Re-Thinking Service Interruptions from Architecture to Operations

A technical paper prepared for presentation at SCTE TechExpo24

Anna Gallegos
Product Manager
Comcast
Anna_Gallegos@comcast.com

Aleksander Hübert
Machine Learning Engineer
Comcast
Aleksander_Hubert@comcast.com

Andy Martushev,
Distinguished Engineer
Comcast
Andy_Martushev@comcast.com

Matthew Stehman
Principal Machine Learning Engineer
Comcast
Matthew_Stehman@comcast.com

Table of Contents

Title	Page Number
1. Introduction	3
2. Approach	3
2.1. Device Grouping	5
2.2. Root Cause Identification	6
3. Capabilities	8
3.1. Implemenation	8
3.2. Use Cases	9
3.2.1. Commercial Power Outages	9
3.2.2. HFC Outages	10
3.2.3. Fiber Outages	11
4. Operational and Business Benefits	12
4.1. Operational Benefits	12
4.2. Business Benefits	14
5. Conclusion	15
Abbreviations	16
Bibliography & References	16

List of Figures

Title	Page Number
Figure 1: HFC Networks Power Grid Footprints	4
Figure 2: Temporally Separated Outages	5
Figure 3: GOLD Data Flow	7
Figure 4: Alarm Processing Flow	8
Figure 5: CPO Duration and Number of Nodes	10
Figure 6: Network Element Root Cause Breakdown for HFC Outages	11
Figure 7: Typical Access Network Fiber Distribution	12
Figure 8: Legacy v GOLD Root Cause & Dispatch Processes	13

1. Introduction

As internet access grows increasingly vital to our daily lives, so do customers' expectations for uninterrupted service. Losses of connectivity are, unfortunately, inevitable due to uncontrollable circumstances like fallen trees or inclement weather. However, to ensure the best possible customer experience it is critical to restore service quickly and provide transparency to impacted customers. To do this, internet service providers (ISPs) must swiftly and accurately detect interruptions and their underlying root causes.

Historically, the nature and causes of interruptions have been difficult to pinpoint. As such, determining root cause has traditionally been a manual and highly specialized process. This poses several challenges to ISPs as manual processes are timely, error prone, resource intensive and less adaptable to technological change. Consequently, Comcast realized the need to rethink this historical approach to understanding losses of connectivity to meet customer expectations.

It was this need that resulted in the creation of Comcast's Global Outage Localization and Detection (GOLD) algorithm. The vision for GOLD was to provide a centralized algorithm that can dynamically detect losses of connectivity and determine their root cause from inception to close. The goal for dynamic detection was to simultaneously understand both the temporal nature and geographic spread of connectivity issues. For root cause identification, the aim was to have an algorithm that can decipher a wide variety of network behaviors to prescribe a simplified path to remediation. More simply, GOLD must be aware of how different network impairments manifest so it can delineate between issues such as a commercial power outage (CPO) or fiber cut as the paths to restore connectivity may differ.

This paper will show how GOLD’s capabilities enable better operations that improve the customer experience. To do this, it provides examples of commercial power and hybrid-fiber coaxial (HFC) where GOLD was used in the detection process. Additionally, it will provide insight into how this customer focused algorithm poses ancillary opportunities for more efficient business management.

2. Approach

GOLD defines an outage as a grouping of offline customers that lost connectivity due to a common root cause. While this definition sounds simple, it requires a more holistic approach to understanding outages when compared to traditional methods.

Historically, when connectivity was disrupted, an analysis was performed at the node level to determine the impacted customers and root cause. This is because device telemetry was typically aggregated to the node. However, limiting analysis to this level of the network often led to mislabeling the outage cause and, subsequently, unnecessary truck rolls. This is because outages are not confined by the logic of data collection and detection systems. Instead, they are dynamic and may be agnostic to node boundaries.

This is especially pertinent for service interruptions caused by CPOs. While HFC networks have a dependency on the commercial power grid and often serve the same customers, their underlying topologies are different: a service interruption caused by a CPO can span node boundaries with the possibility of only impacting a small subsection of a given node.

A sample interplay of the commercial power and HFC networks is depicted in Figure 1, where the HFC network may not overlay one to one on the commercial power grid. In this example, the commercial grid powers a subset of customer premises across four different nodes in the HFC topology. What's more, the outside plant power supply of only one of the nodes is present in the footprint of power grid two. This scenario of non-trivial relationships between the HFC plant and commercial power grid footprints is pervasive and requires advanced techniques to identify the corresponding root cause of a service interruption. But that difference is also the opportunity to exploit.

Figure 1: HFC Networks Power Grid Footprints

While Figure 1 exemplifies how a CPO can span node boundaries. The same can be said for interruptions caused by network specific factors like fiber cuts or firmware update bugs. As such, this ambiguity became a key pain point to solve for when both determining how GOLD would define an outage, and the approach Comcast would take to building the algorithm.

The GOLD algorithm utilizes a two phased approach that is agnostic of node level constructs to generate outages across the entire HFC network, and thus more accurately label the root cause, which is inclusive of both understanding what caused the outage and localization of the problem.

There are two main factors that are central to detecting and labeling an outage. The first is to group devices that are offline into an outage with potentially limited knowledge of their exact connection to the network and the power grid. The second step is to use the patterns of the offline devices in conjunction with telemetry and other data sources to determine the root cause of the outage.

The following sections will expand on these challenges and the approaches GOLD utilizes for outage detection and classification.

2.1. Device Grouping

There are two key concepts for grouping devices together: spatial and temporal closeness. Devices that are offline for the same reason should be both spatially and temporally related. Regardless of the underlying network topologies (cable network and the power grid), the spatial closeness of two offline devices is an important feature of an outage. Next, the temporal aspect is a significant element for outage determination. A group of devices that are near each other losing connectivity at the same time indicates that the interruption was likely caused by the same issue. This is true regardless of the underlying topology.

Devices are grouped spatially based on geolocation. This provides coverage without looking at node level or network topology connections. Meanwhile, temporal dynamics are also important for the creation of outages. This is because most outage causes will affect devices immediately and thus take impacted customers offline within a similar timeframe. Additionally, tracking the temporal component allows outages to be analyzed over time and feed into further analysis of the issue. These temporal insights allow Comcast to identify the following scenarios. First, when power is restored due to a CPO, it can lead to a power surge that causes HFC equipment failures. Secondly, a strong wind gust could cause a power outage in a neighborhood and an hour later another gust may cause an HFC outage in an adjacent subdivision. In both cases, these outages would be considered unique as the underlying cause is different. This is especially important for remediation. Identifying the root cause requires correctly grouping devices. If not, some subscribers will remain offline when the fix is completed.

Figure 2: Temporally Separated Outages

An example of temporally separated outages is shown in Figure 2. In this scenario a group of customers went offline at 10:00 a.m. followed by another neighboring group of customers going offline at 11:30 a.m. Though the outages are nearby, the temporal difference likely indicates differing root causes. A common example of this is a commercial power outage leading to power supplies running on battery and eventually losing power. Figure 2 highlights possible outage patterns in *outage 1* where customers are intermittently online within the boundary which could indicate a power issue and *outage 2* showcases a pattern resembling an HFC outage as all devices are offline within the outage boundary. Cases where a commercial power outage has affected a power supply, and the power supply runs out of battery, we see devices not affected by the loss of power lose connection to the cable network time delayed from the onset of the power issues.

Since the topologies are not always known, this method allows for the creation of outages that can later be applied to all possible outage scenarios and analysis can take place at the level beyond the device grouping. Due to the variation in geography, network challenges, and customer needs, these spatial and temporal groupings allow GOLD to learn from previous outages and tune the algorithms to specific geographies.

2.2. Root Cause Identification

Root cause identification is the next step in GOLD's outage detection process. The device groupings from the previous step are the basis of this process. However, it also requires carefully crafted logic that can synthesize disparate pieces of evidence to point to a root cause.

These pieces of evidence fall into three categories: offline device patterns, device telemetry, and augmented data sources. Each feed into the GOLD system, which generates the two outputs: outage boundary and CPO. This is shown below in Figure 3.

Each outage has its own spatial temporal pattern of devices. This includes whether devices within the node are partially offline, if there are pockets of online devices, and clues into whether the outage is HFC or commercial power. Device telemetry is a supplement when a determination is not immediately clear from the pattern of offline devices. This includes cable modem termination system (CMTS) telemetry, power supply status (commercial power, on battery, etc.), and other network health telemetry that can be used for root cause analysis. Beyond internal network telemetry and spatial patterns, third party and augmented data sources can supplement the identification process and provide independent context to the root causes.

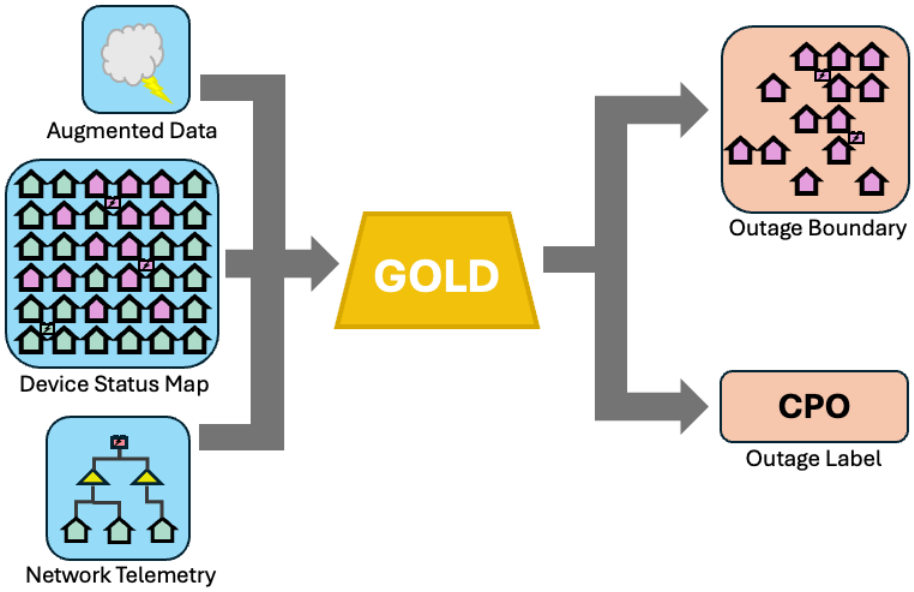

Figure 3: GOLD Data Flow

The data that can be fed into GOLD is flexible given the two-step process of outage determination and labeling. Within the flow of data, devices are grouped utilizing spatial relations and device status. Other sources including network alerts, device least common ancestors (LCA), and external data sources are fed into the system to label grouped outages.

These sources of data are independent of the device pattern and serve as another source of evidence for determinations. They can be tuned to the region, data availability, and necessity. For example, power supply status is a good indicator for electricity grid health when the affected supply indicates a loss of commercial power. The power supplies send telemetry that indicate whether commercial power is online, how long they have been running on battery power, and more fine-grained voltage information that is helpful in making power determinations. The biggest issue is when power affects customers without affecting the power supply or loss of power supply telemetry is not power related. This is where other devices' telemetry can improve analysis. As stated, third party sources can be combined with the existing telemetry to give insights that were previously unavailable.

Root cause analysis being independent from the device grouping stage of the algorithm allows for evolution of the root cause analysis engine to encompass more problem types over time. Additionally, this can be accomplished without rebuilding the system from the ground up when a change is made to either side of the algorithm.

Moreover, once a database of known outage groupings and root causes is created, it can be used to improve classification through machine learning and pattern recognition. Many outages may share the same signature, whether that is the location and timing or the specific features of the telemetry coming in. This allows a new outage to be matched to historical outages with the same characteristics. In short,

quicker outage labeling can be accomplished based on past experiences. The ability to filter, add, and adjust the weighting of inputs into the system, enables both greater accuracy and more rapid evolution of GOLD. Mostly simply, GOLD's modularity allows it to be constantly fined tuned and iterated upon.

3. Capabilities

3.1. Implemenation

Designing a system that can ingest telemetry, compute root cause analysis, and distribute accurate and timely alarms is essential.

The GOLD system ingests a large volume of modem telemetry necessary for analysis. This telemetry data is then associated with customer grouping data, forming the basis for detailed analysis. This process is crucial for accurately producing root cause alarms to be distributed.

A standard schema and distribution system were implemented to manage and distribute alarms across the business. The standardized approach allows for seamless modifications to the analysis engine upstream, ensuring that any updates and improvements do not disrupt the downstream consumers.

The alarms are stored in a central repository for query ability and are also streamed out to allow for downstream system to either consume alarms in real-time off a stream or to query periodically to allow maximum flexibility.

A detailed copy of the analysis results is also stored in a data lake. This approach not only allows for immediate evaluation by various teams but also provides a rich dataset for training future models, thereby enhancing accuracy over time.

The computation framework is designed to be both fast and robust. It can process the entire Comcast footprint in under thirty seconds every five minutes. This efficiency provides substantial headroom for future growth and the management of large-scale events such as hurricanes or other natural disasters.

Below, Figure 4 provides a simplified view of both GOLD's intake and distribution of data. It exemplifies the funneling effect of GOLD. GOLD intakes complex and disparate data to grow and enhance its capabilities; meanwhile, the distribution of the data remains constant regardless of upstream changes.

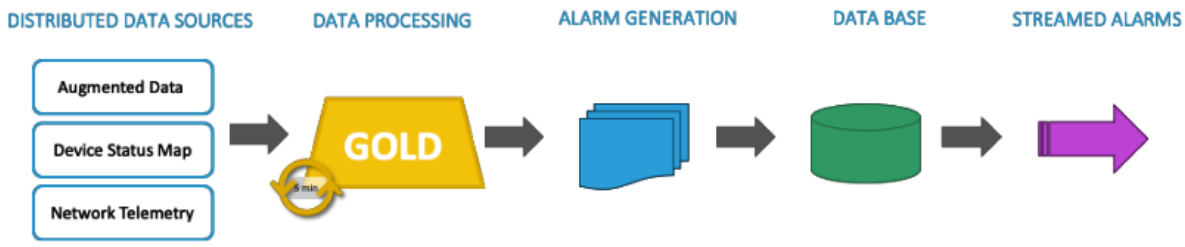

Figure 4: Alarm Processing Flow

3.2. Use Cases

Given the complexity of HFC access networks, technicians are trained for specific sections of the equipment and technologies. Thus, it is imperative that any automation for service outage detection and root cause analysis be as timely and accurate as possible. This allows the correct resource to be deployed to the correct outage source as quickly as possible, ultimately minimizing customer down time. The following subsections will discuss the core outage detection capabilities for GOLD with even more cases in development.

3.2.1. Commercial Power Outages

When there are issues with the power grid, customers feel the impact as soon as the failure occurs. Given Comcast equipment is dependent on the power grid, the impacted equipment's telemetry would reflect this failure. Consequently, the impacted devices can be grouped and labeled as a CPO. The pattern of devices offline, network telemetry, external data sources, and previous outage data inform GOLD's CPO determination.

One important issue that arises during commercial power outages is the interplay between HFC equipment and power. As discussed in Peck 2021, there are several methods for powering network power supplies when commercial power is down. One example discussed is the use of a backup battery. This allows customers to maintain connectivity to the HFC network when there is a commercial power failure. This is important as customers may maintain power in their home due to alternative power generators or other reasons.

With that said, all options outlined in Peck 2021 require either battery replacement or refueling. If this is not completed, the customer temporarily maintained connectivity will drop offline, constituting an HFC outage. This scenario will cause the two groups of devices to lose connectivity at different times. When commercial power failures and network equipment are affecting one another, temporal difference becomes critical in understanding the root cause.

For the second group of customers in the HFC outage, a network technician will be able to restore connectivity by hooking up a generator to the depleted network power supply. For the first group that lost connectivity due to CPO that not only impacted the network power supplies but also their in-home power, they will remain offline until their utility provider restores power.

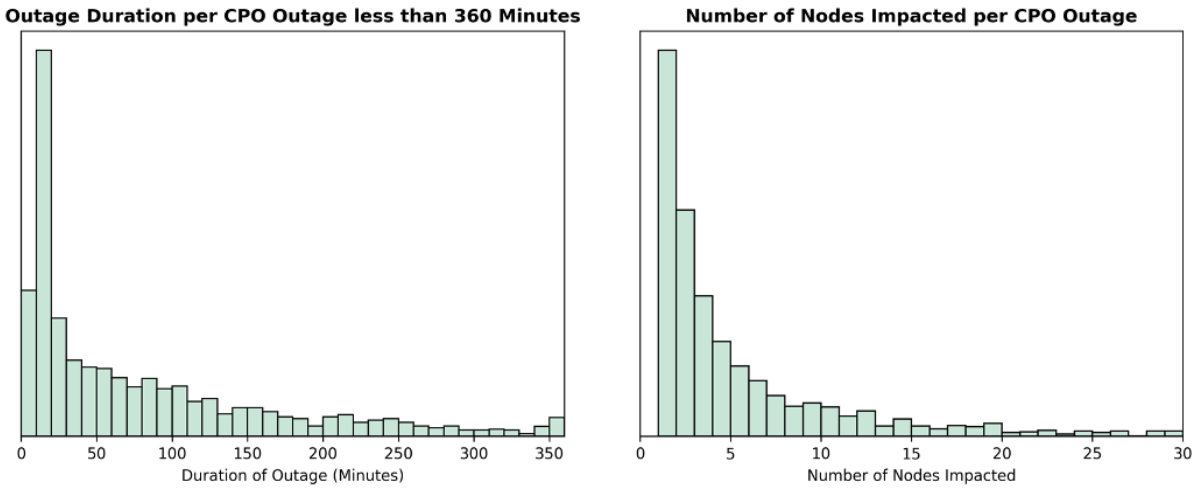

Figure 5: CPO Duration and Number of Nodes

Figure 5 showcases results from a full footprint trial running GOLD. Please note, the distributions are truncated for conciseness with the maximum number of nodes affected by a single outage at 151 and the longest duration being 20,540 minutes (about 2 weeks). As showcased, most commercial power outages are short and affect a small number of nodes. And while large and long outages may be rare, their identification is critical as they are the most urgent to remediate. By leveraging GOLD data, Comcast can dedicate resources as needed to these outages and re-use the historical data for training models.

GOLD also proved useful in detecting quick and small single node CPOs that may be missed by legacy systems that relied upon soaking periods. It is GOLD's ability to ingest and process data in near real-time that provides visibility into short power blips that happen at small scale. In such cases, detection becomes more proactive as there is no longer a need for manual processes or customer calls to be used in their detection. While these blips may not be immediately actionable, they can be used as a leading indicator for possible longer outages in the future. As such, Comcast can take preventive measures to minimize impacts to the customer.

While GOLD enables alarming for issues that may go unnoticed, its grouping logic drives down the alarm volume from CPOs relative to legacy systems. In cases where you see a CPO spanning multiple nodes, field teams would receive an equivalent number of alarms in most cases. However, with GOLD, regardless of the total impacted nodes, field teams would receive one grouped notification based on a shared cause, allowing teams to respond more tactically.

3.2.2. HFC Outages

Many of the benefits related to CPOs can also be translated to HFC outages: different causes can be identified; small outages blips can be detected and tracked; and the number of alarms will be reduced. However, there are other capabilities available for HFC outages.
First is the ability to identify a potential root cause of an issue even when there is a total loss of telemetry,

and devices are unable to communicate. When devices lose their connection and they are not able to transmit telemetry, historical outage data can provide insights into possible causes. This is because retroactively labeling such outages can be useful for training GOLD to recognize a loss of communication as known behavior of certain outage causes.

Additionally, GOLD's ability incorporate LCA information enables HFC determination for smaller scale outages. This is because small scale outages may not provide the requite level of detail to make an HFC determination. Passing the impacted device list through an LCA engine enables GOLD to understand how and if impacted devices share a common parent network element. If so, it could be indicative of an HFC issue. This can give the added benefit of not only determining whether the outage is HFC but also visibility into how far upstream in the network an issue is occurring. Algorithms for efficiently determining an LCA from a list of devices and a network topology graph are outlined in Lutz and Stehman 2023.

Figure 6 illustrates the typical breakdown for the types of network elements responsible for HFC outages. As can be seen, taps are the leading culprit followed by amps and passive devices with only about ten percent of HFC outages impacting the entire node.

This is invaluable information for operations teams as typically for outages technicians begin troubleshooting in some general area near the outage and sometimes even start all the way back at the node. Heading to the exact root cause element of the outage can easily save hours of remediation time as opposed to driving the plant attempting to find the network element of interest.

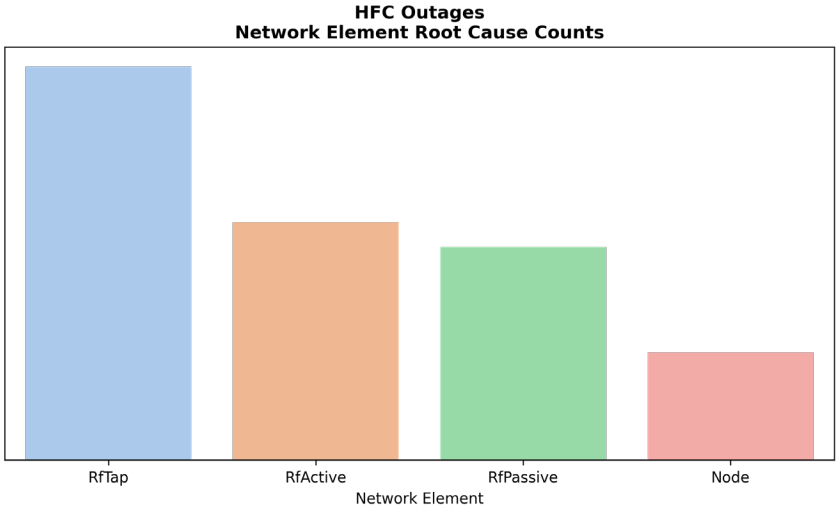

Figure 6: Network Element Root Cause Breakdown for HFC Outages

3.2.3. Fiber Outages

Modern HFC networks rely on fiber optic cables to connect the nodes in the field to the CMTS in the headend. For efficiency, a single fiber sheath leaving the head end will typically house fibers serving

many nodes. This well-known architecture is depicted in Figure 7. Any issues related to these fibers will result in entire nodes going offline, and depending on where the issue occurs, many nodes may be impacted.

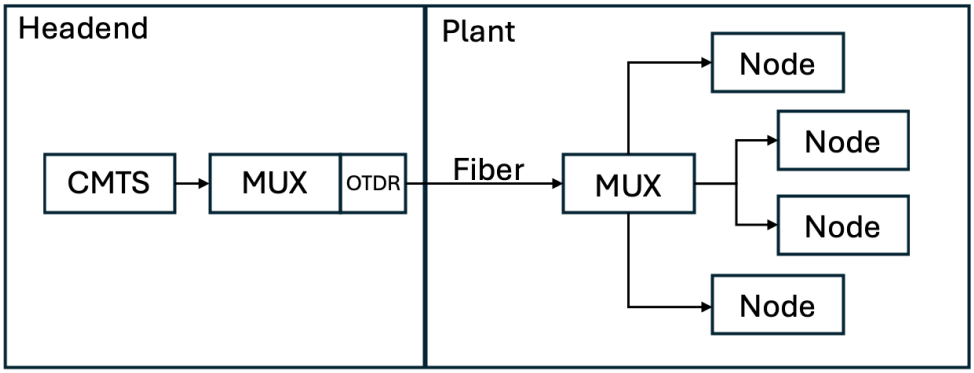

Figure 7: Typical Access Network Fiber Distribution

In node-based outage analyses, a fiber outage may result in many individual node level outages being declared, rather than a single multi-node fiber outage. Legacy processes require a significant effort to determine if the individual outages are all related to the same root cause. GOLD's ability to handle multi-node outages allows it to inherently detect outages where multiple nodes are offline due to a fiber issue between the headend and the node since the issue will meet the spatial and temporal requirements as defined in Section 2.1.

There are two main approaches to determine if an outage is caused by a fiber issue. The preferred approach is to integrate an Optical Time Domain Reflectometer (oTDR) in the headend as discussed in Mutalik et.al. 2020. The oTDR can detect fiber issues in real time - providing alarms for impacted fibers and CMTSs. If an integrated oTDR is not present, a database of the fiber connectivity can also be integrated into GOLD to determine if a fiber issue is the probable root cause. At Comcast, GOLD has successfully detected fiber cuts impacting hundreds of nodes while only generating a single outage alarm. In the previous node-based outage detection, hundreds of individual node-based alarms must be correlated together ultimately resulting in delayed action times and longer than necessary customer impact.

4. Operational and Business Benefits

4.1. Operational Benefits

The GOLD algorithm offers several key benefits for network operations, which include a reduction in overall outage time, better dispatches, and enhanced customer notifications.

The reduction in overall outage time was central to the vision for GOLD and supported by its enablement of better dispatches. As stated, pinpointing outage causes has traditionally been a highly manual and time consuming process. The process had three key focuses that are essential for dispatch: confirmation that the outage is legitimate, determining the broken network component that caused the issue, and isolating where to dispatch the technician. While GOLD provides insight into these three components, its

automated and data-based approach helps evade errors and inefficiencies inherent to the legacy process. By labeling the root cause swiftly and correctly, Comcast dispatches the appropriate fix agent to broken network elements to restore service more promptly. Figure 8 depicts legacy process relative to GOLD, where a red home icon signifies offline devices, and the different outcomes that can result for manual triage. Please note, while Figure 8 exemplifies GOLD's CPO detection capability, this example could also be applied to HFC outage causes like fiber cuts.

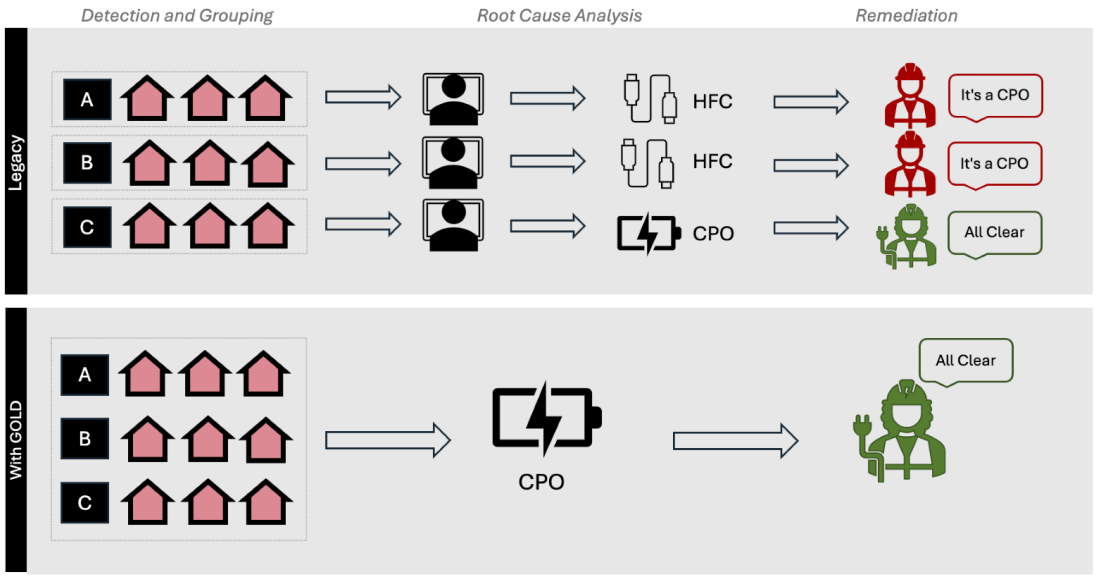

Figure 8: Legacy v GOLD Root Cause & Dispatch Processes

As explained in this paper, some of GOLD's key capabilities include outage detection, outage grouping, root cause identification, and the incorporation of LCA information. Together, they have a cascading effect on improving dispatches and outage duration.

First, detection and grouping logic allows Comcast to understand not only when a loss of connectivity occurs but also the scope of its impact. Understanding how widespread the issue is from detection allows Comcast to label the loss of connectivity as one issue rather than many. This information is important as the geographic spread of an issue can be used to understand its cause. This ability is especially meaningful for multi-node outages as it marks a common point of error in legacy triage processes, where one issue would be treated as numerous independent problems, leading to misidentifying the cause. In such cases, one issue may be assigned multiple fix agents that are incorrectly dispatched.

That said, GOLD's grouping logic enables better dispatches in another way: workforce prioritization. In the case of inclement weather, there can be widespread outages for many reasons like fallen trees that take down cable lines, CPOs, etcetera. Regardless of the reason, widespread outages can put a strain on resourcing, when the number of issues exceeds the available fix agents. In such cases, GOLD's grouping logic provides insight into how many customers are impacted by an outage. This is valuable information for prioritizing work, where issues with the largest customer impact may be prioritized most highly.

GOLD's benefits extend beyond its detection and grouping logic. Its automatic root cause identification is at the heart of better dispatches and a reduction outage time. First, identifying a root cause can be used to bypass soak periods, as it can certify that an issue is legitimate and requires remediation. Secondly, understanding root cause is critical for assigning the right fix agent. This is because some issues, like a fiber cut, require specialized fix agents. Finally, root cause information is also useful for workforce prioritization within and across causal categories. This is because some issues do not require immediate remediation as there are built-in mechanisms to temporarily maintain connectivity. For example, when a CPO is present, network power supplies have a backup battery to maintain connectivity for customers who have power in their home. So, in the case of severe weather events, CPOs where there is ample battery life may be prioritized after work to restore connectivity where cable lines have been taken down by wind or trees. Additionally, alternative and off-grid power sources present a challenge to ISP when there is a CPO. This is because they present an added layer of complexity where a customer may not initially appear to be impacted by a CPO due to the presence of something like a generator, but in reality, they are. In such cases, GOLD will not only detect that these customers are impacted but also this information can be helpful to prioritize truck rolls to set up a generator where there is a high saturation of alternative power sources.

While root cause information is helpful for understanding who and when to dispatch, information about the least common ancestor can help a technician arrive on site faster. For issues isolated below the node, technicians can be provided with exact locations of where to go to fix the problem.

Due to GOLD’s capabilities outlined above, Comcast can better serve its customers by providing the outage information needed for an effective and timely dispatch. That said, GOLD's data can also be used to enhance transparency to the customer while in outage. Most simply, the same insights provided by GOLD to field teams can be leveraged by proactive messaging tools. These tools can use GOLD's output to swiftly notify customers of connectivity losses. Additionally, root cause information can be used to augment estimated time to repair (ETR) predictions. Given connectivity has become an integral part of daily life, meaningful updates regarding connectivity losses are imperative for a better customer experience.

4.2. Business Benefits

While GOLD was born out of a focus on the customer experience, it offers ancillary business benefits. GOLD provides useful analytical data into areas of the network that are most vulnerable to connectivity issues and what is driving them. This information can be used to inform long-term facility planning. For example, in areas that are especially vulnerable to CPOs, Comcast can deploy an advanced power supply strategy to pre-emptively avoid power related issues. Likewise, in areas that are experiencing rapid growth, GOLD can provide insight into whether capacity constraints are a threat to connectivity and, if so, inform Comcast's distributed access architecture (DAA) deployments. As GOLD continues to expand its root cause identification capabilities, so does its ability to inform facility and network planning.

From a technical standpoint, GOLD aligns with Comcast's goal for rapid innovation. This is because GOLD's modular, yet centralized nature simplifies change management. On the backend, GOLD’s modularity allows it to expand its root cause detection inputs without fundamentally changing the core logic. On the front end, GOLD’s data distribution remains constant: one common access point, one standardized data schema. In this way, GOLD has a funneling effect. It takes in complex and ever evolving network telemetry and outputs simplified messages that can be actioned upon by various users.

This means as GOLD's capabilities expand, its users can quickly reap the benefits without completing an equivalent amount of integration work.

5. Conclusion

As outlined in this paper, the GOLD algorithm is a fundamentally different approach to understanding outages. It leverages device grouping and root cause logic to generate alarms for outages that are temporally, geospatially, and causally related. It was born out of a focus on the customer experience, where not only understanding the scope but also the cause of connectivity losses supports a reduction in overall outage time and customer transparency. However, GOLD enables both better facilities planning via its insights and Comcast’s goal for meaning and rapid innovation via its technical implementation.

Abbreviations

CMTS	cable modem termination system
CPO	commercial power outage
DAA	distributed access architecture
ETR	estimated time to repair
GOLD	Global Outage Localization & Detection
HFC	hybrid fiber-coaxial
ISP	internet service providers
LCA	least common ancestor(s)
oTDR	optical time domain reflectometer

Bibliography & References

Lutz, B. & Stehman, M. (2023). Graph Algorithms and Real Time Telemetry for Intelligent Plant Operations. *SCTE 2023 Fall Technical Forum*.

Mutalik, V., Rice, D., Spanbauer, R., Capuano, S., Gonsalves, R. & Gaydos, B. (2020). It's 10 PM: Do You Know Where Your Wavelengths Are? Continuous and Pervasive Monitoring of Optical Assets in the Access Domain. *SCTE 2020 Fall Technical Forum*.

Peck, T., Frankhouser, J. (2021). Ensuring HFC Network Resiliency During Extended Utility Outages. *SCTE 2021 Fall Technical Forum.*

