300\ watts}{I} = 0 \tag{5}$$

Multiply through by I

$$(I^2)(3\ \Omega) - (90\ volts)(I) + 300\ watts = 0 \qquad (.6)$$

Solving for I using the quadratic equation

$$I = \frac{90 \pm \sqrt{\left((-90)^2 - (4)(3)(300)\right)}}{(2)(3)}\ amps \tag{7}$$

Resulting in the two solutions I = 26.2 A and I = 3.8 A. We will ignore the larger solution as this is more current than can be typically supplied by a line power supply and chose the smaller solution. Plugging into equation (X.3) we find the voltage at the optical nodes to be

$$V_L = 90\ volts - (3.82A)(3\ \Omega) = 78.54\ volts \tag{8}$$

We can also calculate the total power being delivered by the line power supply, P_{LPS},

$$P_{LPS} = (90\ volts)(3.82\ A) = 343.8\ watts \tag{9}$$

Note that 43.8 watts is being dissipated as heat by the 3 ohm DC loop resistance of the coaxial cable.

Now consider a second scenario. In this scenario node splitting has replaced the four optical nodes of the previous scenario with six optical nodes. Additionally, the six optical nodes contain advanced electronics such as DAA electronics and consequently each optical node dissipates 100 watts. This results in a total power utilization by the six optical nodes of 600 watts. This scenario is shown in Figure 6.

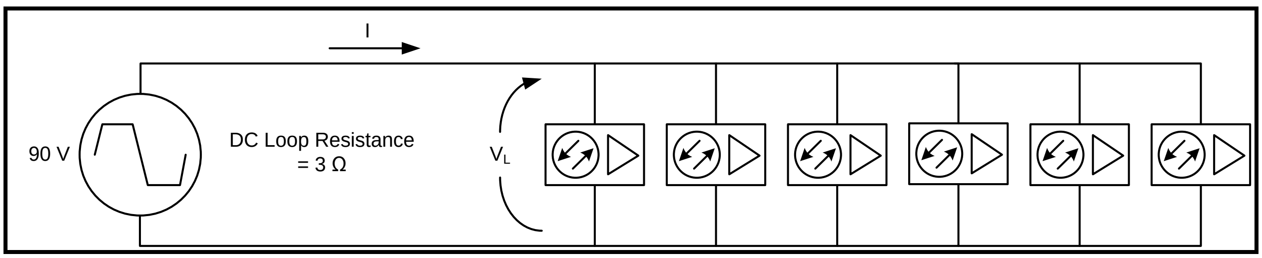

Figure 6 - Six Optical Nodes Connected to a 90 Volt Line Power Supply

It is possible to do an analysis of this scenario that is similar to the previous analysis.

$$600\ watts = (V_L)(I) \tag{10}$$

Rearranging

$$V_L = \frac{600\ watts}{I} \tag{11}$$

Considering the voltage drop across the DC loop resistance of the coaxial cable we can calculate the voltage across the optical nodes, V_L, by

$$90\ volts - (I)(3\ \Omega) = V_L \tag{12}$$

Rearranging

$$(I)(3\ \Omega) - 90\ volts + V_L = 0 \tag{13}$$

Substituting V_L using equation (X.2)

$$(I)(3\ \Omega) - 90\ volts + \frac{600\ watts}{I} = 0 \tag{14}$$

Multiply through by I

$$(I^2)(3\ \Omega) - (90\ volts)(I) + 600\ watts = 0 \tag{15}$$

Solving for I using the quadratic equation

$$I = \frac{90 \pm \sqrt{\left((-90)^2 - (4)(3)(600)\right)}}{(2)(3)}\ amps \tag{16}$$

Resulting in the two solutions I = 20 A and I = 10 A. We will ignore the larger solution as this is more current than can be typically supplied by a line power supply and chose the smaller solution. Plugging into equation (X.12) we find the voltage at the optical nodes to be

$$V_L = 90\ volts - (10\ A)(3\ \Omega) = 60\ volts \tag{17}$$

As compared to 78.54 volts in the previous scenario. The additional power loading has caused the voltage at the actives to drop by 18.54 volts. We can also calculate the total power being delivered by the line power supply, P_{LPS},

$$P_{LPS} = (90\ volts)(10\ A) = 900\ watts \tag{18}$$

When we compare the second scenario to the first we see that the power required by the optical nodes increased by a factor of 2 in the second scenario with respect to the first. However, the power required from the line power supply went from 343.8 watts to 900 watts due to the additional coax I^2R loss with the higher current. The power required from the line power supply increased by factor of *2.6*.

This analysis is not intended to present the results for any given specific architectural case. Rather, it is intended to illustrate the nonlinear relationship between the power required by active devices in the network and the resulting power delivered by line power supplies.

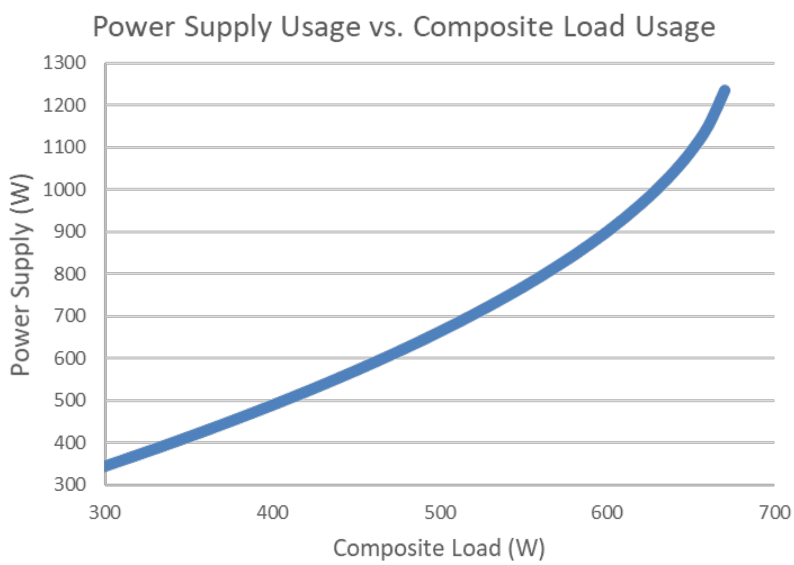

Figure 7 - Power supply usage over load variance for cluster example

This simplified analysis illustrates a more general conclusion. Figure 7 shows the dynamic between the varying loads and power supply usage for the example topology. We note that the power supply output varies increasingly faster than the load variations, particularly as the load reaches higher values. Most operators would prefer to continue to use existing line power supply locations and existing powering architectures wherever possible. But as new system designs and technologies increase the total power utilization of active devices in the outside plant by a factor of two in our example, the power required from the line power supplies may increase in a non-linear manner by a factor that is greater than two due to the DC loop resistance of the coaxial cable supplying the power to those active devices. One slightly mitigating factor is the efficiency of line power supplies increases as the load increases. However, this may not be enough to offset the additional losses in the coaxial cable DC loop resistance.

It is important to note that the DC loop resistance will be unique for different node clusters. For this reason, it is not really practical to write completely effective generic rules for how to deal with the dynamic between nodes and power supplies in a deployment without carving out clusters as the quanta for system solutions. Beyond the identification of clusters however, generally stronger tools are needed to manage the complex nature element relationships. The next sections of the paper guide us through this process.

Cloud-Based Adaptive Power Manager

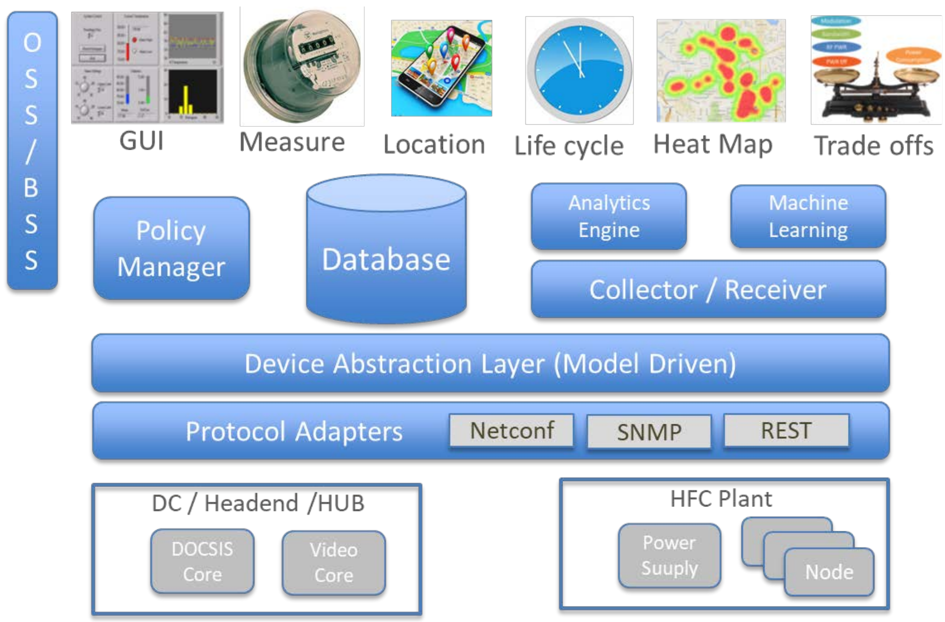

Figure 8 - Cloud-Based Adaptive Power Manager

The previous sections show that the power and performance dynamic of a cluster of nodes is complex enough that simple tools would not be able to extract its most optimal states. It is also the case that what is needed is some form of centralized view of the cluster, with knowledge of all parts of the system and ability to infer how their own sensitivities affect overall group targets, and more importantly, to be able to adjust power consumption of individual devices within the cluster when possible so that the current drawn by the entire cluster can be reduced and thus also the power consumed. Luckily this challenge comes to us at a time where the tools to create such a solution exist, and perhaps even with infrastructure already deployed, in support of other parts of the network. We then build a power manger in the context of the cloud, with otherwise generally understood architecture and approach, see Figure 8. Below we expand on this part of a cloud solution.

1. General Assumptions

Individual network elements must have power monitoring and reporting capabilities. Data may be communicated via a “push” model, whereby the network element sends telemetry data to a receiver in the management system, or a “pull” model, whereby the network element is polled periodically by a collector agent. An intelligent adaptive power management might consist of the following (or similar) components, deployed as micro-services in a cloud-native environment, as shown in Figure 8. Note, there have been discussions in this direction within the energy management community, but in this paper we aim to add more detail around the particular structure and function of supporting applications, (Sandoval, 2016)

2. Protocol adapters

It is expected that different network elements, from different vendors, would use a set of standard protocols to command and communicate power monitoring data and provide a method of control. Therefore, protocol adapters for protocols such as SNMP, REST, and NetConf would be required to communicate with network elements.

3. Device Abstraction Layer

A model-driven device abstraction layer would map device-specific application programmer interfaces (APIs) to a standard set of APIs used within the management system. This is where we assume the Yang model and/or management information bytes (MIBs) for various devices would be translated into generic descriptions. This effort is of course facilitated by the work of the SCTE sponsored APSIS specifications, where robust models have already been defined, or can be further refined. (SCTE, 2018)

4. Collector / Receiver Service

This service or group of services would receive telemetry from network elements sophisticated enough to push power monitoring data up to the management system. For elements that only support a pull model of data retrieval, the collector would periodically poll the desired data from the device. In the illustration, the collector / receiver service collects data from a power supply and distinct types of nodes in an HFC cluster, as well as data from various flavors of associated CMTS and quadrature amplitude modulation (QAM) devices in the cable headend or data center. The collector writes the raw data to the database as-is, with little or no processing of the data itself.

5. Analytics Engines

The analytics engine refers to environment that is particularly structured to handle the scale and workload generated by analyzing big data. If broadly deployed the data that is coming from node clusters could become very large very quickly and the analytics engine would provide a set of tools to organize the nodes cluster information read by the plant and allow us to do relational calculations. We expand on this is in the next section.

6. Data Base

The data base is formally a part of the analytics engine as it is the analytics engine that would determine the format for storage distribution and organization of data. Because this is a "big data" exercise there are several qualities to the storage of data that are necessary, like the ability to process in parallel, allow for low commodity hardware and interact with a robust resource manager. We expand on this in the next section.

7. Machine Learning

The machine learning module is really part of the analytics function but we call it out separately. It takes in sensor data, the capabilities of the equipment and statistically tracks both state data from cluster components and makes predictions and recommendations from applications according to programmed rules. In time this training data will allow for the creation of other algorithms such as those in Section 9 below, whose relationships would otherwise be too obscure to conjure up directly, to help maximize the relationship of power consumption and power performance of the cluster. This is exactly where machine learning can be most useful, and additionally, part of the vision for this module is to overcome the inevitable problem of fully calibrating parts in design and manufacturing before adding to the cluster, both for varied descriptions of power and performance and over time. Machine learning algorithms are often ideal for making predictions and recommendations from incomplete or slightly inaccurate data.

8. Policy Manager

The reduction of power, or the maximizing of performance via the execution of applications just described needs a higher perspective to serve the interests of the wholistic system. These decisions might have to be done in a case by case basis, for example in conjunction with the service level agreement (SLA) of a particular set, or individual end-line customers within the cluster. The policy manager is in charge of executing this higher perspective and must have a view of the other components in the system and their priorities, and this would include other packet cores like the CMTS.

9. Applications

There are a number of possible top line applications. These applications would access relevant information stored in the database to perform their respective functions. Perhaps the most basic function would be to provide a graphical user interface for the rest of the system to operators in the back-office. Another application could take the actual responsibility of measuring the system and providing energy consumption data at a granular and summary level to corporate sustainability offices. Another application could provide the physical location of clusters and another application provide a heat map for power efficiency improvement opportunities. Another application could provide a description of status within the life cycle of products in a cluster, making recommendations on how to optimize upcoming plant design over time, or even do so in modeling before a section of the plant is built. Another application could provide a real time savings calculation in terms of watts or dollars for the cluster—not shown in the picture. We do expect there would be an application that could adaptively modify the element states toward energy savings. There would be an application that makes suggestions on settings of nodes according to their respective power supply capability, towards overall power reduction or performance improvement. Finally, there could be apps that go beyond the settings of the node and extend to plant and network configurations via comparison between clusters. In general, the applications space is an open canvas for creativity and multi-vendor differentiation.

10. OSS / BSS

The operation support system (OSS) and billing suport systems (BSS) are the executive coordinators of the whole operation. They determine policy and the ultimate experience of the end line customer. Ultimately, the work of the power management environment must work within the confines set by the OSS and BSS. This is where the policy manager plays a crucial role making sure that any changes in the system do not violate higher directives.

Predictive Analytics

Predictive analytics (PA) has two main targets, the first is to take large sets of data and process it into meaningful relational information and the second is to learn from the experience of that data and anticipate meaningful directions that allow for positive action on the system. In our case the big data set is the power state information for clusters and their components in an HFC plant. The data could be as small as one node cluster, or as large as all the clusters in the national footprint of an MSO. In our case the relational information would be between the node components in a cluster and its line power supply, or between clusters, or comparisons of particular type of nodes in all clusters. There is no limit to what sort of relations we can study once the data is present. For the anticipation part of PA we could look to avoid failures due to future power inefficiencies and take action to prevent them, or avoid any performance shortcoming when there is power available to make it better. There is also no limit to the actions taken from anticipation learned. Some of the actions could include recommendations for the challenges pointed out in the "Use Cases for Adaptive Power Using APSIS", (SCTE245, 2018)

While a full treatment of PA would be very extensive, for our particular goal of highlighting the general principles within the context of HFC power management, we focus on three main functions: data storage structure, analysis of data for execution, and prediction of future situations via machine learning. We do this by presenting a practical example. While there are various frameworks to conduct this exercise, we rely here on the infrastructure of an open-source predictive analytics engine called Apache Hadoop. Hadoop stems from a project by Google to organize storage and do computation on big data. It is a now a common tool. (Note that in the case of HFC node clusters, our data analysis does not have to be real time so Hadoop is a good fit.) The framework for Hadoop is shown in Figure 9.

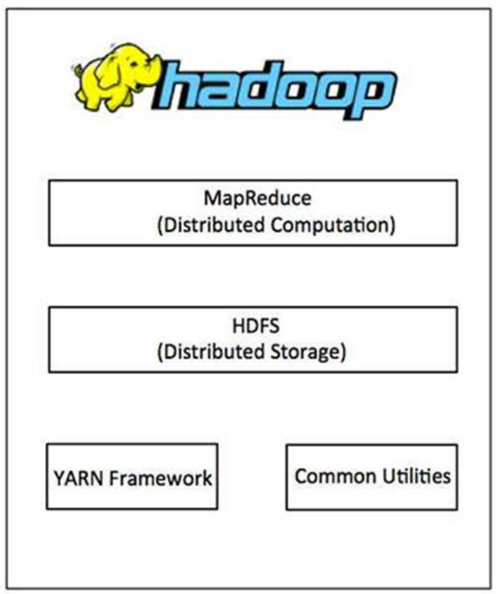

Figure 9 - Hadoop analytics engine structure

The first basic component of this analytics engine is its storage framework, the Hadoop Distributed File Storage (HDFS), (EdurekaHDFS, 2017). As the name implies, when data is stored it is done so in a distributed way, with a management layer or descriptive metadata and separate data clusters that are generic enough that can run on commodity hardware. A more concrete example will be shown below.

The second component is a sequence of algorithms that do analysis on data sets by distributing computation and coalescing results, these are the MapReduce functions (EdurekaMapReduce, 2017). In MapReduce one can execute known or program custom function on data sets. A more concrete example will be shown below.

The third component is Yet Another Resource Negotiator (YARN) which is an operating system of sorts for the architecture of data storage and its usage, (EdurekaYARN, 2017). This resource manager organizes the data clusters for storage and makes available data sets for computation and parallel processing, along with its scheduling.

There are also other common utilities that are present to facilitate the work of HDFS and MapReduce and while very helpful in practice we do not cover them here.

1. Heat Map Application

As an example, we sketch the heat map application on the PA platform. The heat map application would show the state of interaction between a line power supply and its subtending nodes in and HFC node cluster. This effectively shows whether a group of nodes is within the power envelope of its main supply or not.

1.1. Data Storage

Our first task is to take and organize and store our large data set from the field. In Figure 10 we show this data organization as would be done in HDFS. We note that there are independent raw data clusters and a higher level of metadata, which has description for the raw data clusters. These raw data clusters have redundancy and can be executed on commodity hardware. The management and organization of these raw data cluster sets is done by YARN. In our particular example the raw data clusters have the power state information for a power supply and its subtending nodes. Interestingly, we can organize the data clusters to coincide with HFC node clusters somewhat simplifying the steps that follow in MapReduce. The metadata in this example then just keeps a list of the types of components in the clusters, in our case a list of the power supplies and nodes is available.

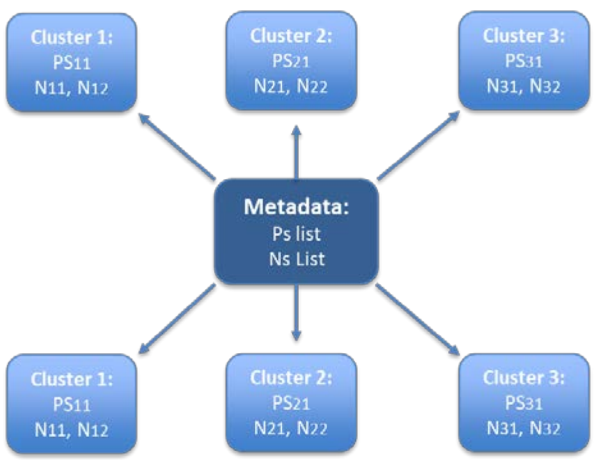

Figure 10 - Distributed data storage of clusters

1.2. Data Analysis

We show our data analysis on a framework that represents MapReduce. Our goal is to compute the comparison of the power available by a line power supply and the composite power usage of its subtended nodes.

In Figure 11 we show the progression for analysis. First is the identification of a data set as input. Note this data set is identified but not retrieved. In our case we need a list that includes power supply and nodes for HFC clusters 1-6. Then this data is organized (split) into usable groupings for distributed computation.

In our case the grouping matches the HFC node cluster for simplicity, so each node cluster's information is a set, line power supply for its two nodes, for six HFC node clusters. Next is the mapping itself, which is the execution of functions in a parallel manner on the data sets identified. In this case it is a logic identifier between the power supply capability and the addition of power consumption for the two nodes it serves. If the line power supply capability is more than its nodes usage then the value "0" is returned. If the line power supply capability is equal to the nodes usage then the value "1" is returned, and if the line power supply capability is less than the nodes usage the value "2" is returned. This is done in parallel for all six groups. The shuffling exercise is the organization of results from computations done in parallel. In our case it coalesces the power supplies that are well within their operation capacity, the power supplies which operate on the very edge and those that are being overworked or beyond their means. The reducing function then gives the actionable data we seek. In our case which power supplies we should monitor and which power supplies need immediate attention. This data is then exposed by the Heat Map application in a graphical representation of colors, yellow clusters if they need monitoring and red if they need attention.

Figure 11 - MapReduce Functions for Heat Map Application

This albeit simple example shows the thought progression behind using analytics for actionable intelligence. With the data at hand it is just a matter of creativity for how to use it.

2. Predictions

The predictive part of predictive analytics comes for the application of machine learning. Formally machine learning is a subset of the broader field of artificial intelligence which allows for some newer technologies like self-driving cars. In our case we prefer it because it allows us to apply statistical methods on data to learn from experience and predict the future behavior of a system, the system in our case being a node cluster.

Below we provide a simple example of using machine learning progression to predict what will happen to a cluster's line power supply output when two new nodes are added, as could be the case for an HFC node cluster when the HFC plant evolves to DAA, for instance.

In Figure 12 three plots show the progression of a very simple machine learning sequence. The very top diagram is borrowed from Figure 6 which shows physical layout of a node cluster with a line power supply, unique DC loop resistance, and six subtending nodes with particular power usage load.

The left graph shows a scatter plot of data learned by the system. In our case this is the power relationship between main supply and composite load. This data of course is available because of the analytics framework we have built already. The relationship mentioned above for instance would be built from a MapReduce module made to return corresponding line supply power and additive load.

The middle graph shows a regression for the data learned; this step is the application of some statistical tool for the data available. In our case we use a straight forward linear regression and extend its outcome beyond the data available. This is our prediction model then because as we look to add two nodes we can estimate the behavior of the system as two more nodes are added, as represented by the red circles on the line. We can also estimate the range of possibilities for these points, per the confidence level of the regression, thus the error bars on the points. Note that there are various statistical tools, beyond linear regression, available depending on the machine learning package being used, but generally with some mathematical treatment many relationships can be reduced to approximate linear forms. Thus simple linear regression can be quite common.

Finally, our last graph plots the actual relationship of two real nodes added to the system. Now it is time to test the model proposed by our machine learning algorithm. We see that in comparison to the regressed line the two real new points are not that far off, but we also note that the regressed line is not perfect, and so we can go through the process again and again until the user is satisfied with the fidelity of the line to the system. Once there is confidence in the model it can be used as actionable intelligence. In our case we can use the model to predict the behavior for other clusters of similar topology that need to add more nodes. We can use it to model a system before deployment. We can even use it to find and build optimal topologies for new builds. The possibilities are many. Note that the fine-tuning capability of these predictive techniques is what makes them ideal tools for node clusters because as we saw earlier the energy consumption signatures will be unique for different node clusters.

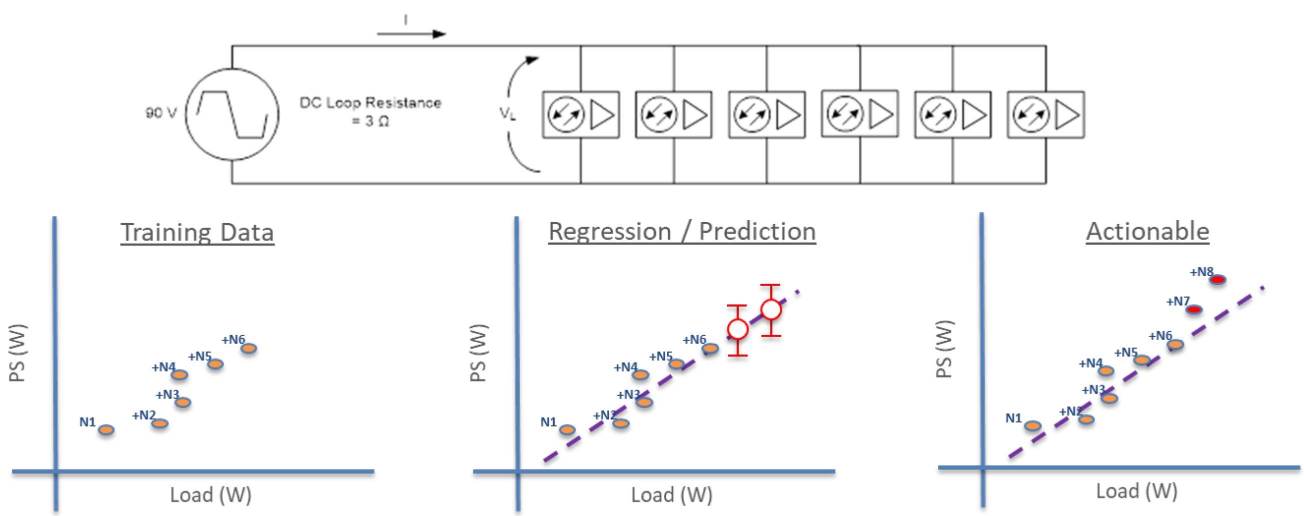

Figure 12 - Machine learning progression for adding nodes to cluster example.

Power-Managed HFC Lifecycle

The interesting takeaway about having a tool that can give insight into the power relationship of the HFC plant is that we can use it throughout the plant's lifecycle, see Figure 13. It can be used for modeling capacities and topologies even before it is built because we can anticipate outcomes with high confidence. It can be used for set up and create birth certificates for plants in deployment. It can be used to monitor and react in the day-to-day operations. It can help to manage service windows for the plant, it can help predict and identify end of life to power sensitive products. It can change the perspective of nodes and components according to expected power consumption, and can use items that are best in a system, not necessarily just efficient at one setting on its own. And it can also anticipate and assist in the sourcing of replacements parts. Overall it can have a very positive impact on the energy consumption and performance of HFC plant moving forward.

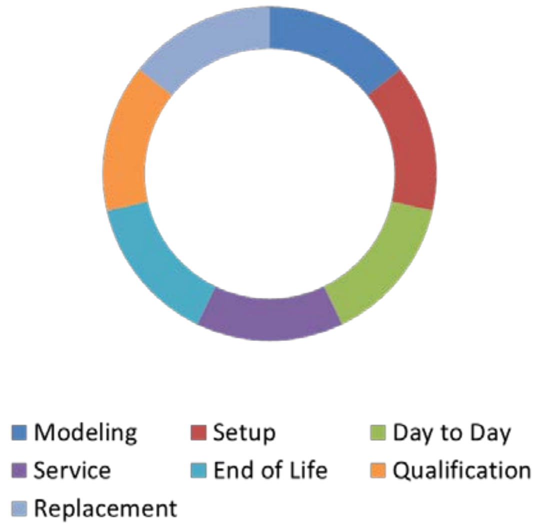

Figure 13 - Function of an adaptive power manager throughout the HFC plant.

Recommendations

The deployment and use of an adaptive power manager can be facilitated and cause several changes in the way the industry would approach the opportunity for HFC plant energy conservation measures and their relationship to network performance. Below we list several recommendations for the MSO community to take in this direction.

- Think of power in manageable chunks: we called them node clusters.
- Create robust and open power data models for <u>all</u> legacy and new nodes technologies.
- Centralize power management via cloud infrastructure.
- Allow predictive analytics to do the heavy lifting.
- Specify node power limits in context of cluster capabilities, not individual targets based on other averages.
- Qualify new node product in context of performance in clusters over time, not solely on individual performance.

Some form of power management is likely inevitable due to the otherwise large energy consumption and performance problems its absence would entail. The above recommendations would facilitate such an implementation and create an interoperable environment ripe for its fast development.

Abbreviations

ASIC	Application-Specific Integrated Circuit
BSS	Billing Systems Support
CIN	Carrier to Composite Intermodulation Noise
CMTS	Cable Modem Termination System

CNR	Carried to Noise Ratio
CPE	Customer Premise Equipment
DAA	Distributed Access Architectures
DC	Direct Current
DOCSIS	Data Over Cable Service Interface Specification
DPD	Digital Pre-Distortion
FDX	Full Duplex DOCSIS
FPGA	Field Programmable Gate Array
HDFS	Hadoop Data File System
HFC	Hybrid Fiber Coaxial
I2C	Inter-Integrated Circuit
IC	Integrated Circuit
ISBE	International Society of Broadband Experts
MSO	Multiple System Operator
OLT	Optical Line Termination
OSS	Operations Systems Support
OTN	Optical Transport Network
PA	Predictive Analytics
PHY	Physical Layer
QAM	Quadrature Amplitude Modulation
REST	Representational State Transfer
RF	Radio Frequency
rms	Root Mean Square
SCTE	Society Of Cable Television Engineers
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol
SNR	Signal to Noise Ratio
YARN	Yet Another Resource Negotiator

References

Chong, K. (2018, July 6). *Director of Product for Quorvo.* Retrieved from Broadband Tech Report: https://www.broadbandtechreport.com/articles/2018/07/growing-the-hfc-pipe-to-10-gbps.html

EdurekaHDFS. (2017, July 17). *HDFS Tutorial: Introduction to HDFS & its Features*. Retrieved from Edureka.com: https://www.edureka.co/blog/hdfs-tutorial

EdurekaMapReduce. (2017, July 17). *MapReduce Tutorial – Fundamentals of MapReduce with MapReduce Example*. Retrieved from Edureka.com: https://www.edureka.co/blog/mapreduce-tutorial/

EdurekaYARN. (2017, July 17). *Hadoop YARN Tutorial – Learn the Fundamentals of YARN Architecture*. Retrieved from Edureka.com: https://www.edureka.co/blog/hadoop-yarn-tutorial/

K. Sundaresan, e. a. (2016). Applications of Machine Learning in Cable Access Networks. *Spring Technical Forum*.

Mitchson G., H. D. (2016). Measuring and baselining power consumption in outside plant equipment and power supplies. *SCTE Cable-Tec Expo.* www.SCTE.org.

Sandoval, F. (2016). *APSIS and Open Daylight*. Retrieved from https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=8&ved=0ahUKEwjI89qFrancAhWBc98KHT11BqEQFghSMAc&url=https%3A%2F%2Fwiki.opendaylight.org%2Fimages%2Fc%2Fc4%2FAPSIS_and_OpenDaylight_-_light.pptx&usg=AOvVaw1sHqHxs0z4inNNYEfRL8ZK

SCTE. (2018). *SCTE Energy Management Program*. Retrieved from SCTE: https://www.scte.org/SCTE/Areas_of_Interest/Energy_Management/SCTE/Areas_of_Interest/SCTE_Energy_Management_Program.aspx?hkey=fa20e5a1-38c1-444e-84ee-e5f35541d6bd

SCTE245. (2018). *SCTE 245 2018, Use Cases for Adaptive Power Using APSIS.* Retrieved from www. SCTE.org: https://www.scte.org/SCTEDocs/Standards/SCTE%20245%202018.pdf