Access Network Operations Savings Through Extending Automation and Orchestration Beyond Remote PHY

A Technical Paper prepared for SCTE/ISBE by

John Holobinko
Director Access Networks Strategy
Cisco Systems, Inc.
Lawrenceville, GA
Phone: 770-236-1123
jholobin@cisco.com

Ron Zimmerman
Director, Technology
Cox Communications, Inc.
ron.zimmerman@cox.com

Todd Greene
Product Manager, Marketing
Cisco Systems, Inc.
greenet@cisco.com

Table of Contents

List of Figures

Introduction

The basic design of optical nodes has gone unchanged for over a decade. This paper explores the notion that rethinking the design and function of the optical node, specifically by incorporating intelligence and leveraging automation, can potentially have a material effect on cable operations expense (OpEx) and network availability.

Content

1. Problem Statement

Cable access networks are built with a combination of optical nodes and amplifiers, coaxial cables, taps, connectors and drop cables. By their very nature, cable access networks are more OpEx intensive than passive optical fiber networks (PONs). For example, the passive elements of a coaxial distribution network include many separate pieces of cable, taps and connectors that join these together. Connectors may loosen; water seals may fail; cables may be dented resulting in impedance changes. All of these deleteriously impact network performance and require corrective maintenance. While not immune, fiber plant is less susceptible to such issues.

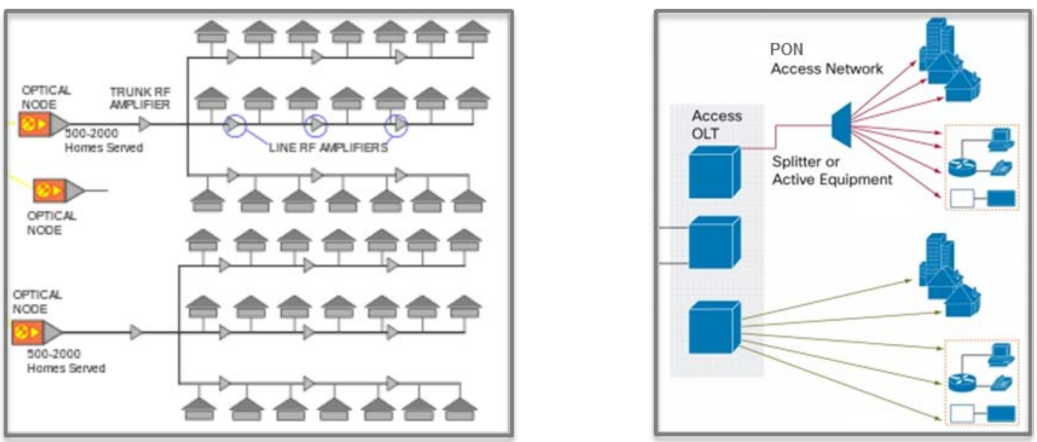

Figure 1 – Cable Access Network versus PON Network

Even small cascade/no cascade cable access networks remain more costly to maintain than PON networks.

In a cable access network there are many "cable specific" potential failure points including:

- Node (N+0) or Node + Amplifiers (N+M)
- Taps – water ingress, etc.
- Hardline cable – dent, rodent damage, shield issues
- Drop cable – cut, damaged, improper connections
- RF connectors – loose, water, damaged

In a PON, the number of potential failure/performance impacting mechanisms is significantly less than current HFC networks. The major failure mechanisms in PON are:

- Optical network terminals and optical line terminals (ONTs + OLTs)
- Fiber cable cut
- Optical connectors

However, PONs are much more expensive for cable operators to implement for serving existing subscriber areas than cable network upgrades to N+1 or N+0 architectures. Put into perspective, in PON there is a dramatic trade off of much higher CapEx in exchange for lower OpEx.

As discussed below, given new advances in bandwidth achievable by implementing DOCSIS 3.1 and 1.2GHz bandwidth, cable systems can potentially deliver 10Gbps downstream. With emerging full duplex DOCSIS (FDX) technology, the upstream is potentially capable of 5 Gbps. Therefore, the cable system challenge vis a vis PON is not as much bandwidth as it is to shrink the operational cost of the cable access network to approach the operational cost of the PON network.

Figure 2 – The Chasm: Cable Acess OpEx vs. PON OpEx

2. Postulating a Change in Node Thinking

We began our study with a question: Can automation and additional intelligence be incorporated within the cable access network in such a way as to reduce the gap between the cable access network and PON relative to outside plant installation and maintenance?

One of the current means of reducing maintenance is a radical reduction in the number of active devices between the cable headend and the subscriber. Reducing amplifier cascades to N+2 will typically provide sufficient bandwidth per subscriber, while N+0 networks have the additional benefit of having only a single active device between the headend and the subscriber, albeit for a much higher capital investment. As evidenced by announcements for N+0 network plans over the last year, more cable operators appear willing to trade off this higher capital expense (CapEx) to obtain gains achieved in OpEx reduction within a low/no cascade cable access network.

In addition to N+0 benefits, significant technology advances are coming to the cable access network. In the near future, a cable network implementing Remote Phy and Full Duplex DOCSIS (FDX) in N+0 architecture across 1.2GHz of spectrum will provide more bandwidth to a given group of subscribers (a service group) than a classic gigabit PON network (GPON) serving the same subscribers. However, even as these advances will deliver a dramatic technical advance in bi-directional bandwidth capacity, the operations costs of an HFC network still remain significantly higher than most of today's installed legacy PON networks, which are 1G EPON or 2.5G x 1G GPON.

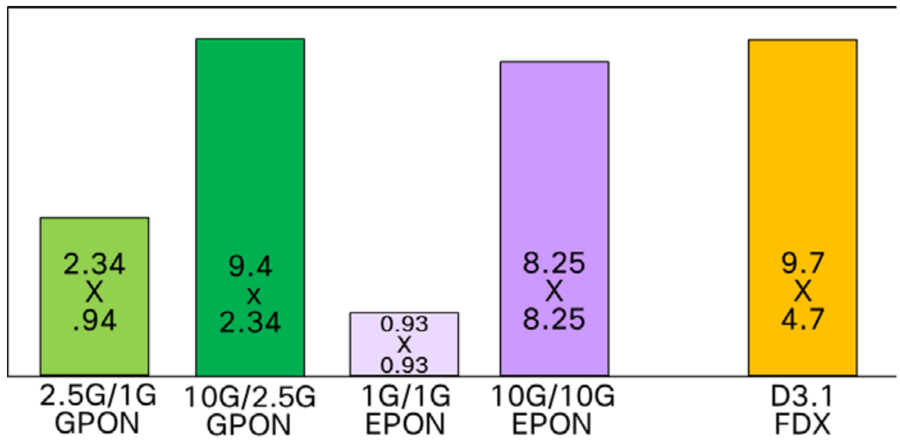

Figure 3 - Various PON Effective Throughputs versus DOCSIS 3.1 with FDX in 1.2GHz System

The performance and reliability of a cable access network remains highly dependent on the skills of the RF technicians who are tasked with installing and maintaining the nodes (and amplifiers where applicable), passives, taps and coaxial cables. Even with proactive network management (PNM) tools, the overall installation, plant operation and problem determination of the cable network remains dependent on the skill of the cable technician. Remote PHY device (RPD) technology can mitigate some of the issues previously associated with the analog fiber portion of the cable plant, but it also brings operational complications. Ironically, an N+0 network employing 10GE optics connected to RPDs is in some ways more of a challenge to maintain than a traditional hybrid fiber coax (HFC) network employing analog optics, because the use of some network diagnostic tools are either not supported or simply become too expensive. (The reason for the latter is that in a rebuild to N+0, the number of nodes in the network increases by an order of magnitude.)

In a low/no cascade cable access network, the node becomes the central control point for the service group. The design and function of HFC optical nodes has remained relatively the same since their introduction into the network outside plant over twenty years ago. A cursory look at node performance and failures would conclude that nodes are very reliable components of the network, yielding little potential for improving cable access network costs and availability. However, what that does not consider is the node as a remote information and remote control point within the cable RF distribution plant. Therefore we posed the following question:

3. Can a node be redefined to positively impact cable network OpEx

Can the node be designed in such a way as to have a material effect on plant power consumption? What if the concept of the node could be redefined in such a way as to enable more control and problem determination of key plant issues and failure mechanisms? Based on these questions, we set out to determine whether a rethinking of the node and its functions could have a material effect on operations costs and network availability.

For many years, the cost of automating nodes and providing a smattering of additional information remotely from nodes to hubs was so expensive as to make wide spread deployment impractical in most systems. In North America and Europe, only a few operators have even implemented the ability to report back basic information such as optics status, and the ability to attenuate a return path leg (also known as a "wink" switch) as a means of isolating some cable system problems. But, what if by rethinking the node, automation technology and intelligence could be directly incorporated into the node design, offsetting certain costs and making it possible to build an advanced node at approximate cost parity with traditional "dumb" nodes?

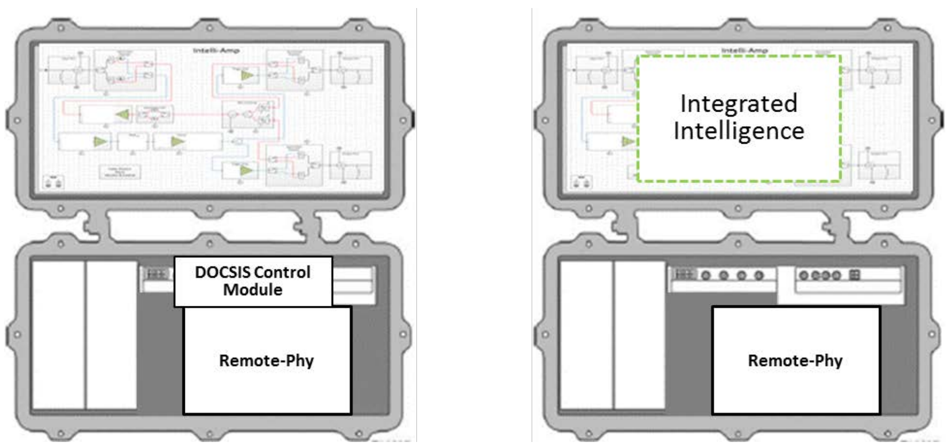

Figure 4 – Traditional node with DOCSIS transponder vs. node built with integrated intelligence

4. Potential Cost Savings of An Intelligent "No-touch" Node

Importantly, to determine if enhanced node functionality can have an impact on network reliability and services availability requires examining the major or areas of OpEx cost for the network.

1. Power Costs
2. Operations Costs and System Availability
3. Lost Subscriber Costs (i.e. customer retention)

4.1. Power savings

For an N+0 architecture to be the most capital cost efficient requires redesign of the system to maximize the number of subscribers able to be reached by a node, and for the node to achieve the highest output power possible on every leg. While the number of households passed (HHP) per node will vary based on home density, a North American average is 40-70 HHP per node.   However, homes are not laid out linearly. Densities and topologies vary. In 1.2 GHz N+0 cable access network designs Cisco has performed, we find that not every node will utilize all four outputs. The average number of outputs is approximately 3.4 versus 4.0. In other words, in a 50,000 HHP design with 50 HHP per node, there will be 1000 nodes, and a total of 4000 outputs, of which 600 will not be used. This means that with traditional nodes, the operator will be paying for substantial power consumption across the N+0 footprint that is unnecessary over the life of the system. When systems are redesigned for N+0, they are mostly designed for 1.2 GHz operation with the goal of 10 year lifetimes before they reach traffic saturation. Operating at 1.2 GHz is substantially more expensive than operating at 860 MHz or 1.0 GHz, yet the operator may not utilize the full bandwidth for a number of years. The difference in power consumption between 860 MHz, 1.0 GHz and 1.2 GHz is also substantial.

With today's traditional high output nodes there is no means to turn off unused RF outputs, nor is there the ability to run an RF output hybrid at less than full power (i.e. full bias current and bias voltage). Therefore these savings are currently unobtainable with traditional nodes. However, what if the node were intelligent and control over all of its internal operating parameters could be done remotely without touching the node? Specifically, if the node were designed to control both "on/off" plus variable bias voltage/ current of each output of each output hybrid amplifier, what are the potential savings?

We calculated the power savings of the unused outputs across the footprint using the ratio of 3.4 active ports per 4 port optical node. For power savings from lower initial spectrum use, we assumed five years (the estimated system half-life as defined by the cable operator) as the point when a system would require full use of 1.2 GHz bandwidth. We then calculated the power in those first years that is wasted based on two scenarios: with no traffic occurring above 860MHz or 1.0 GHz. For systems that pay for metered power versus bulk rate, the potential cost savings of lower power consumption is substantial. Using a cost of $0.10 per KWH and a total footprint of 1 Million HHP results in the following total potential combined power savings from turning off unused ports and operating the remaining ports at reduced power until they are fully enabled at the five year point:

Table 1 - Computed Power Savings From Automated Control of Output Power Hybrids

10 Year Total Savings @ 860 MHz	$2,993,000
10 Year Total Savings @ 1.0 GHz	$2,620,700
10 Year Average Savings/Node 860 MHz	$179.58
10 Year Average Savings/Node 1.0 GHz	$157.24

4.2. Operations Costs Savings and System Availability

We address OpEx savings and systems availability simultaneously as they are interrelated. Operations costs savings are derived by determining a reduction in the time and effort required to detect and remediate system issues. Faster detection, remediation, or rectification of issues reduces the time that the system is not operating or services are operating in a degraded state, resulting in higher network availability. Higher network availability is a factor that corresponds to reduced customer churn.

System/services availability is dependent on at least three major factors:

1. The number of hard cable access network failures
2. The number of degraded performance issues in the access network
3. The time to detect and correct hard failures and degraded performance issues times the number of subscribers per event

In today's cable networks, more customer complaints originate from degraded performance than hard network failures. As stated previously, the number of hard failures in an N+0 network is substantially less than in a cascaded amplifier network. Hard failures are limited to node failures and coax/fiber cuts. By far, degraded performance incidents dominate service calls related to the access network.

Cable plant issues that degrade return path performance represent some of the most difficult problems to resolve. They can result in slow internet performance, or intermittent loss of service as cable modems lose registration, re-register, and lose registration repeatedly. Even with next generation Remote Phys installed in nodes, the operator is challenged to determine which leg of the node has the problem. PNM can be used to somewhat isolate the problem to a localized area of the cable plant, but PNM can only narrow down the likely causes of the problem, versus pinpoint the exact cause in every case.
Return Path Affecting Problem Detection, Mitigation and Resolution

Today, usually as a result of a customer triggered trouble ticket, technicians are sent out to the service area. The first course of action is to visit the node in order to get a reference signal. Then the technicians go from the node to the location where the suspected problem exists. A few examples include loose connector, water in tap, rodent damage to cable, dented cable, etc. To confirm that they have located the problem either requires an additional person at the headend or a special diagnostics program. The technician corrects what s/he believes the problem to be, then waits for feedback to see if the problem is corrected. In some cases, they also use a diagnostic program that interfaces to the cable modem termination system (CMTS). This multi-step corrective action takes time and resources.

How could functionality in the node be changed to significantly reduce the effort to address return path problems?

1. *Sense the problem before the customer trouble ticket even occurs.* The node should have the ability to take measurements of return path waveforms for each individual return path at the node RF port, independent of the RPD, and send these to a software package such as an enhance PNM system with big data analytics in order to perform trend line analysis and detect issues before the customer calls. Such software goes beyond a traditional PNM system.
2. *Enable the problem to be mitigated before the technician is dispatched.* Waveforms for each node leg should be remotely visible and recordable and the node should provide a capability to remotely attenuate each individual return path leg, so that in the case of impulse noise, customers on that leg may still operate with reduced performance instead of total loss of service. In the case where the problem is too severe, the node should be able to turn off only the offending leg, so that the other 75% of the homes attached to the node remain at full service.
3. *Eliminate the technician's need to visit the node before going to the suspected problem location.* The node should provide a remote spectrum analysis function and eliminate the need for the technician to calibrate test equipment. Note that spectrum analysis differs from spectrum capture.
4. *Enable the technician to validate the problem and the solution from the point in the plant where the problem originates.* This requires a real time spectrum analysis function at the node on the specific leg affected. It is not sufficient to provide a static "snapshot" of the spectrum. The technician will want to view the waveform live, therefore the node function should be able to

provide multiple images per second with sufficient resolution to confirm both the problem and resolution.

Notably, the measurement capability of the node must enable the technician to be able to obtain all of the information necessary obtained previously at the node test points, without the need to go to the node. Therefore, the intelligent no-touch node needs to provide the primary functions of a technician's hand held spectrum analyzer and to do so remotely. Ideally, a secure application will enable the technician to do remote inquiries of the node. Typical spectrum analyzer functions that should to be provided include the following measurements:

- Full band spectrum capture and display
- Partial band spectrum continuous sweep (at least once every 100 msec)
- Channel power / Total composite power (TCP) Measurement
- Max. hold, Min. hold, average, clear write
- CCDF (complementary cumulative distribution function used for peak to average power)
- MER / BER (modulation error ratio/ bit error rate)

While the number of truck rolls per year per 1000 subscribers is based on a number of factors, and the cost per truck roll varies by cable operator, what we can say is that the time in the field per call will be reduced. We estimate that the time to diagnose and remediate can be improved by at least by at least 25% per truck roll for these types of problems over traditional systems. Given that a truck roll costs between $100 and $150 for most operators, the potential savings over the life of the node are potentially substantial.

In addition to truck roll costs, there is the cost of service calls to customers when any service degradation or outage happens. While the cost per call is significantly less, t number of potential calls may result from a single event.

4.2.1. Subscriber Losses Due to System Issues (Customer Retention)

Network outages and degraded performance are key contributors to subscriber losses. Correspondingly, we know that the ability to detect problems sooner and to mitigate any disruption of services to the smallest number of customers and the shortest overall time has a correlation to customer retention. While we do not yet have sufficient operating data to measure the impact of mitigating these issues faster, it is clear that these abilities will positively affect retention. Given that the enterprise value of a cable customer today is between US$4,000 and $5,000 it is quite possible that benefits of customer retention may actually be higher than OpEx savings provided by intelligent, touchless nodes.

What about today's nodes outfitted with a remote phy module? Don't they provide the same functionality and the same benefits? They do not. The table at the end of this section summarizes the functions we have identified, versus what is provided by traditional nodes, or a traditional node plus RPhy.

Performance and Reliability Issues Related to RF Technician Actions

Another cause of service problems is ironically the result of RF technician behaviors. A technician may be called to a subscriber residence to address a service issue. When the problem is not easily resolved, in hope of a fix some technicians will go back to the node and raise the downstream output level by removing and replacing the attenuator pads in the node housing. Not only does this change system behavior, changing the pad value creates a system outage for all customers on the RF leg while the pad change is being made. The new pad value changes the RF output level, overriding the original system design and impacting the RF levels on all of the homes attached to that leg of the node. Over time, these changes usually result in additional customer service calls. In the operator employs PNM to detect level imbalances, a trouble ticket may be issued to send a technician to the node and revert the pad levels to those specified in the network design, and in doing so create another service outage. Because of these outages, the reversions need to be done in the overnight maintenance window, resulting in additional costs. The cost of these truck rolls and any associated additional customer calls are unnecessary, additional OpEx expense. The total unnecessary OpEx cost involves the original time of the technician to modify the RF output level and the time and additional truck roll required to revert the system back to its design specified level. In some systems, this is a major cause of needing to perform periodic system sweeps.

An additional expense of traditional nodes is the cost of maintaining passives inventories for technicians. While the passives themselves are not very expensive, the time and effort to maintain a complete set of pad (and equalizer) values, as these are consumed in the field. This is not a trivial matter, as technicians only tend to ask for new values when they have entirely consumed the value of a pad and need that value for the job. It can result in wasted truck rolls when a technician discovers that s/he does not have the right pad or equalizer value, plus the time spent ordering missing tap values on a rush order.

How could node functionality be enhanced to eliminate this type of technician behaviors, eliminate the pad and equalizer parts inventory issue, and simultaneously reduce OpEx?

1. Eliminate all pads and equalizer plug-in accessories from the node design – the node electronics must be designed to self-align and not require manual adjustment by the technician. Any adjustments must only be able to be done electronically and remotely without the need for any plug-in pads or equalizers.
2. The operator must be given the choice to prevent any changes to node settings by technicians, or at their option, to be able to authorize only certain technicians to modify values on certain nodes while automatically logging and time stamping any change made to node settings. By providing absolute control over node operation, unauthorized changes to levels are eliminated.

By automating set up and control of the node, and eliminating the need for any plug-in accessories, the potential OpEx savings are elimination of accessories management, plus avoidance of truck rolls necessary to rebalance systems to their original design settings. But most importantly, for each unauthorized level change, two outages are eliminated: The first one by the rogue technician and the second one for restoring the RF levels to the correct design values.

Other forms of diagnostics

Some problems are highly difficult to detect because they are intermittent – either occurring in unpredictable time intervals or very infrequently. If a particular node leg is suspect, polling of the suspected leg at very short intervals for long periods of time is not practical today, due to the heavy load it puts on the DOCSIS platform and the impact on data throughput. If the node were able to do full band or partial band capture and its polling rate is adjustable, a large number of samples could be stored in local memory and downloaded to the hub at the operator’s convenience. This polling could be done completely independent of the CMTS and without impacting customer DOCSIS performance. But to do this requires intelligence in the node in the form of a processor, local memory, plus circuitry to measure the spectrum at critical points within the node.

Simplifying Initial Installation and Eliminating Set-Up Errors

As addressed previously, the ability to set levels initially lower in order to save power and then later remotely change levels to support full spectrum operation results in considerable power savings. Importantly, for a system that is going to be rebuilt to N+0, there will be between ten and fifteen times the number of current nodes in the network, and all must be installed at the same time. Typically, a contractor hangs the node, while a second visit is made by a technician to install the node and confirm its proper operation.

If the node functions can be automated using intelligence, then the node installation can be made automatic, eliminating the need for the skilled technician at installation. If the contractor can install the fiber connectors and RF power, the Remote PHY can advertise itself to the network and self-install, followed similarly by the node if it has the requisite intelligence. Geolocation data and serial numbers can be used to identify the specific node, and a look up table used to download the specific RF design values specified for the node location. No previous set up of the node is required and there is no need to deal with any pads or equalizers at installation. This will result in considerable savings at installation as well as insuring that the proper node configuration is made 100% of the time.

Ideally, all of these communications functions should be made controllable using a standards based interface and standards based communications. In nodes that employ DOCSIS Remote PHY and the Open PHY standard this means utilizing Ethernet at the transport layer and using a Netconf/Yang model to provide an open interface to third party software.

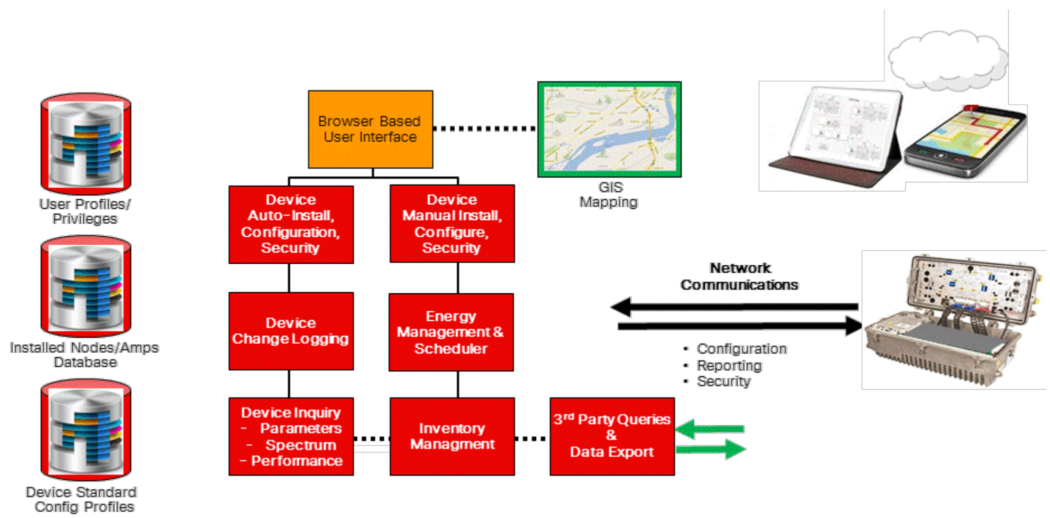

Figure 5 - Touchless Control of Intelligent Node

Follow-on Operation with Advanced Modulation Schemes

No operators are fully utilizing the entire 1.2GHz spectrum in their downstream in N+0 systems today. While it is feasible to calculate required RF levels for full spectrum utilization, in practice it may be possible to lower levels, or do other adjustments in particular areas of the spectrum. As FDX is implemented, there will be an additional learning curve on signal levels and settings. Finally, there may be improvements in customer premises equipment (CPE) that change the definition of what ideal levels are in both the forward and return path of the system.

With traditional nodes, the ability to redefine levels, slopes, and other node operating parameters is severely handicapped by the need to visit each individual node to effectuate any changes, and every change results in a network outage for those customers attached to the node under adjustment. In an N+0 network, the sheer number of nodes makes any such adjustment across the footprint non-feasible.

With an intelligent, no-touch node, he ability to remotely set and adjust node levels and other parameters individually, regionally, or globally provides a sort of futureproofing for future learnings. None of this is possible with traditional nodes.

Table 2 - Comparison of Advanced No-Touch Node vs. Traditional Node Table

	Advanced Node with Intelligence	Traditional Node	Traditional Node with 1 x 1 RPD	Traditional Node with 1 x 2 RPD
Waveforms remotely visible on a per leg basis	YES	NO*	NO, per combined returns only	NO, 2 legs combined at a time
Return path level and turn off control per leg	YES	NO	NO	NO
Remote adjustment of signal levels, slope, while eliminating pads or EQs	YES	NO	NO	NO
Capture of waveform and other data without DOCSIS performance impact	YES	NO	Limited information	Limited information
Remote power savings and power on/off for each RF Power Hybrid	YES	NO	NO	NO
Ability to control access to RF section, log changes, etc.	YES	NO	NO	NO
* Only possible when equipped with internal cable modem and controls at significant additional cost				

5. Next Generation Intelligent No-Touch Nodes in non N+0 networks

While the benefits of complete automation are reduced in a network with amplifier cascades, there remain benefits that are not available with traditional nodes. The ability to see return path levels on each leg enables each return path to be examined independent of the other return paths. This is beneficial because in today's nodes, two to four return paths are typically RF combined before sending to the RPD or burst receivers in the CMTS. Additionally, for integrated CMTS's, very often the return paths from two to four nodes will be combined into a single service group, meaning that the burst receiver sees the simultaneous return from up to 16 return path legs. Being able to see each individual return path signal at the entry point to the node and after amplification gives up to 16 times the resolution as a traditional return path service group solution.

The ability to attenuate one return path to minimize the number of subscribers impacted by a high noise event remains valuable. The ability to bring back additional telemetry for PNM allows better problem determination. For example, in an N+1 system where the amplifier is a single port line extender, spectrum display and measurement of return path inputs and forward path outputs at the node can provide additional information for diagnosing problems at the line extender or even the cable beyond.

Conclusion

We believe that the industry is on the verge of the first major change to node design in years. Next generation intelligent, no-touch nodes will integrate digital control, intelligent signal adjustment and more granular plant measurement capabilities such as rapid sweep spectrum analysis per leg (not to be confused with static spectrum display) within their basic functionality. This will require an ability to develop these nodes via new digital design techniques such that the cost of these nodes to the operator will be on parity with traditional non-intelligent nodes. Complementing these intelligent nodes and the RPDs installed inside them will be automation and orchestration software that will automate installation, control technician access, and thereby remove variables out of the technician equation. These new capabilities promise a significant impact on reducing field problems, outage times, thereby improving OpEx and customer satisfaction and reduce the PON operations advantage.

Abbreviations

BER	bit error rate
CapEx	capital expense
CMTS	cable modem termination system
CPE	customer premises equipment
DOCSIS	data over cable service interface specification
EPON	Ethernet passive optical network
FDX	full duplex DOCSIS
FEC	forward error correction
Gbps	gigabits per second
GHz	Gigahertz
GPON	gigabit passive optical network
HFC	hybrid fiber-coax
HHP	households passed
Hz	Hertz
ISBE	International Society of Broadband Experts
MER	modulation error ratio
MHz	Megahertz
OpEx	operations expense
OLT	optical line terminal (PON network element)
ONT	optical network terminal (PON network element)
PNM	proactive network management
PON	passive optical network
RF	radio frequency
RPD	remote PHY device
SCTE	Society of Cable Telecommunications Engineers
TCP	total composite power