THE WiFi Happiness Index (WHIX)

A Technical Paper prepared for SCTE by

Krithika Raman
Vice President Product Engineering
Comcast India Engineering Center LLP,
Chennai One SEZ, 5th Floor, North Block in Phase II,
Module 7 & 8 - Chennai One,
Pallavaram-Thoraipakkam 200 Feet Road,
Thoraipakkam, Chennai - 600 097,
Tamil Nadu, India
+919152007903
Krithika_raman@comcast.com

Charles Moreman
Executive Director of System Architecture
Comcast Cable
1800 Arch Street, Philadelphia, PA 19103
Charles_ Moreman@comcast.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

1. Introduction

During the 2018 SCTE Cable-Tec Expo, panelists in a "birds of a feather" RDK session introduced the concept of a "Wi-Fi Happiness Index," abbreviated WHIX. Described as a weighted, realistic view of Wi-Fi behavior in a home, the WHIX index may involve more than 50 parameters, and very large quantities of machine-derived data. Since then, the WHIX concept has grown in breadth and depth, and is the topic of this paper.

The paper will cover what it takes to deliver per-device WiFi metrics, normally gathered across the industry, that contribute to an aggregate / whole house WiFi health assessment; what factors shape high and low WiFi performance; algorithmic detail regarding the non-linear aspects of whole-home "happiness;" and the value of customer feedback in informing health-centered WiFi metrics.

The paper will also cover best practices that can be proactively applied to WHIX assessments, in terms of recommended next steps. For instance, in situations where the WHIX score may indicate a suboptimal WiFi experience, a series of complementary algorithms can be run to evaluate different recommendations for improving the WiFi connection. For example, an algorithm can be run to determine if Wi-Fi extender "PODs" are recommended for this WiFi network. Attendees will also learn how WiFi health assessments can be applied to triage situations, and for engineering-level troubleshooting of new software releases.

1.1. The WiFi Happiness Index

WiFi is complex: The user experience of someone using WiFi is affected by multiple different WiFi conditions, occurring simultaneously. Even a WiFi subject matter expert can find it difficult to consider all the various WiFi conditions, assess the health of the WiFi network and determine how these conditions interact to affect the user experience. WHIX can be designed to process this complexity, with a goal of making it easy to understand whether a customer is having a good or bad experience. WHIX can also provide details on what is affecting each customer's WiFi experience.

WHIX can measure how well a WiFi network is working by modelling the WiFi user experience. It can provide details on specific WiFi conditions and how they correlate to good or bad WiFi connectivity. WHIX can assess the WiFi user experience at two levels: Per client device and for each WiFi network as a whole. Specifically, it can measure the WiFi connection on every client device in every network, every hour. It may also use this client device information to measure the user experience of each WiFi network. As API-based software, it may run in the cloud and can be integrated into other internal tools and UIs, including those used by field technicians to find and fix customer problems.

This may be used to assess the user experience for each specific WiFi client device. WHIX may also make it possible to understand the overall user experience provided by that WiFi network. Together, the per-device and per-network assessments may inform and model the WiFi user experience in a way that is intuitive and comprehensive.

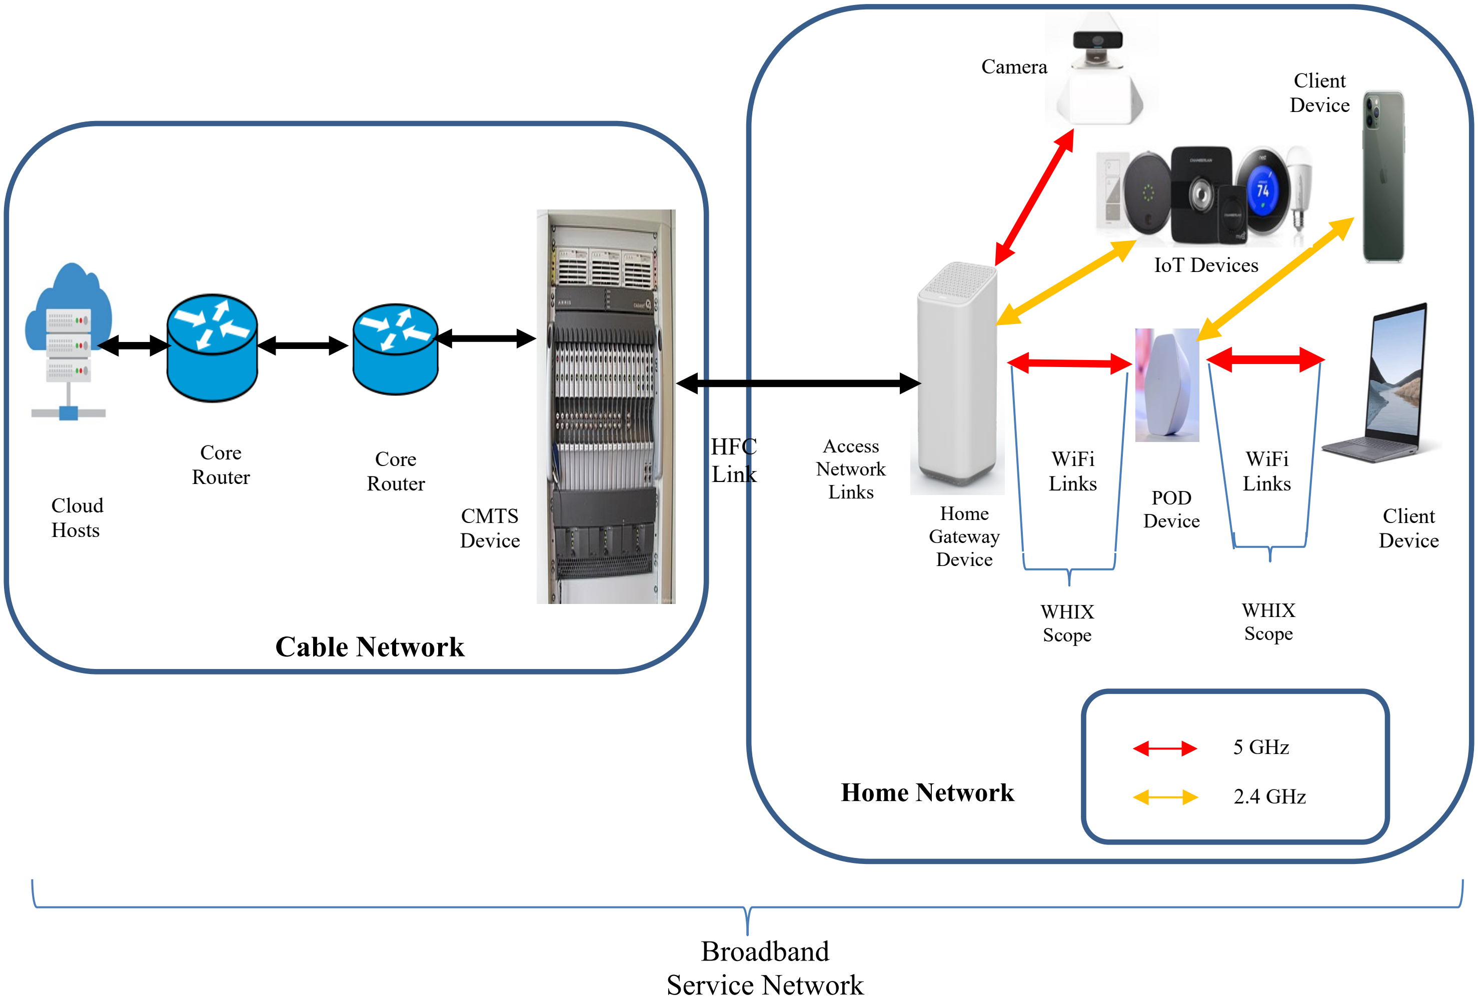

Figure 1 - The scope of the WHIX in an end-to-end broadband network

Figure 1 shows a typical end-to-end residential WiFi network. The cloud hosts the WHIX server components. The cable network component shows the typical network components, such as Core Routers and Cable Modem Termination Systems (CMTSs.) The "last mile" is over Hybrid Fiber-Coax and is connected to the Home Gateway Device (a WiFi router). Together, the Home Gateway and the WiFi extender "PODs" provide the WiFi service.

The WiFi router and WiFi extender PODs may use either dual-band (2.4 GHz + 5 GHz) or triple-band (not shown; 2.4 GHz + 5 GHz + 6 GHz) WiFi technology to provide high-speed connectivity. The WiFi extender (POD) is optionally used to improve the signal strength of WiFi in specific locations of the house. The various client devices shown in Figure 1 include a mobile phone, laptop, camera, and home automation devices, including a touch screen and various sensors. The Home Gateway device is connected to the POD, client laptop and to the mobile device(s) on the 5GHz channel; the touch screen and the camera over the 2.4 GHz channel.

The scope for the WHIX WiFi assessment is all the WiFi connections (both the 2.4 GHz and the 5GHz) that form the end-to-end path between the WiFi gateway and each WiFi client device. Other links, including the core network and access network, are not in scope for the WHIX WiFi assessment.

WHIX uses a machine learning (ML) model to process raw telemetry data and create WHIX “scores” that can model the user experience. Each score is based on a 0 to 100 range. 100 may indicate the best possible user experience, and 0 may indicate the worst possible user experience. In practice, most scores fall on a sliding scale between these two extremes. Different industry implementations could use different score ranges based on the granularity of the score that is desired and the data available.

Figure 2 shows a sample of score interpretations.

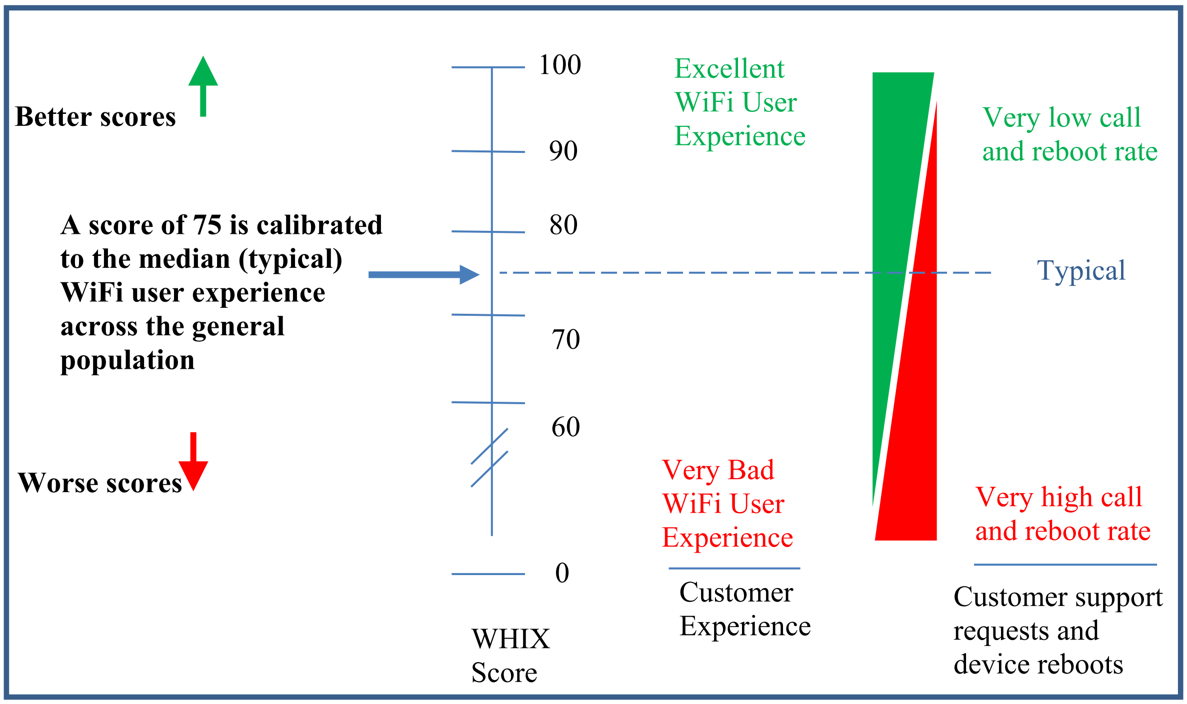

Figure 2 - The WHIX Uses Human-Readable Labels

The WHIX machine learning (ML) model can be trained using billions of historical WiFi data points collected over time from millions of WiFi networks. The ML model can be calibrated to provide a score of 75 when the user experience matches the median (typical) user experience. Scores greater than 75 may indicate a better than typical WiFi user experience. Scores less than 75 may indicate a worse than typical WiFi user experience.

A low WHIX score generally indicates a poor user experience and often results in a higher rate of customer support requests to the connectivity provider.

WHIX Scores may be mapped to human-readable labels. These labels may include:

- Excellent (scores >= 90 and <= 100)
- Very good (scores >= 80 and < 90)
- Good (nominal) (scores >= 70 and < 80)
- Bad (scores >= 60 and < 70)
- Very Bad (scores >= 40 and < 60)
- Extremely Bad (scores >= 20 and < 40)
- Worst (scores >= 0 and < 20)

Multiple factors simultaneously affect the user experience on a WiFi network. The WHIX assessment can consider both positive and negative factors to create these scores. When the positive factors are more significant than negative factors, the WHIX assessment score may be greater than the median score of 75. Conversely, when negative factors outweigh the positive factors, the WHIX score may be less than 75.

2. Telemetry data used by WHIX

One of the most common problems seen with WiFi is the change in the signal strength at different locations within a residence. WiFi signal strength can be a problem when the signal is too weak -- and can also be a problem when it is too strong. Weak signal strength can be caused by longer distances between WiFi devices and by obstacles that impact signal strength. Weak signals can be difficult for the receiver to reliably discern from background noise and interference. Very strong signals can sometimes occur when WiFi devices are too close to each other. Strong signals can overwhelm the dynamic range of some WiFi receivers and cause errors. WiFi tries to adjust for poor signal strength by changing how the WiFi signal is modulated. A lower signal modulation index can improve the reliability of the WiFi reception, but can have adverse effects, like lower throughput and higher channel utilization for a given bit rate.

Another common WiFi problem is high WiFi channel utilization. This is often expressed as the percentage of airtime used on the WiFi channel. The available airtime on a channel is finite, yet is used for multiple reasons. Airtime is used by the WiFi stations and access point(s) (APs) in the WiFi network being assessed; by WiFi networks in neighboring homes/businesses; and by non-WiFi RF sources such as baby monitors and microwave ovens. When a greater percentage of the available airtime is consumed, WiFi devices must contend with other WiFi devices to access the WiFi channel. As a result, the WiFi channel utilization increases. High WiFi channel utilization can cause the following undesirable scenarios:

1) *It becomes more difficult for WiFi stations and access points to send traffic.* During periods of high channel utilization, WiFi devices often wait for other WiFi devices to finish transmitting before the begin sending their data.
2) *There's a greater chance that multiple WiFi devices will attempt to begin transmission at the same time*. This can cause collisions, such that none of transmissions can be received.
3) *It can cause higher latency*. Higher latency means that packets incur longer delays. This reduces the user experience of delay sensitive applications like interactive games.

It follows that lower WiFi channel utilization can cause improved latency, improved throughput, fewer dropped packets -- and an improved customer experience.

Many other factors can affect someone's WiFi experience. These include downlink and uplink PHY rates, MCS (Modulation and Coding Scheme) levels, the background RF level (also called the "noise floor"), and Signal to Noise Ratio (SNR). Also impactful: The rate of packet re-transmissions on each client device; the rate of lost packets (when all retransmissions fail) on each client device; service availability (percentage uptime) of the access points in the WiFi network; and the ability (or inability) of WiFi clients to sustain the WiFi connection (measured as "rapid reconnects"). The total number of WiFi clients on each WiFi channel can impact performance, as can the rate of WiFi password failures, the configured maximum channel width, the operating channel width for each client device, media access delay for transmitted packets, the ability (or inability) of the network to steer each client device to a closer AP or steer that client device to a better WiFi band, QOS configuration/usage -- and more. It's a long list!

These WiFi factors tend to affect different WiFi client devices in different ways. This is because different types of client devices use WiFi very differently. For example: WiFi cameras use more uplink bandwidth (in the direction toward the access point) than downlink bandwidth. This makes WiFi cameras more susceptible to problems that impair uplink data transmission. WiFi media players and set-tops use more downlink bandwidth, so these are more susceptible to downlink-related problems. Other devices, like IoT sensors and thermostats, use very little bandwidth and can successfully operate at low PHY rates -- but must have extremely reliable connections. Still other devices, like game consoles, tend to be more sensitive to latency. WHIX may use a mapping of client device "fingerprint" data to categorize client devices into one of multiple WiFi device categories. WHIX may consider the WiFi device category when analyzing the WiFi user experience for each client device.

These factors can work in combination to affect multiple WiFi links used to create the end-to-end connection for each client device. For example: when a client device is connected to a gateway that provides both a WiFi Access Point and the home's WAN service connection, then only a single WiFi link is used in the end-to-end path between that client device and the WAN connection. In the case where a client device is connected to an access point provided by a WiFi extender POD, additional WiFi link(s) can be used to backhaul traffic to the WiFi gateway that includes the WAN service connection. In this case telemetry data from multiple WiFi links may be analyzed to understand the user happiness for that client device.

This telemetry data may be collected by both main WiFi gateway access points and WiFi extender POD access points. WHIX may import this data and use it to measure how well WiFi is working for each client device and for each WiFi network as a whole.

3. WHIX algorithm overview

The algorithms that can be used to create a WiFi Happiness Index may use a supervised ML training model based on this WiFi-related telemetry data. Telemetry data may be sourced from each of the Home Gateways and WiFi extender PODs, in millions of homes, and aggregated as anonymized high level statictics to train the ML model.

Client types spanning nine WiFi categories may be derived from device fingerprint data (see Table 1). In a typical iteration, the training model may examine >4K specific WiFi conditions over these nine different client device types. This model may form the core of the WHIX analysis and feed the outputs of the WHIX API.

Table 1 - Types of WiFi client device categories assessed by WHIX

WiFi Client Device Types Categories Assessed by the WiFi Happiness Index
1. Broadband Devices: Smartphones, laptops, tablets
2. Media Players: Smart TVs, streaming video devices, streaming dongles
3. Media Sources: Cameras, streaming servers
4. IoT / Printers: Thermostats, lightbulbs, sensors, and networked printers
5. Managed Set-tops
6. Managed WiFi Extender PODs
7. Audio Devices: Smart speakers, connected audio devices
8. Network Devices: Customer-owned WiFi extenders, etc.
9. Game Consoles

4. Training the Machine Learning Model

Multiple feedback mechanisms can be used to determine WiFi conditions associated with user happiness (or pain) with WiFi connectivity. These feedback mechanisms can include User Contact Rate statistics and Unscheduled Reboot Rate statistics.

Customers contact their service provider for many different reasons. Some contacts are for questions about billing or if the customer wants to upgrade/downgrade service. Other contacts occur when the customer needs help to resolve a WiFi-related problem. As an indication of WiFi-related user happiness, customer contacts may be filtered for contacts related only to connectivity problems. After filtering, connectivity-related contact rates may be compared with connectivity-related contact rates for the general (total) population of all customers. A higher rate of connectivity-related contacts may be associated with a poorer WiFi experience. A lower rate of connectivity-related contacts may be associated with a better WiFi experience.

Gateway reboots can occur for a variety of reasons. These may include “scheduled” reboots that coincide with software upgrades during a late-night maintenance window. Reboots can also occur during “unscheduled” times. Unscheduled reboots can occur as a result of power outages, software “self-heal” mechanisms and because a customer reboots the device. Many customers may reboot their WiFi gateway when experiencing a WiFi related problem. A higher rate of unscheduled gateway reboots may be associated with a poorer WiFi experience. A lower rate of unscheduled gateway reboots may be associated with a better WiFi experience.

WHIX ML models may be trained using statistical feedback from both user contacts rates and unscheduled reboot rates. Different ML models can be trained for different WiFi client device categories.

First, each WiFi-related condition (WiFi criterion) may be mapped into a range of values called a "bin". Bins may be constructed based on the distribution of telemetry data, so that enough Boolean "matches" occur within a bin range (where possible) to have sufficient data to avoid noise. For example: a bin may be used to characterize channel utilization (CU) criteria in the ML model might include all CU telemetry values greater than 22% and less than or equal to 25% for the client device category of "Media Player" devices.

Next, correlation may be measured for each WiFi criteria bin with customer contacts related to connectivity. For each bin (range of values for a single WiFi criterion), the percentage match rate in the population that made contact with the service provider for connectivity-related issues within the previous 24 hours may be calculated. This may be calculated by dividing the number of criteria matches for this bin range (which occur in the customer population that made a connectivity-related contact in the previous 24 hours) by the total number of telemetry reports received for any value of this WiFi criterion.

Then, the "Criteria Match Difference" (CMD) may be calculated to compare this match rate by this criteria match rate in the total general population. For example, if the criteria match rate for this bin range in the population with connectively-related contacts is 10% higher than the criteria match rate for this bin range in the general population, then Criteria Match Difference = 110%. CMD values equal to 100% may indicate that when conditions match this criteria bin range, there may be no effect on customer happiness. CMD values higher than 100% may indicate a poorer customer experience compared to the general population. CMD values lower than 100% may indicate a better customer experience compared to the general population.

The same process may also be applied using feedback from unscheduled reboots. This may establish correlation between each specific WiFi criteria bin range and customer happiness as indicated by unscheduled reboots.

This process may be repeated for more than 4,000 combinations of WiFi criteria bins across the 9 client device types.

Based on this correlation data, an algorithmic weight may be calculated for each of the 4,000 WiFi criteria bins. Positive weights may indicate WiFi conditions that correlate to WiFi user happiness that is better than the general population. Negative weights may indicate WiFi conditions that correlate to WiFi user happiness that is worse than the general population. Zero weights may be used to indicate correlation with the typical WiFi user happiness level (same user happiness as the general population). For each WiFi criteria, the weights for all bins in the ML model may be normalized such that positive weights offset (balance) negative weights across the total population.

The ML models may be re-trained for every WHIX software release. This may recalibrate WHIX to support and track with all new WiFi Gateway models, extender PODs models and software versions. Each new WHIX software release may typically also expand the ML model, to include new WiFi telemetry data that may become available since the previous release.

5. How WHIX uses the Machine Learning Model

After the ML model may be trained and validated, it can become available for use by WHIX. The following is a description of how WHIX could use this ML model.

On each telemetry interval (currently each hour), WHIX may use the ML model to assess the WiFi conditions for every client device in every home. First, all available WiFi telemetry can be collected from the WiFi gateway and from all extender PODs in that home. Telemetry for device fingerprinting can also be collected and used to map each client device into one of the 9 WiFi device categories. The WiFi telemetry data may be used to create a Boolean match for a specific bin range of each WiFi criteria.

- For example, if fingerprint data maps client device "A" to the WiFi device category of "Media Player" then the section of the ML table for Media Player devices can be used.
- Then all WiFi telemetry may be compared to the bins defined in the ML model. For example: telemetry for 2 GHz Channel Utilization is collected for all access points used in the end-to-end path for client device "A". If the telemetry indicates 10% channel utilization on the WiFi gateway access point and 42% channel utilization on a WiFi extender POD used in the end-to-end path for this client device, the larger (worst) telemetry value from these 2 WiFi link may be used.
- In this case, a channel utilization telemetry value of 42% may match the 2 GHz channel utilization criteria bin may be defined as "greater than 40% and less than or equal to 45%". Once this Boolean match is determined, the weight for this bin may be used to indicate the user impact of this individual WiFi criterion. In this case, the weight may be -2 for this bin range of Channel Utilization for media player devices. A weight of -2 may indicate that this condition may correlate with a very slight negative impact to the user's WiFi experience.

Next, the process may be repeated for all other telemetry data that was collected on all WiFi link(s) used for the end-to-end path to client device "A". Each additional telemetry data value may be processed by the ML model to add an additional weight for each WiFi criterion.

All weights for all criteria matches for client device "A" may then be summed and added to a value of 75 that can be used tonormalize the WHIX score to the typical score in the general population. This calculation may provide the WHIX score for client device A for the most recent hour. If the WHIX score is greater than 75, this may indicate that WiFi conditions on the links used for this client may provide a better-than-typical user experience. If the WHIX score is less than 75, it may indicate that WiFi conditions provide a worse-than-typical user experience. Previous scores may be retained to support historical views and trend analysis for this client device.

Next, the WiFi network for that home as a whole (WHIX account score) may be evaluated. This process may use the client device scores in that home to calculate an overall WHIX account score for each hour. This may involve analyzing the client device scores to determine which client devices provide the worst WiFi experiences in that hour. The set of these worst devices may then be used, along with an adjustment to recalibrate the median (typical) account score to 75, to create the WHIX account score for that hour.

WHIX can be analyzed over time. Various statistical analysis methods may be used to aggregate WHiX time series data and characterize the WiFi user experience over extended time periods. Statistical methods that provide the best predictor that the customer will/will-not contact the service provider and/or reboot their gateway device within the next 24 hours may have the best correlation with user WiFi happiness.

6. Benefits of WiFi Performance Monitoring

The WHIX concept may enable service providers to determine the quality of each customer's WiFi experience, per connected device and for every customer's account as a whole. This can enable service providers to see the most sample interval (based on available telemetry data) and also see a timeline that shows historical views, including any times when the customer's WiFi experience was materially better or worse. Data may also provide indications of what caused the customer's WiFi experience to be better or worse for each hour. The detailed data may indicate the root cause of what affected the customer's WiFi experience, and can be used to troubleshoot and remediate any problems. WHIX output data may be provided via an API that supports integration into various user interface tools.

Incremental software upgrades dispatched to Home Gateways and WiFi extender PODs may also have an impact on WHIX scores due to changes that have been incorporated into a software version. By monitoring highlevel statistics based on WHIX scores, service providers can determine if the WiFi quality has changed. If the WiFi quality has degraded, coincident with a software upgrade, the problem can be proactively identified, triaged, and remedied with new software that includes the appropriate fix. This may reduce customer contact rates and may help provide a quick fix to a potential customer-impacting problem.

Monitoring the WiFi Happiness Index may help in troubleshooting difficult issues that normally surface over time. It may also help to narrow down whether problems are intermittent or persistent.

WHIX can also be used with other cloud-based tools to create automated recommendations to improve the WiFi user experience for each home. This can include automated recommendations to perform specific targeted action(s) designed to improve the WiFi network in each individual home.

7. Conclusion

This paper illustrates a comprehensive concept for monitoring and assessing the user's WiFi experience at the per-user and per-client device level. The concept provides details that can be used to troubleshoot common customer WiFi problems. WiFi Happiness Index (WHIX) can be based on a Machine Learning algorithm that models how various WiFi conditions can affect user happiness with each of multiple WiFi client device categories, ranging from laptops and phones to IoT devices, for every account holder, every sample interval.

By using this concept, technician and care agents may identify specific client devices and specific accounts that are may experience WiFi-related issues, and, importantly, may recommend solutions to solve those problems.

WHIX can also be used to measure the aggregate user WiFi experience in WiFi gateway and extender POD software releases. This can be measured in early trial groups and/or in lab test environments. By comparing the aggregate WiFi experience on successive software releases, service providers can identify hidden software problems and prevent problematic software releases from being widely distributed. Conversely, it can be used to identify software releases that improve the user's WiFi experience and/or confirm the customer experience benefit from correcting previous software bugs.

Abbreviations

CMD	Criteria Match Difference
CMTS	Cable Modem Termination System
CU	Channel Utilization
HFC	Hybrid Fiber-Coax
ML	Machine Learning
RDK	Reference Design Kit
WHIX	WiFi Happiness Index
WiFi	Wireless Fidelity

Bibliography & References

802.11ax Wiki article: https://en.wikipedia.org/wiki/IEEE_802.11ax-2021

802.11ax spec: https://standards.ieee.org/standard/802_11ax-2021.html

WiFi Alliance: https://www.wi-fi.org