TVWS Network Resilience in Rural Towns: How External Factors Impact Network Performance and Social Factors Impact Network Longevity

Vaasu Taneja
University at Albany
United States
vtaneja@albany.edu

Munthir Chater
University at Albany
USA
mchater2@albany.edu

Tony Comanzo
University at Albany
USA
comanzo.anthony@gmail.com

Petko Bogdanov
University at Albany
USA
pbogdanov@albany.edu

Mariya Zheleva
University at Albany
United States
mzheleva@albany.edu

Abstract

Equitable access to information and communication technologies has become a basic human right. However, residents' ability to access these technologies significantly lags behind in areas with spotty or missing Internet connectivity. This issue is particularly pronounced in rural areas, where rugged terrain, long distances and sparse populations make it challenging to provision connectivity using high-speed wired backhaul. As a result, fixed wireless backhaul and last mile solutions, such as Television White Spaces (TVWS) and the Citizen Broadband Radio Service (CBRS) have emerged to connect populations where fiber and cable technologies are not economically viable. While fixed wireless can substantially lower the cost for Internet access, it often suffers performance issues that are poorly understood. Some factors affecting performance may be internal, stemming from network architecture, integration and frequency allocation. Other issues are external including weather and tropospheric effects, and the proximity of other in-band spectrum users, which in turn, may affect signal propagation properties or cause external interference.

This paper aims to shed light on the performance of Internet access networks that operate in the TVWS bands. We collaborate with a rural town in the northeast of the United States, where a TVWS network was used for 5 years — from 2016 through 2020, to deliver residential connectivity. Our data comprises a longitudinal trace of link layer and transport layer data. We also employ two datasets that capture historic weather and tropospheric ducting conditions. We show that higher wind speeds, precipitation, and temperatures correlate with worse network performance on the TVWS network. We also find that antenna cross-talk and non-ideal antenna orientation will negatively impact network performance. We study socio-economic factors leading to the network's termination. In light of our findings, we recommend potential improvements to rural network topology design and management.

CCS Concepts

• **Networks** → **Network performance analysis**; Network manageability; Network reliability; • **Human-centered computing** → *Empirical studies in ubiquitous and mobile computing*.

ACM Reference Format:
Vaasu Taneja, Munthir Chater, Tony Comanzo, Petko Bogdanov, and Mariya Zheleva. 2024. TVWS Network Resilience in Rural Towns: How External Factors Impact Network Performance and Social Factors Impact Network Longevity. In *The 13th International Conference on Information & Communication Technologies and Development (ICTD 2024), December 09–11, 2024, Nairobi, Kenya.* ACM, New York, NY, USA, 20 pages. https://doi.org/10.1145/3700794.3700811

ICTD 2024, December 09–11, 2024, Nairobi, Kenya

ACM ISBN 979-8-4007-1041-4/24/12
https://doi.org/10.1145/3700794.3700811

1 Introduction

The Rural-Urban digital divide refers to the lack of broadband connectivity for rural areas compared to urban ones. In the US, the FCC reports that 28% of rural residents did not have access to fixed broadband at speeds at or better than 100/20 Mbps in 2022 [18]. Internationally the ITU reports that only half of individuals in rural areas use the internet compared to 81% of those from urban areas [61]. This gap is a problem as rural populations do not have equitable access to online services, including health and educational resources as well as opportunities for economic development [2, 60].

Researchers and policy makers have attempted to address this divide in a number of ways. From the policy side, policy-makers around the world have created universal access programs and funds to provide an economic incentive for rural development. Due to rural communities being inherently less dense than urban ones, the cost of bringing broadband connectivity to residents by traditional methods is higher [26, 38]. The ITU explains, "In rural area with low population densities, providers are deterred from investing in telecommunication projects since building broadband infrastructure is expensive, takes a long time to pay back and the returns from a small customer base are perceived to be unattractive" [60]. The universal access programs grant funding to provide broadband in otherwise neglected areas. Recent legislation in the United States established the BEAD act [36], which awarded $42.45B to states and providers to build out broadband infrastructure in rural areas that otherwise might not be seen as profitable. The fund has thus far supported over 270 grant recipients across 56 states and territories [3] and has substantially contributed in shrinking the Internet availability gap between rural and urban areas. Counterpart programs exist globally including in Europe [13], South Africa [62], and Asia-Pacific[58].

Researchers have developed multi-step approaches to designing and implementing broadband network solutions in rural areas [16, 59]. Due to different cultural and societal practices, as well as challenges with terrain, infrastructure, equipment maintenance, and budget constraints, these approaches differ from those of urban areas. Consequently, alternative network architectures and technologies arise to address the urban-rural digital divide. Communities and have used Mesh Networks [23], Low-Earth Satellites [6], Low Powered Radio (LoRa) [65], Long Distance Wifi [40], and Television White Spaces [67]. Each solution has its own benefits, drawbacks, and deployment challenges.

Our research concerns one of these alternative network types: Television white spaces (TVWS), which is seen as a viable alternative to provide fixed wireless services for rural broadband [8, 67]. As television broadcasts decline, much of the spectrum band allocated for television is left unused. These unused bands are called white spaces, hence the name, "Television white spaces". Rural communities have utilized TVWS for their own internet networks [32, 33, 35]. Although, these networks have many upsides due to the open spectrum and usage on the lower frequency UHF band, there are still some questions about their viability when deployed for real world networks. When television network solutions are deployed at-scale, they face additional challenges. External factors like weather, topography, and cross-talk may impact the networks in ways that the network architects do not expect. In addition, there may be challenges with network maintenance, funding and community support. These additional challenges mean that potential solutions to the digital divide may need to be rethought in order to be successful. For our research, we examined a TVWS network deployment in Thurman, NY, a rural town in the northeast United States. Our goal was to assess the network's performance and the reasons for its ultimate termination. We studied network design choices, such as antenna collocation and relative location between client and base station, and evaluated their effects on performance. Additionally, we analyzed network data to determine how external factors impact TVWS networks considering atmospheric effects and weather conditions. Based on our analysis, we recommend potential changes and additional considerations when deploying networks in rural areas.

This paper makes the following contributions:

- Our work is the first multi-modal study that combines weather and atmospheric data with multiple years of TVWS network traces. We find that (1) higher wind speeds, precipitation rates, and temperatures correlate with worse network performance on the TVWS network. (2) TVWS networks performance is not affected by minor tropospheric ducting effects in the northeast United States. However, areas further south have more widespread ducting that could effect a network.
- We quantify the transport layer performance of users across the network and discuss its impact on user experience.
- We analyze antenna cross-talk and non-ideal antenna orientation and find a negative impact on the network's bit-error-rate. Based on our findings, we see a trade-off between the performance of the network with the cost of network deployment and maintenance.
- Using the minutes from community town-board meetings, we study the impact of socio-economic factors on network longevity, and make the following recommendations for rural network topology and management. (1) Ensure maintenance of technology does not rely on a single person. (2) Before deployment, carefully survey community support and enthusiasm for the given project. Mixed support will inhibit chances for long-term success.

2 Related Work

2.1 TV White Space in Rural Areas

Researchers across the world in Africa [8, 19, 35, 39], Europe [5, 44, 63], Asia [15, 24] and North America [12, 67] have explored the potential for employing TVWS in rural communities to improve network connectivity and access. They consistently discuss three primary benefits to TVWS technology. First, they refer to the amount of available spectrum. Research across the globe shows that UHF bands for TVWS are consistently available in rural areas [5, 15, 26, 39, 63, 67]. While there are some variations in number of channels available and some challenges with ensuring that users do not infringe on primary users, these variations and challenges have been consistently overcome. The available spectrum relates to a second benefit of TVWS: the lower cost. As a secondary user on the UHF band, TVWS network providers do not need to pay for their spectrum allocation. This means that the cost to enter into a community is lower. In combination with lower cost infrastructure, these networks can provide more communities access to internet for a lower cost than traditional networks [15, 26, 39, 67]. Finally, researchers consistently cite the better propagation distance and penetration of the UHF band as a third motivation for utilizing this technology [5, 8, 12, 15, 19, 35, 39, 44, 63, 67]. Although there has been significant research into the applicability of TVWS networks in rural areas, there is less work analyzing deployed networks. Our research fills this gap. We hypothesized that TVWS deployments, with their fixed modulation and coding rates, are potentially susceptible to performance deterioration in cold climates due to less hospitable weather and less ability to adapt to lower SNR. Of the work on deployed TVWS networks, no research considers the impact of snowy climates on network deployments and few have analyzed data encompassing as long of a time period as we did. Finally, while research has been done into theoretical deployment designs to minimize interference among and between network access points and incumbents [5] [67], our research fills in a gap by analyzing the direct effect of base stations' interference with each other on real users' network performance.

2.2 Performance Evaluation of Rural Community Networks

Researchers have tried to leverage wireless architectures, spectrum sharing, and new technologies to create networks to meet their communities' connectivity needs. As there is a large diversity of challenges for rural communities, these approaches tend to be designed to address challenges of a specific community or set of communities. The connectivity solutions are evaluated based largely on the following metrics: energy efficiency [14, 25, 68], reliability [32, 55, 68], network statistics [27, 28, 41, 64, 68], cost [27, 28, 68], and community impact [32, 64, 68]. These metrics all relate to the core challenges of networking in rural areas.

For our research, we analyzed a rural community network deployment in the northeast United States. This TVWS network was funded with grant money from the state government. Our research team was not involved in the deployment nor management of the network, but was given access to link layer and transport layer data over the course of multiple years. Our research is the first long term (multiple years) analysis of a TVWS deployment with network traffic from users. We evaluated the network based on its network statistics, reliability, and resilience to external effects. Using the network clients' bit-error-rate (BER) and signal-to-noise ratio (SNR) over time, we studied how resilient the network is to weather effects. We also considered factors impacting network longevity. A key concern for this network were social factors relating to the network's budget and the community's acceptance of the network. These factors impacted the management and maintenance of the network. Although researchers cite energy concerns as a challenge in rural environments [14, 25, 68], that was not a challenge for this rural community which had a reliable power grid. By analyzing this unique set of factors that needed to be balanced, our research makes a new contribution to the field.

2.3 External Effects on Wireless Networks

TV White Spaces technologies operate in the Ultra High Frequency (UHF) spectrum. Researchers have previously studied how weather effects impact the radio signals in UHF. Previous work found that the received signal strength for UHF radio decreases as temperature, humidity, and pressure increase, and that wind has no effect on signal strength [7]. In another study of UHF signals, the researchers concluded that received signal strength increases as temperature increases and pressure has no clear effect. They found that signal strength decreases as humidity increases [57]. A third study found that wind and precipitation have a negative impact on signal strength [29]. Finally, in another paper, researchers state that precipitation has a minimal effect on signal attenuation when tested on 2.1 GHz signals and rainfall over a four year study [34]. Based on these studies and their contradictory findings, we think that other environmental factors and deployment specifics may influence how weather affects signal strength. These studies were conducted in areas with different climates and environments, and their findings may be limited to the environment that they were studying. None of these papers discussed weather impacts in temperate climates like that of the northeast United States, and they do not study the impact on active TVWS networks. There is one recent study that considered how external effects impacted a TVWS deployment in a farm setting [42]. They found that higher placed antennas were preferable and that humidity may impact signal strength due to how crops absorb water at humid times of the day [42].

Unlike previous work which focus on signal strength in given conditions, we quantify the impact of these external factors directly on TVWS clients' BER and SNR. This allows us to determine how atmospheric effects and spatial orientation of the network architecture impact the usability of the network.

3 Network Overview and Approach

We now provide background on the history and architecture of the community network we study. Furthermore, we provide details about our data collection and processing efforts.

3.1 The Community Network

Brief History of the Network. The construction of the TVWS network in Thurman, NY was initially funded by the ConnectNY Broadband state grant of 200,000 dollars [1]. The grant stipulated that the town must maintain the network for five years and the grant would provide funding for the initial setup of the network [53]. This initial setup was done by a network engineer whom the town contracted to design and deploy the network. Residents who wanted to get internet services through the network paid an initial setup fee and then 50 dollars per month for its use [22]. The town of Thurman turned on the Television White Space network in 2015, and residents started to use it the following year. The network was managed by the town's board which meant that as the board and their respective priorities changed, so too did the level of support for the network. Throughout its lifespan from 2016 through 2020, there were difficulties maintaining equipment and funding the network. It ultimately operated at a loss for multiple years before being shut down in the beginning of 2021 [54].

Network Topology. The TVWS network is comprised of six base stations serving 34 clients throughout its life span. Fig. 1 shows the network topology. Note that only 29 of the 34 clients are shown in Fig. 1 as we do not have the location for the other five active clients. Also, note that the network was located in a mountainous terrain, but we do not have data about the antenna heights or potential obstacles obstructing each client. Base stations 1-1, 1-2, and 1-3 are a tri-sector deployment that covers 360 degrees, whereas the remaining three base stations are on a 120-degree sector each. Each TVWS base station operated on a 6 MHz bandwidth channel that was chosen by the network operator to minimize interference with the other base stations and with any nearby television stations. Although station channels had a bandwidth of 6 MHz, the base station hardware could only utilize up to 4.5 MhZ of this channel [56]. The gateway router for the network was connected to the tri-sector base station tower. The other three base stations sent data to the gateway routers through Point-to-Point transport links. For base stations 4 and 5, data was transmitted from the station to the gateway using a 900 MHz Point-to-Point transport link. Base station 6 used a 900 MHz link and an additional TVWS link operating at 482 MHz (See Fig.1). This 482MHz Point-to-Point link used channel bonding on two neighboring channels and had different hardware so it could reach 10 MHz of bandwidth. The network backhaul meant that clients for stations 1-1, 1-2, and 1-3 sent their data

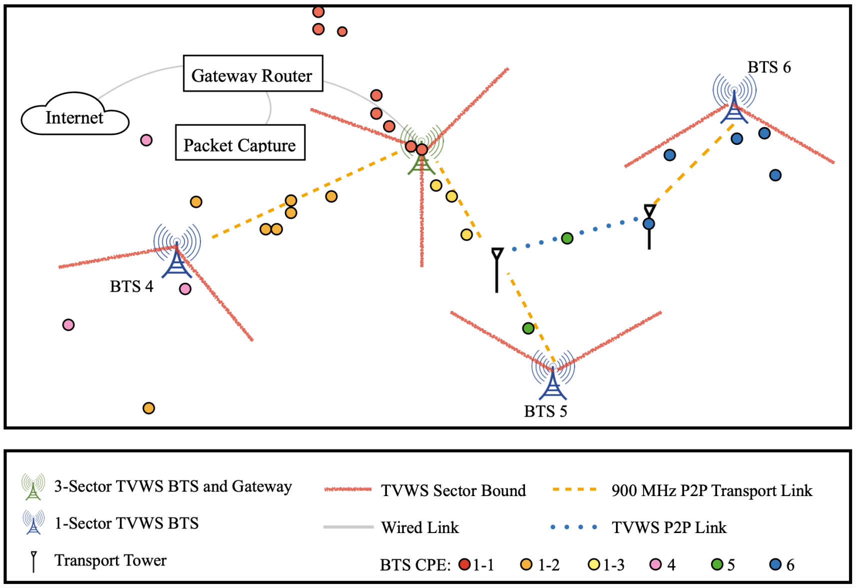

Figure 1: This is a map of Thurman's TVWS network. Base stations 1-1, 1-2 and 1-3 are on the tri-sector tower. Clients (CPEs) are color-coded based on their paired base station. Packet capture was done with a port mirror at the network gateway that sent network traffic to a monitoring computer for capture and storage.

Figure 2: On top, a photo of the antennas for the tri-sector tower (base stations 1-1, 1-2 and 1-3). Below that, a photo of a client antenna.

directly to the gateway. Clients of stations base stations 4, 5, and 6 sent data to their base stations which then sent the data to the gateway router through the Point-to-Point links.

Network Hardware. The rural town's network used hardware developed by Carlson Wireless Technologies. The base station and CPE terminals were RuralConnect® Generation II equipment [56]. The base stations had a transmit power of 28dBm with Single Input Single Output (SISO) antenna systems. The CPEs had a transmit power of 27dBm, a beamwidth of 45 degrees, and used SISO antenna systems too [56]. The equipment manual states the operating temperature for the equipment is between -30C and 55C and up to 95% humidity [56]. The backhaul TVWS Point-to-Point link used equipment from Redline instead of Carlson Wireless. It had a transmit power of 18dBm and had capability for Multiple Input Multiple Output (MIMO) antenna systems.

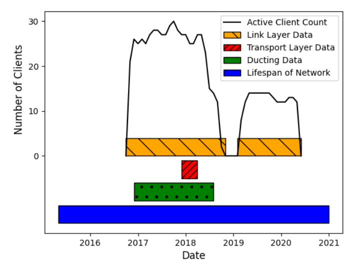

Figure 3: Timeline of the network and our datasets. The client count is based on our link layer data. The period in 2019 with 0 clients is therefore due to missing data. We have weather data for the entire lifespan of the network.

3.2 Data Collection and Processing

Throughout the life of the TVWS network we collected two types of data: (i) link layer traces from the base stations, and (ii) transport/application layer traces at the network gateway. We also retrieved general weather data for the entire lifespan of the network from a nearby weather station[4] and tropospheric ducting data from DX Info Centre[21]. As each data set was gathered from a separate source, we have data from different time periods. In Fig.3 we present the time periods for the datasets that we have used. Our link layer dataset encompasses data from varying number of clients over an approximately three year time period with a brief gap in the middle. The client number changed as clients joined or left the network, or as they malfunctioned and were repaired. We have a window of about four months of transport data from November 2017 to March 2018. The ducting data spanned from December 2016 to August 2018. Finally, as mentioned we have weather data for the entire lifespan of the network.

Furthermore, our data is of different granularity. Specifically, the link layer metrics were reported by the network approximately every six minutes, transport layer data was collected on a fine per-IP packet granularity, weather data was available at a 10 minute granularity, and the tropospheric ducting at 3 hour granularity. We unified the granularity of the individual datasets to 10 minutes in order to facilitate multimodal analysis. In what follows we detail our data preprocessing.

Link Layer Data. At the data link layer, we collected and analyzed data from the network spanning a 39 month period from 2016 to 2020. For each client link, the TVWS base station reports six metrics: aggregate bytes sent, aggregate bytes sent with errors, and the average SNR in the uplink and downlink directions. Each metric was reported at a cadence of approximately 6 minutes. For each client and for every ten minute period, we aggregated the total bytes sent, and bytes sent in error. If a raw 6-minute period overlapped with two 10-minute bins, the aggregate bytes were allocated proportionally to the amount of temporal overlap. From these two quantities, we calculated the BER as the ratio of erroneous to total bytes sent in a 10-minute period. Additionally, we calculate the average SNR on a 10-minute cadence considering all SNR reports that fall within a 10-minute period.

Transport Layer Data.

At the network gateway, we collected pcap traces comprised of the first 120 bytes (i.e. the packet header) for each packet traversing the network. This approach was useful in providing sufficient information about the network performance without infringing on users' privacy, as the actual data they exchanged was not captured by our measurements. These traces facilitate our evaluation of transport layer protocol performance and user experience in Section 4. To parse the data, we developed a command line tool utilizing Tshark, the command line counterpart of Wireshark. Our goal was to assess packet-level and flow-level data from the network. We selected specific Wireshark fields for network analysis (see Table 1), stripped the relevant data from the .pcap files, and converted the data into .csv files. We used the converted .csv files to calculate statistics on a per-packet or per-flow basis. Our analysis specifically focuses on performance of applications that use reliable TCP transport, which amount to 90 percent of the overall traffic seen on the network. Therefore, we use fields like 'tcp.analysis.retransmission' to indicate a TCP packet being retransmitted, 'frame.len' for a packet's size, and 'tcp.time_delta' for the time passed between packets in a TCP flow. To understand the packet flows, we filtered for TCP conversations using the Tshark flags, '-r <filename> -N dm -q -z conv,tcp'. With this filter, we could separate each flow and calculate the total flow counts and sizes.

Weather Data. Weather conditions have been shown to affect the propagation properties of UHF frequencies, as discussed in Section 2. Since TVWS networks operate in these bands, we sought to understand the effects of weather conditions on the network performance. To this end, we used data from the NYS Mesonet, a regional weather monitoring system which has 126 stations [4]. We utilized data from Chestertown, the weather station closest to the TVWS network at about 10 miles (16 km) away. Although the weather may vary over that distance, we assume that the geographic features and proximity between the two points are close enough for that weather data to be suitable for our analysis. For every ten minute interval between November 2016 and January 2021, the weather station provided data about weather features such as temperature, precipitation, humidity, pressure, and wind speed. Using the timestamps of the weather data and our network data, we joined the data tables together. The combined dataset supported analysis of the network performance as a function of the weather conditions.

Tropospheric Ducting Data. Tropospheric Ducting is an effect that can occur under certain atmospheric conditions and can impact radio waves. In a vacuum, radio waves will travel in a straight line. However, in the atmosphere, these waves will travel in an arc as they pass through the air before escaping into space. Under certain atmospheric conditions, they may get refracted enough to bend back towards Earth causing potential interference to distant base stations using the same frequency range [45]. This effect, known as tropospheric ducting, can cause interference and potentially impact the quality of a TVWS network.

To measure the potential impact of tropospheric ducting on this TVWS network, we utilized historical ducting data from DX Info Centre[21] for the period December 2016 through August 2018. The data set contains a series of maps (JPEG images) of the eastern US that were color-coded based on the estimated refraction over a given area. Each pixel on the map covered an area of about 40km^2. As the network could fit within a pixel, we assumed that the intensity measurement across the network would be the same.

The color intensities of the "tropo-index" were integer values ranging from zero, meaning no ducting, to ten, meaning widespread ducting, though local topography could impact the reliability of this measurement [21]. For each map, we converted the color value at the town's location back into its intensity value. Each map in the set represented a three hour time period, and we assumed that the intensity was consistent over this time.

4 Overview of Network Usage

In this section we discuss the quality of the transport layer of the TVWS network. We present different network metrics for the transport layer of the network and assess how well the network is supporting users. We found that performance across the network's base stations was relatively similar and that users of the network had better experience in the downlink direction than the uplink.

For our transport layer analysis, we used data that we collected from November 2017 through March 2018. We analyzed packet and flow data to understand the general performance of clients on the network. In Table 2, we present the packet and flow sizes separated by base station. We also present packet statistics. The majority of the network traffic came from clients of the tri-sector base stations (1-1, 1-2 and 1-3), with the most traffic belonging to station 1-1. Despite the higher traffic, clients of these base stations performed similarly to other clients in the network.

We see that downlink packet and flow sizes are about an order of magnitude greater than uplink packet and flow sizes. Other papers have also observed this discrepancy between packet and flow sizes in the uplink versus downlink direction in various contexts [30, 66], which we may expect since Internet users tend to download more than they upload. In the last row of the table, we present the percentage of control packets in our traces. Control packets are small packets that contain information that the network uses to manage traffic. We define control packets as any TCP ACK, FIN, SYN, or retransmitted packets and counted control packets by using Wireshark's built-in fields. We found that approximately 27% to 41% of base station packets were control packets meaning the network spent significant resources maintaining connections with some base stations (e.g. base station 5 with 27.58% control overhead) doing better than others (e.g. base station 4 with 41.4% overhead).

Next, we present uplink and downlink cumulative distribution functions (CDFs) for the round trip time (RTT). The RTT is the amount of time it takes for a network request to be sent and then a response to be received. Therefore, this metric can be used to

Table 1: Wireshark fields used.

Wireshark Field	Description
frame.len, tcp.len	Packet size and size of the TCP data.
ip.src, ip.dst, ip.srcport, ip.dstport	Source and destination IP addresses and ports. Fields were used to identify specific flows.
tcp.flags	Field was used with the flags ACK, FIN, and SYN to identify control packets.
tcp.analysis.retransmission	Filters for retransmitted packets.
tcp.analysis.ack_rtt, tcp.analysis.bytes_in_flight	Round trip time and bytes in flight. Fields were used to calculate the respective distributions in Section 4

Table 2: Transport Layer Data for Clients. Nov 28 2017 through Mar 28 2018. RTT is the round trip time. Percent Retransmission refers to the percent of packets that were corrupted or lost and needed to be sent again.

Base Station		1-1	1-2	1-3	4	5	6
Num. Active Clients		8	5	3	2	2	4
Total Packets x 10^8		21.8	9.57	12.5	3.94	2.66	11.2
Total Flow Count x 10^6		35.8	2.76	5.50	2.99	2.64	3.99
Average Packet Size (KB)	UL	.191	.146	.149	.196	.205	.260
	DL	1.41	1.46	1.43	1.32	1.28	1.40
Average Flow Size (KB)	UL	15.0	17.1	11.3	10.7	5.86	28.5
	DL	152	246	147	78.6	47.7	196
Percent Retransmission	UL	1.31	0.64	0.96	2.12	1.69	1.46
	DL	5.3	6.44	4.94	5.58	4.88	4.03
Percent Control Packets		37.67	35.62	33.54	41.4	27.58	38.78

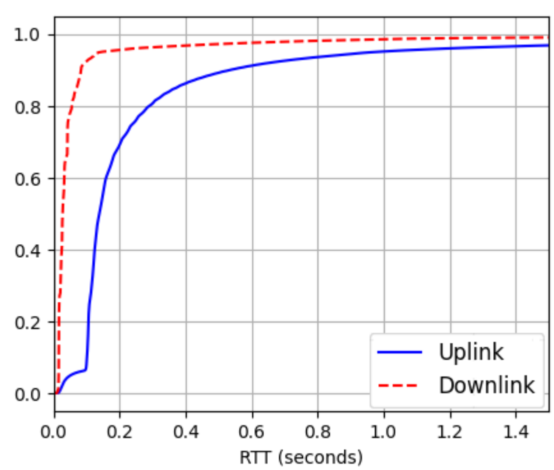

Figure 4: CDFs for uplink and downlink Round Trip Times. The uplink CDF has a long tail and reaches 100% past 5 seconds.

understand the timeliness with which the network is able to handle user traffic, which has a direct impact on the performance of the TCP protocol. Specifically, longer RTTs lead to TCP timeouts and a subsequent inability for the protocol to maintain high user throughput. In our data, RTT is reported with every TCP acknowledgement (ACK). ACKs can occur in both directions within a flow: the initial flow establishment request is ACK-ed by the server and then data packets are also ACK-ed by their recipient. Thus, we differentiate uplink RTT from downlink RTT based on the direction of the ACK. This helps us evaluate the timeliness of session establishment and data exchange. For the TVWS network, we saw that the median uplink RTT is about 140ms, whereas that in downlink is about 30ms. These values are good RTTs for most common applications. However, considering the CDF for the uplink RTT (corresponding to Internet downloads), there is a long tail that continues past 1.5 seconds for 3.1% of packets and past 5 seconds for 0.5% of packets. Many applications require round trip times at least below 300ms [10], to function without users noticing the impact, but 20.3% of packets had uplink round trip times above this threshold.

Finally, we consider bytes in flight to evaluate the performance of the TVWS network at the transport layer as well as evaluate how clients' link layer quality affects their transport layer performance. Bytes in flight is a measure of how many bytes have been sent by TCP without being acknowledged yet. It is closely tied to the window size as larger window sizes will lead to larger quantities of bytes in flight, and subsequently, a higher user throughput. Note that higher quality links will have larger window sizes and in turn higher quantities of bytes in flight.

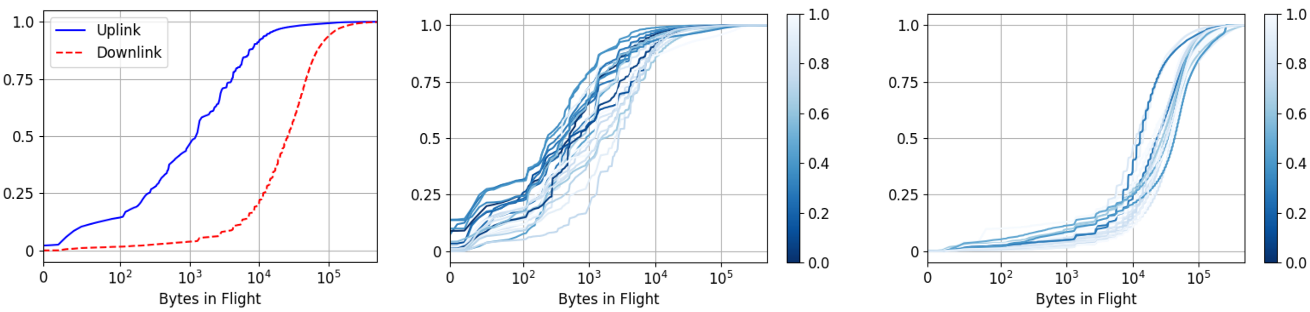

(a) Combined Network BIF

(b) Uplink BIF Colored by Frequency of Time Meeting Up BER Threshold

(c) Downlink BIF Colored by Frequency of Time Meeting Up BER Threshold

Figure 5: CDFs for Bytes in Flight. 5a shows the uplink and downlink bytes in flight for clients grouped together. 5b and 5c show CDFs for individual clients in the uplink (b) and downlink (c) directions. They are color coded by the percentage of time the client meets the 10^{-5} threshold for BER; a lighter color means that the client more frequently had a low BER.

In Fig. 5a, we present CDFs for bytes in flight for the entire network in the uplink and downlink directions. Notice that the downlink bytes in flight tended to be larger than those in uplink. This is expected as there is usually more downlink traffic than uplink. We also see that the median bytes in flight in the downlink direction was over 10^4 while in the uplink direction it was about 10^3. This difference concurs with the results in Table 2; the average downlink packet sizes are about an order of magnitude larger than those in the uplink direction. Furthermore, referring again to Table 2, there is a higher retransmission rate in the downlink direction. Therefore, the higher bytes in flight quantity likely leads to the more frequent time outs or packet losses in the downlink direction.

Last we consider how bytes in flight is related to BER. For this comparison, we color code each client's bytes in flight CDF by the percentage of time its BER is below a threshold of 10^{-5}. We used 10^{-5} as our threshold because a BER above that threshold is said to start to impact application performance [9]. We did this analysis separately in the uplink and downlink directions (Fig. 5b and Fig. 5c). In both traffic directions, we see that as the link layer quality increases, the bytes in flight tends to be larger. This trend is expected since higher quality links mean that the window size can grow larger and more bytes can be sent quickly. In the uplink direction, we see more variance between the best and worst bytes in flight curves as opposed to the downlink direction. This discrepancy between uplink and downlink is due to a difference in maximum allowable transmit power at the base station and client [17]. In the uplink direction, clients can not send as powerful of signals so they are more susceptible to geographic and weather factors impacting the clients' performance. We see that there are some outliers to these trends. Some clients that generally had worse BER, had larger quantities of bytes in flight. We suspect that these outliers are due to differences in the applications used by individual clients, and the respective offered load.

5 End-to-End Network Performance Evaluation

For our End-to-End Network Performance Evaluation, we investigated the impact of several factors on network performance. Our goal was to examine the resilience of an active TVWS network to multiple effects over a period of multiple years. The shortcomings or strengths of the network could provide insight into how to better deploy and maintain networks in rural areas.

5.1 Overview of Clients

From 2016 to 2019 the TVWS network had 34 clients across six base stations. Referring back to Fig. 3, at its most busy, the network had about 30 clients while at other times it had fewer than 15. Some of the clients were on the network for over two and a half years while other clients were active for a shorter duration of time. In section 6, we discuss some of the reasons that the activity of clients varied, why there was a significant drop in the number of clients in 2018, and how this drop relates to the termination of the network.

In addition to this wide range of activity among network clients, there were also differences in individual link performance. We measure link performance using BER and SNR. The BER, as explained above, is a measure of the fraction of corrupted bits among all transmitted bits over a connection. The SNR is a measure of the signal strength over the noise floor, and can be measured at both the client and base station. SNR can be impacted by distance from sender, link setup, environmental factors, and network hardware. Typically, a better SNR will lead to a better BER.

In Fig. 6 we present CDFs for the uplink and downlink BER and SNR of each client. Note that the x-axis for the SNR figures is reversed to ease viewing across figures; higher performing graphs for BER and SNR will all lie to the left. Although many clients consistently had good BER under 10^{-5}, there were also some that rarely had good performance. Those clients which rarely had a high quality link were potentially in a more challenging terrain due to obstacles in the way of their antenna. In our subsequent

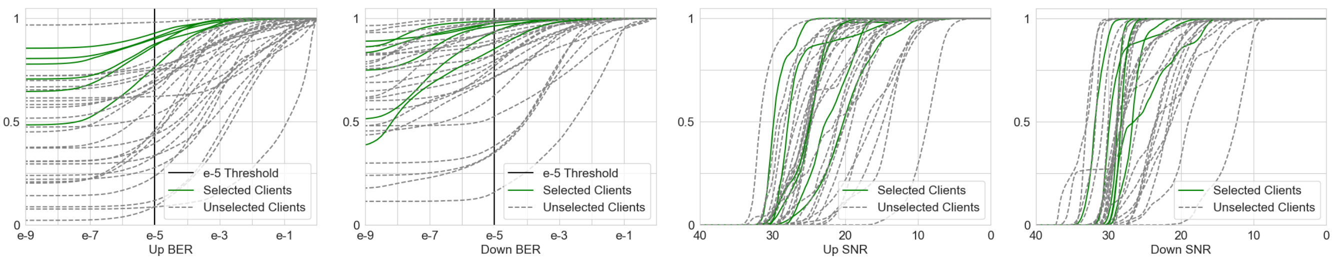

(a) Uplink BER **(b) Downlink BER** **(c) Uplink SNR** **(d) Downlink SNR**

Figure 6: CDFs for all clients' Bit-Error-Rate and Signal-to-Noise Ratio in uplink and downlink directions. Each line is the CDF for an individual client. For time of day (Section 5.4) and atmospheric effect analyses (Section 5.5), we use the subset of 6 clients in green. To avoid bias from client setup or irregular activity, we select the client subset based on their activity and the percentage of time they met a threshold for good uplink BER.

analysis of temporal and atmospheric effects on network performance (Sections 5.4 and 5.5), we only focus on clients with a well performing link. This allows us to isolate weather factors from any other performance inhibitors. The green plots in Fig. 6 show clients who were considered in that analysis.

5.2 Angle and Distance

First, we explore the effects of client location with respect to the base station on performance. As we noted previously (Section. 5.1), there was wide variation in client performance at the link layer. Although terrain and obstacles will have an impact on wireless link performance due to path loss, we do not analyze the impact of these environmental factors as we did not have client antenna height, which precluded us from computing terrain profiles. Instead, we focus on client distance and antenna orientation, as we have accurate records of these factors. Since each base station sector covers 120 degrees, we hypothesize that both the distance and angle with respect to the base station will have an effect on the client performance. We measure distance as the straight line on a map connecting base station and client. Our distance analysis includes 27 CPEs for which we had locations. We define angle as the angle between the base station's antenna direction and the direct line connecting a base station and the client. We do not know the exact angle of the CPE antennas, but we assume the CPE antenna is optimally aligned in the direction of the base station. To test our hypothesis about angles, we use the 25 client links for which we have information about both the client location and the orientation of the base station. We considered both the uplink and downlink directions for this analysis.

In Fig.7 we plot the SNR as a function of angle and distance. Each point on the graph shows the average SNR for a single client. Blue and red represent uplink and downlink respectively. The size of the dot indicates the number of data points we have for each client. We also plot the linear regression to show a general trend for the observed data. Our results show that the mean SNR tends to decrease as angle and distance increase. We see the slope of the linear regression for the average uplink SNR over distance is more negative than the slope of that for average downlink SNR. This trend is due to the broadcast power in the uplink direction being weaker than that in the downlink direction. Terrain and obstructions from the environment (the area is densely covered with trees) may explain why the decrease in average SNR was steep over only a few kilometers. We see one client at over 3 kilometers away has uplink and downlink SNR values above that of many of the other closer clients. This client may have had fewer obstructions blocking it.

Fig.8 shows the link quality of clients across distance and angle in terms of the percent of time the BER is below the threshold of 10^{-5}. Again, blue and red represent uplink and downlink and the dot size indicates the number of data points we have for each client. Our results show that links with a smaller angle with respect tot he base station's main direction tend to be more reliable. However, as the distance increases (Fig. 8b), the downlink reliability increases too. This is somewhat surprising, as typically, the performance of wireless links deteriorates with distance. However, we note that the client links within this network are relatively short – up to 3km – compared to the distances for which TVWS is designed. Thus, we conclude that although the SNR will decrease over distance, the overall consistency of a connection is not negatively impacted by distances within a few kilometers further emphasizing the efficiency of TVWS technology over wide areas.

5.3 Cross-talk

Antenna cross-talk occurs when base stations collocated at the same tower interfere with each other. For our target TVWS network, base stations 1-1, 1-2, and 1-3 shared the same antenna mount and were split into trilateral sectors. We hypothesized that this arrangement may make the three stations' clients susceptible to decreased network quality due to antenna cross-talk, even though the three sectors were set to different frequencies. More specifically, the upload performance of clients on one base station may be negatively affected by concurrent download traffic on the adjacent base stations.

To measure the presence of cross-talk, we compared the uplink BER for clients on a given base station sector as a function of the download traffic quantities on the other two collocated sectors. To isolate effects due to cross-talk from other issues that may affect

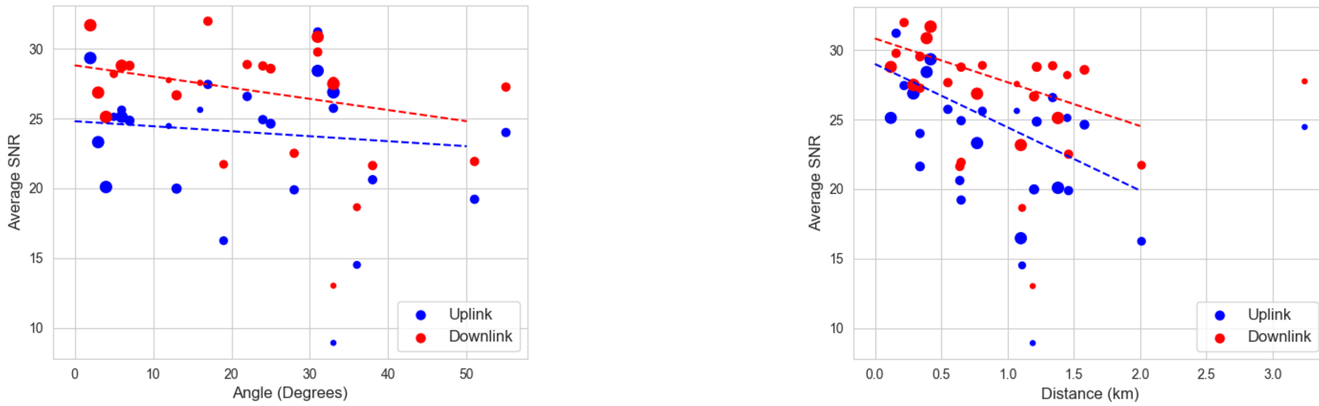

(a) Angle **(b) Distance**

Figure 7: SNR for the TVWS network for each client sorted by distance and angle. Each point shows the mean SNR for one client. The linear regression is weighted by the amount of data for each client.

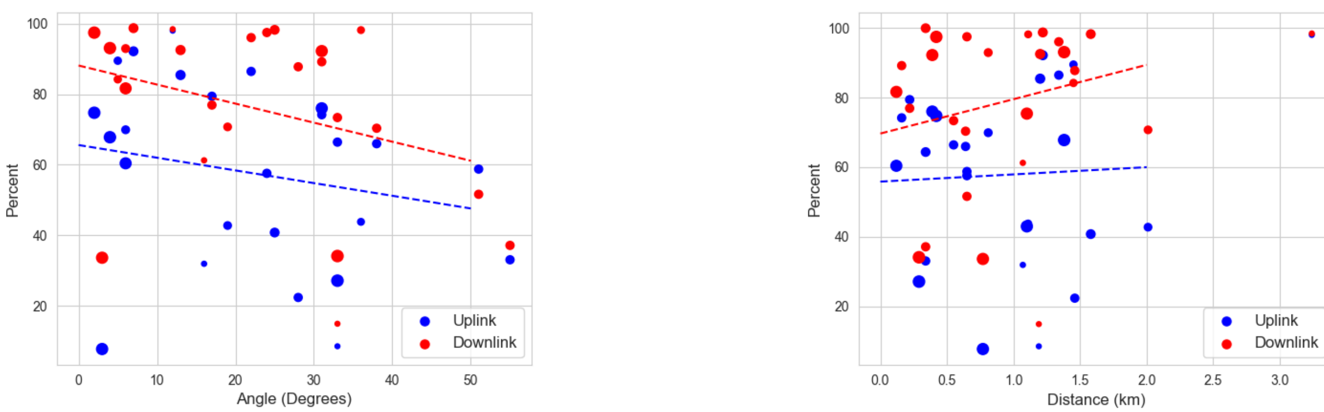

(a) Angle **(b) Distance**

Figure 8: Percent of time clients have an UP BER better than 10^{-5} sorted by distance and angle. The linear regression is weighted by amount of data for each client. This plot informs the consistency at which clients had a high performing link.

performance, we only consider well-performing links for this analysis. A well performing link is one whose uplink BER is below 10^{-5} at least half the time. We report our results in Fig. 9. The triangles are the means for each boxplot, the dark blue lines are the medians, and the box itself the range of the 25th to 75th percentile.

In Fig.9 we see that the median Up BER increases as the quantity of bytes downloaded in adjacent sectors increases. The mean is relatively stable due to the exponential nature of BER values; a few data points with BER close to 1.0 will heavily skew the mean as the remaining points are orders of magnitudes smaller. From this analysis, we concluded that antenna cross-talk negatively impacts the performance of clients on the TVWS network. For future investigation, more granular data is needed to more precisely determine the amount of alignment between uplink and downlink data.

5.4 Time of Day and Time of Year

We now move to discussing factors related to temporal and atmospheric effects. For the rest of our analysis in this section (5.4 and 5.5), we use a select subset of clients. Specifically, to ensure that our data was not biased by client setup or irregularly active clients, we only used data from clients that had regular activity and good performance. By regular activity, we mean clients that had data for all 12 months of the year. By good performance, we mean clients that most frequently met the uplink BER threshold of 10^{-5}. We also checked that we did not analyze data from days in which a client link was malfunctioning. We identified these malfunctioning links by searching for multi-day periods in which the daily median downlink BER was above 0.001, multiple orders of magnitude greater than typical for a functioning link. In addition, we checked that these periods of bad performance were occurring on different dates so as to rule out a storm that would impact the clients more

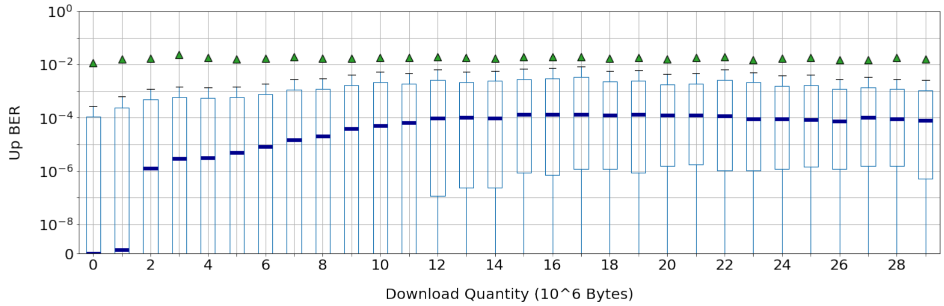

(a) Uplink BER S.1-1

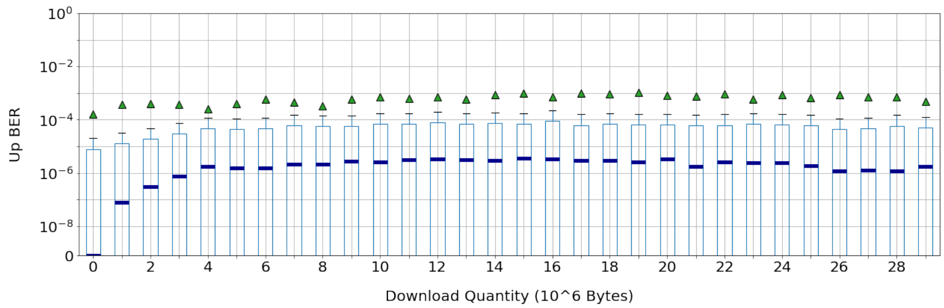

(b) Uplink BER S.1-2

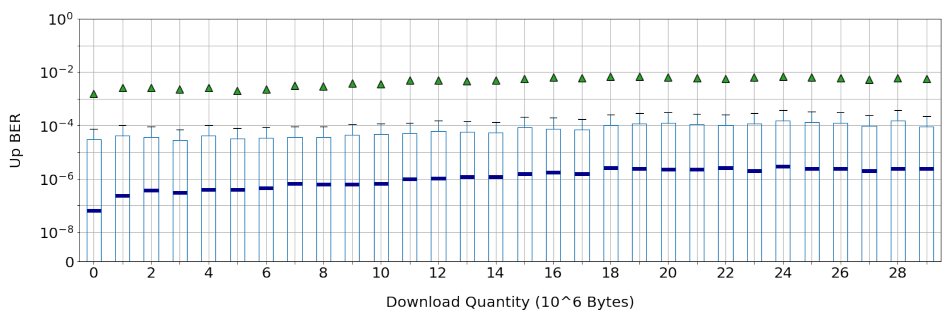

(c) Uplink BER S.1-3

Figure 9: Uplink BER as a function of download quantity on adjacent sector for base stations 1-1 (a), 1-2 (b), and 1-3 (c). Each box plot displays one station's clients' Up BER distribution for an amount of traffic downloaded on adjacent antenna sectors. The triangles are the means and the blue bars are the medians. Approximately 80 percent of the datapoints fall within the boxplots shown. The remaining datapoints are spread over a wide range of boxplots at higher download quantities, each of which have few datapoints.

uniformly. With this criteria, we selected six client representing 20.0% of our total data. For these clients, there was only one five day long period in July, 2017 in which one of the final six selected clients was malfunctioning. We removed that data from our analysis.

We begin by assessing if time of day or time of year correlates with changes in network performance. Our goal for this analysis was to measure the consistency of the network performance in terms of SNR and BER throughout the day and year. For time of day, we expected to see better performance at night as there would be less network contention. If there were a drop in performance during the day, it may indicate that the network was struggling to keep up with users' demand during the day. To test our hypothesis we compared BER and SNR performance on an hourly basis for clients of the TVWS network.

We present our results in Fig.10. Each box plot shows the BER and SNR distributions for all data that fell within that hour (in Eastern Standard Time). As expected, performance was best during night time hours with all BER values close to 0. The average BER remained relatively stable at all hours. However, the average BER is heavily impacted by outliers.

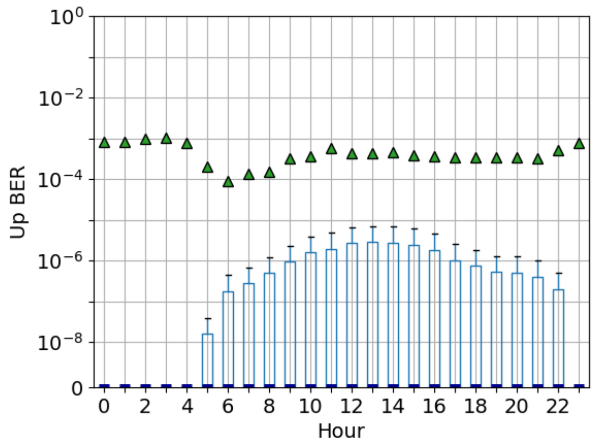

(a) Uplink BER

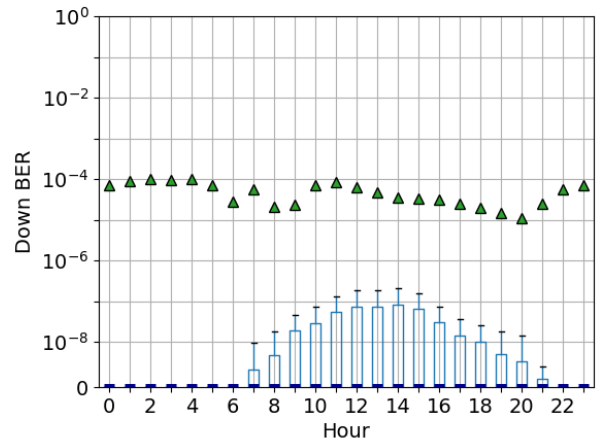

(b) Downlink BER

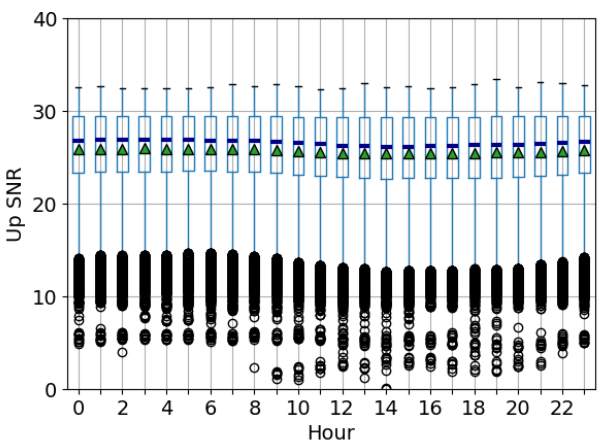

(c) Uplink SNR

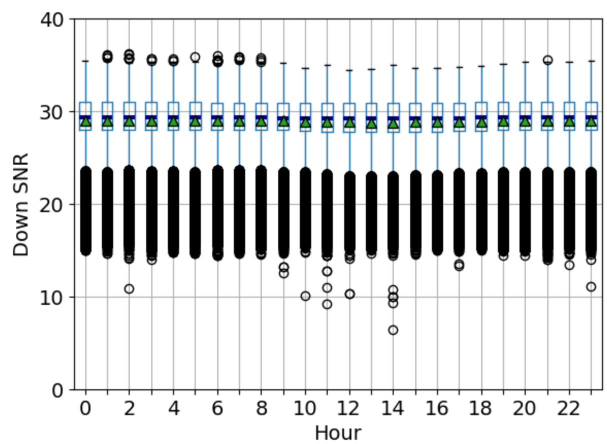

(d) Downlink SNR

Figure 10: BER and SNR for the TVWS network split by hour. Each boxplot shows the BER and SNR spread for our data collected within that hour of the day.

In contrast to BER, SNR was stable during both day and night, indicating that the physical link's performance was consistent regardless of the time of day. Thus, indeed, the performance deterioration observed during the daytime, is due to network contention rather than link deterioration. Along with network contention, this deterioration can also be attributed to cross-talk as we discuss in 5.3.

Next, we analyze seasonal effects by repeating the same experiment, but this time separating data by month of the year instead of hour of the day. We expect to see variations in performance across months of the year if seasonal changes impact the network. We present our results in Fig.11. The SNR plots makes a small arch dipping in performance closer to the summer months. We hypothesize this dip is due to thicker foliage and for weather-related reasons that we discuss in section 5.5.

In contrast, the BER improves in the summer months. Although the median BER remains stable throughout the year, the 75% marks of the box plots for spring and late autumn rise meaning that those months less consistently have high quality links. We also see that the mean BER follows an arch with the lowest values in late summer months. The contrast between the median and mean BER could be explained by heavy storms that impact the mean without affecting the median significantly. In section 5.5 we discuss correlations between atmospheric phenomena and network performance that relate to the patterns we see with the BER and SNR.

5.5 Atmospheric Effects

We last consider the impact of specific atmospheric effects. The goal of our atmospheric data analysis was to understand if climate conditions were affecting the TVWS network performance. These

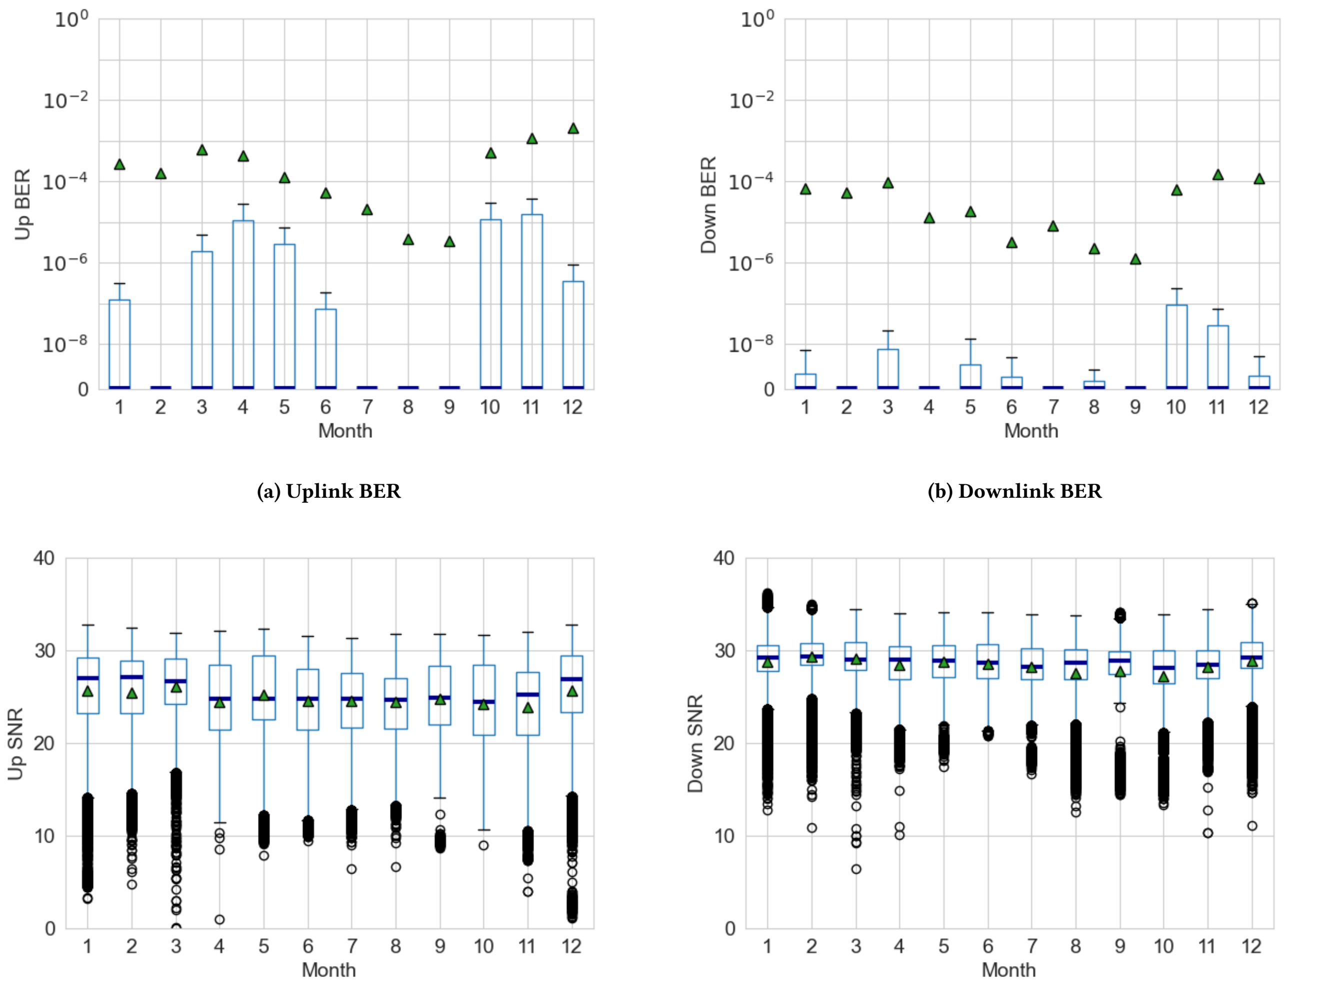

(a) Uplink BER

(b) Downlink BER

(c) Uplink SNR

(d) Downlink SNR

Figure 11: BER and SNR for the TVWS network split by month. Each boxplot shows the BER and SNR spread for our data collected in one specific month of the year.

findings can inform proactive network adaptation in response to changing climate conditions. Using weather data (see 'Weather Data' in Section 3.2) from the Mesonet, in tandem with our network data, we analyzed correlations between network performance and weather patterns. Our hypotheses, based on prior research in weather effects on UHF wave propagation [7, 29, 57], were that hotter temperatures and more humid weather would negatively impact network performance. We also hypothesized that precipitation and high winds would lead to worse performance due to obstruction of line-of-sight, vibration of the network infrastructure, and richer multipath scattering environments with swaying foliage. For our analysis of these atmospheric effects, we used the same subset of six network clients which had both good performance and high activity in order to minimize confounding factors related to setup and poor link quality.

Wind Speed. We begin with our analysis of the effects of wind speed. The Mesonet weather station records the average wind speed over every ten minute interval of the day [4], so we used this metric for our analysis. Our hypothesis was that greater wind speeds would lead to worse performance due to more movement of the tower's infrastructure and surrounding trees. For this experiment, we divided our dataset based on the National Weather Service's categorizations of wind speeds: Calm to Light Breeze, Gentle Breeze to Moderate Breeze, and Fresh Breeze or Faster. These correspond to wind speeds under 8mph, between 8mph and 19mph, and over 19mph [43], respectively. We present our results in Fig.12. We see that BER performance drops significantly while the SNR drops slightly in the windiest time periods confirming that indeed, higher winds negatively affect TVWS network performance. We expected the changes for SNR to more closely match the pattern for BER, but

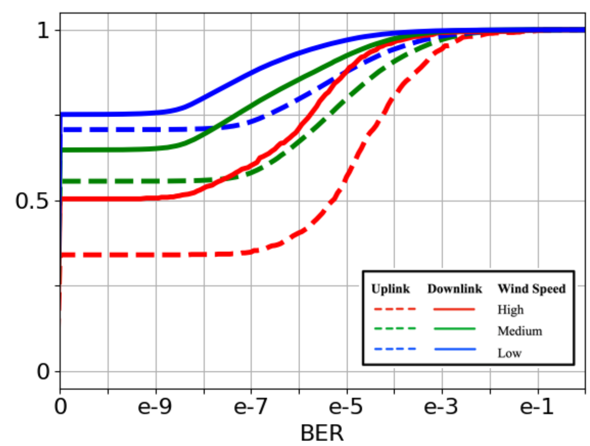

(a) BER

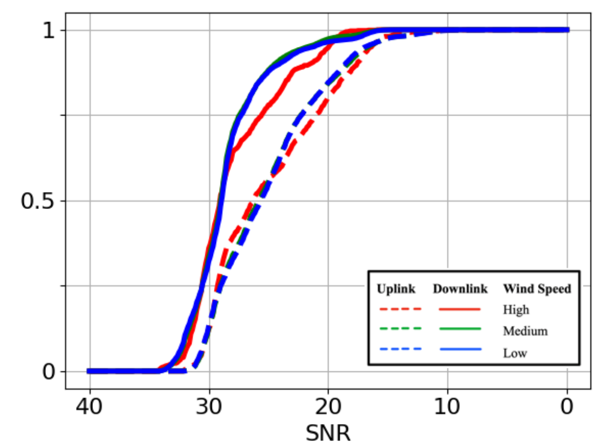

(b) SNR

Figure 12: BER and SNR CDFs for wind speed. We split the data using the National Weather Service's categorizations of wind speeds. 0.2% of data fell into the high wind speed category (Fresh Breeze or Faster), 21.7% in medium (Gentle to Moderate Breeze), and 78.1% in low (Calm to Light Breeze). For SNR, the medium windspeed CDFs are not visible as they overlap with the low wind speed CDFs.

the difference for SNR was not as noticeable as that for BER. This may be due to the lower granularity of SNR data. Additionally, we did not experience many high-wind conditions and as a result, our network data for the windiest conditions was limited.

Temperature. For our temperature analysis, we hypothesized that especially hot weather and especially cold weather would correlate with worse performance. This hypothesis is based on the equipment's ideal temperature operating range which has both an upper and lower threshold for best performance of 50 and -30 degrees Celsius [56]. Although, the temperature never reached those temperatures, we thought more extreme weather may still affect the network. To test this, we split our dataset into three sections: the hottest ten percent of data, the coldest ten percent of data, and the middle 80 percent. We present our results in Fig.13.

We found that the network performed best during the coldest weather. It performed worst during the hottest weather, but the difference between the hottest weather and mild weather is less dramatic. Prior research showed that hot weather led to worse signal propagation [7], but this research was conducted in a hot climate and did not measure the effects on a TVWS link. For the TVWS network, the temperature was rarely above 21 C which may explain the similar performance between mild and hot weather. This trend of better performance in colder weather held even when we isolated for season and time of day. In Fig. 14 we show the BER trends across multiple different temporal slices. We select for only January nights between 2am and 4am and June afternoons between 2pm and 4pm. January 2am to 4am represents one of the coldest periods during a low traffic time of day and June 2pm to 4pm is one of the hottest periods and in a higher traffic time of day. Across these different scenarios, the BER and SNR are better during the colder periods. This leads us to conclude that colder weather continues to be preferable for TVWS network performance across multiple different environments.

Precipitation. We next consider the impacts of precipitation on BER and SNR. We expected that performance would be worse during periods with high precipitation, as it obfuscates the line of sight and also moves infrastructure and trees. For our precipitation analysis, we split the data into three sets. We compared performance without precipitation to performance during light rain and during heavy rain. Heavy rain is defined as rainfall at a rate above 7.5mm/hour [37]. We considered any precipitation above freezing temperature as rain. Periods of light rain represent 5.6% of this time period and heavy rain represents 0.1%. Our results are in Fig.15. Based on this analysis, we conclude that TVWS is impacted by precipitation and especially heavy rain. Heavy rain will impact performance and may inhibit usage of the network.

Separately, in Fig.16, we consider the effect of precipitation when temperatures were below freezing. Periods of light precipitation account for 7.2% of data during periods below the freezing point. We had no data for heavy precipitation. We found minimal impact of freezing precipitation on the network performance compared to periods without any precipitation. Therefore, we conclude that the network can handle periods of regular snowfall well.

Humidity. For humidity, we expected more humid weather to lead to worse network performance. This hypothesis was based on prior research which found a negative correlation between signal propagation strength and humidity [7, 31, 42, 57]. "Per the Rural-Connect specification [56], it is safe to operate the equipment at up to 95% non-condensing humidity. Thus, we split our data to low and high-humidity using 95% as a threshold. Our data does not allow separation of condensing vs. non-condensing humidity. From these results, (Fig. 17), we see that the network performs similarly across the two humidity levels, with high humidity being slightly

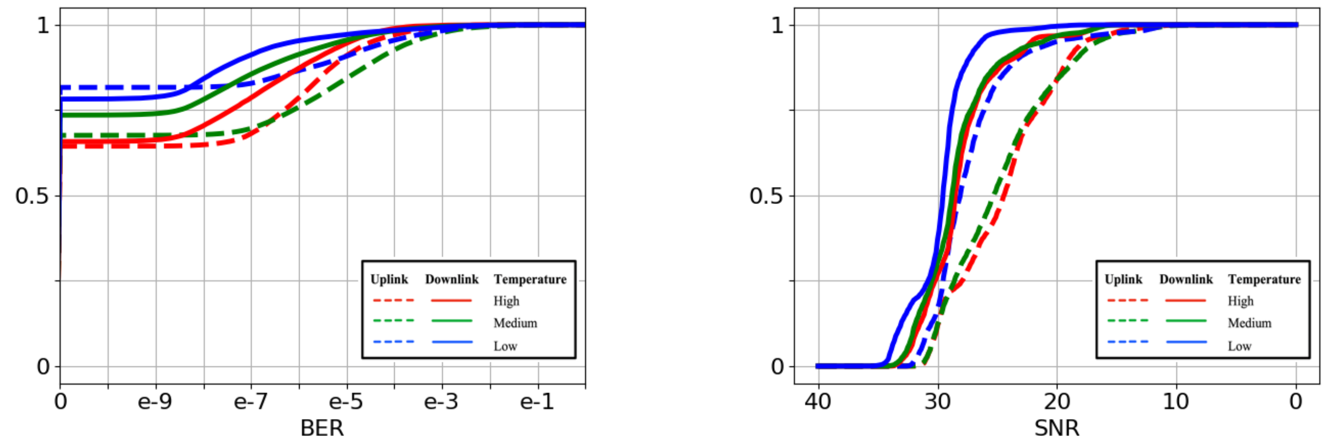

(a) BER

(b) SNR

Figure 13: BER and SNR CDFs for temperature. We split the data based into three groups, the 10% of data at the warmest time, 10% at coldest, and middle 80%.

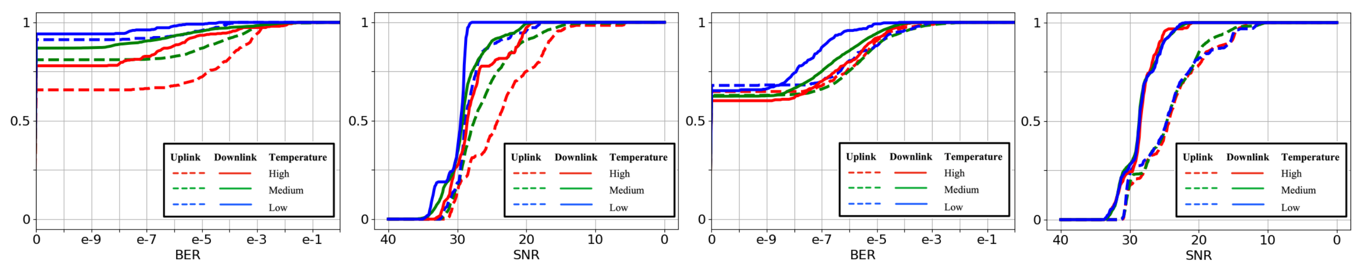

(a) BER Jan. 2am to 4am

(b) SNR Jan. 2am to 4am

(c) BER June 2pm to 4pm

(d) SNR June 2pm to 4pm

Figure 14: CDFs for clients' BER and SNR in uplink and downlink directions grouped by temperature across two different scenarios. In each scenario we split the data into the warmest 10%, coldest 10%, and middle 80%.

better than low. Thus, these results show that the level of humidity does not degrade the network performance in temperate climates.

Tropospheric Ducting. Using ducting data from DX Info Centre [21], we analyzed whether clients performed more poorly during time periods with tropospheric ducting. For this experiment we used the same subset of clients as with the weather experiments. We grouped the data from these clients together and compared the distribution of their BER and SNR when ducting effects were or were not potentially present. 95 percent of refracting effects were from June through October, so we only used data from this time period. For the time period of 2017-2018 in the network's town there was minimal ducting in general; the "tropo index" was usually at 0 and only had 11 occurrences at a value above 3 over the 21 month time period. We present our results in Fig.18. We split our data into two sets based on if the tropo index was above 0 or at 0. We found no evidence of tropospheric ducting having an effect on the performance of the network. In warm-humid locations that have more severe tropospheric ducting, we may expect ducting to still have an effect. Fig.19 is a map showing the percentage of time in which there are potential ducting effects. From it, we see that the northeast and midwest of the United States have less ducting while coastal-southern areas have more ducting.

In addition, in Fig.20 we present a measure of the average tropo index throughout December 2017 to August 2019. We see the same pattern as in Fig.19. These figures show how topology and geography effect ducting patterns. The Appalachian mountain range has less ducting effects than the warmer-more humid regions along the coast. From these figures, we can conclude that TVWS networks in colder climates or in mountainous regions will not be impacted by tropospheric ducting because those areas have minimal ducting and minimal ducting has no demonstrable effect on performance. However, TVWS networks in warmer climates might need to consider ducting to optimize performance.

In summary, we have considered the impact on link-layer performance of multiple types of atmospheric factors. We found that certain climatic conditions are better suited for TVWS than others. Lower temperatures are correlated with better performance, while high wind speeds and precipitation are correlated with worse

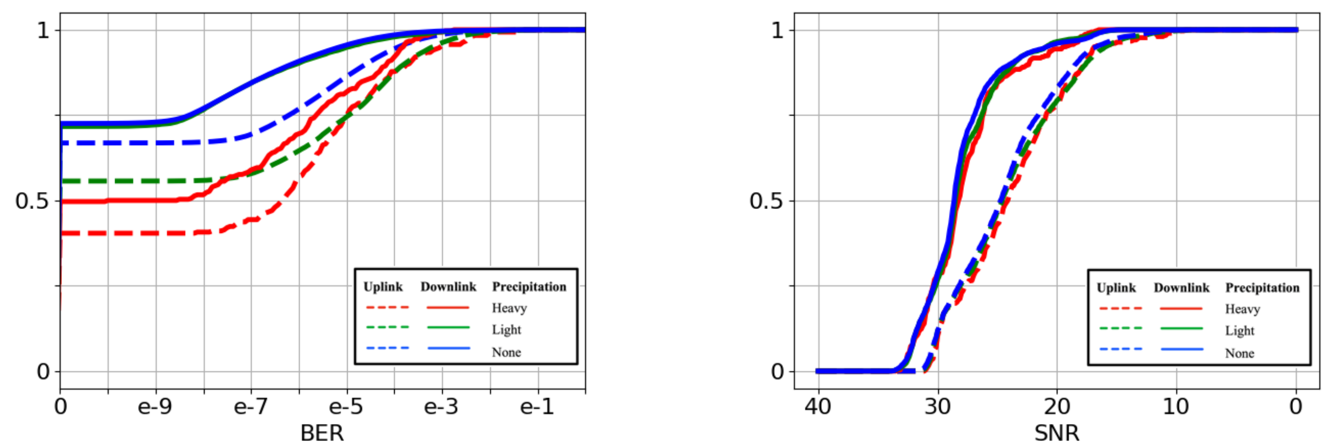

(a) BER　(b) SNR

Figure 15: BER and SNR CDFs for above-freezing precipitation rate. We compared times with no precipitation to times with light rain and heavy rain, assuming any precipitation above the freezing point was rain.

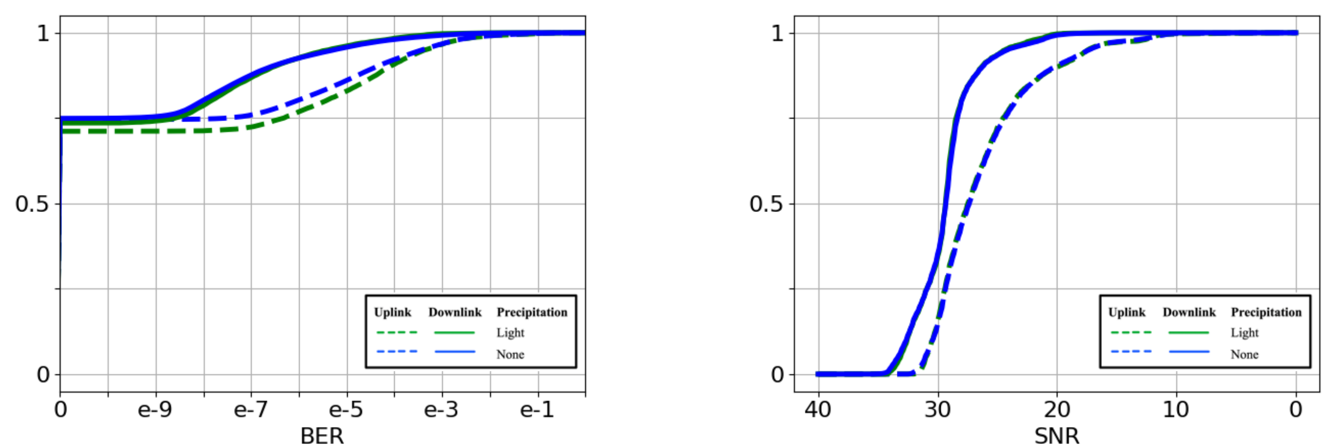

(a) BER　(b) SNR

Figure 16: BER and SNR CDFs for below-freezing precipitation rate. There is no CDF for heavy precipitation below the freezing point due to limited data. The SNR CDF for light precipitation is not visible due to it overlapping with the CDF for no precipitation.

network performance. Therefore, TVWS may perform well in cold environments with protection from wind. In contrast, warmer areas with higher wind speeds and heavier precipitation would be less ideal for TVWS.

6 Network Adoption and Termination

We now move to discussing the social factors impacting the longevity of the TVWS network. This TVWS network was online for about 5 years with about 4 years of active clients before being shut down. Based on our research, the network termination was primarily due to inconsistent management and maintenance of the network, as well as concerns about funding.

6.1 Over-Estimation of Enthusiasm to Adopt new Technology

Initially, the town's TVWS network had support from residents and the town board. At a town board meeting in January 2014, the clerk read a pile of letters from residents in support of the project. In addition, at a July 2014 board meeting, the board shared the results of a survey sent to residents. They state that over 80 percent

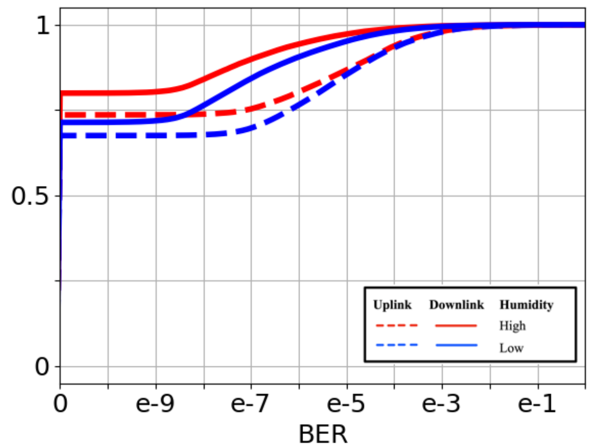

(a) BER

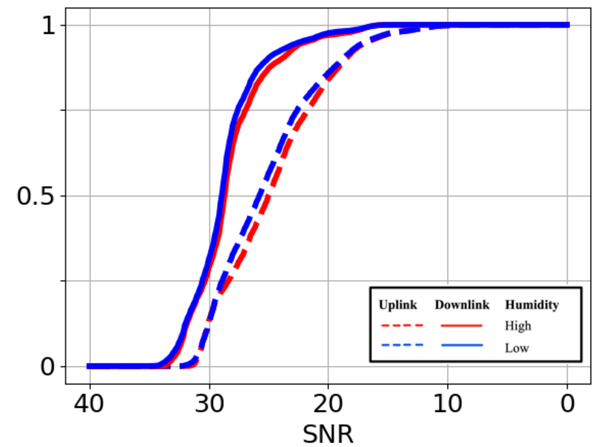

(b) SNR

Figure 17: BER and SNR CDFs for humidity. We split the data into low and high humidity using using 95% as a threshold, based on the manufacturer specification for ideal operating environment. We saw that high humidity had no major impact on the network performance.

of surveyed residents were willing to pay the installation fee and monthly charge for internet [22]. However, this initial interest did not result in residents actually signing up for the network. Although the town had hoped that over 90 households would join the network, the network at its largest only had about 30 [47].

Furthermore, during the initial construction of the network, the town ran into problems relating to the placement of equipment on people's property. These problems compounded preexisting antagonism and skepticism towards the network in the community [46]. When a new town board was elected into office, they were much more hesitant to work with the prior network engineer. Their general attitude towards the TVWS network project was less positive, and they were slow to provide support for the network. These changes led to delays in maintenance and inconsistencies in management by the town board.

6.2 Concerns about Funding

One of the primary concerns regarding the network was a lack of funding. Since fewer residents than expected joined the network, the network was making less money than initially expected. Furthermore, the town also had trouble collecting money from residents who were on the network [11, 48, 52]. On the residents' side, they expressed frustration in the billing process [50]. The concerns about funding affected the maintenance of the network as board members were wary of providing any extra money to support the network or fix problems. This led to a deteriorating network.

6.3 Difficulties Maintaining the Network Infrastructure

For the residents using the TVWS network, there were delays when they needed technical support [51]. Initially, the network had a designated operator, who had also designed and deployed the network. However, this individual was not always employed by the town. During the periods with an operator, clients were satisfied and the network moved towards being financially sustainable [52]. At the other times, there were delays and challenges with fixing equipment [51]. Without a constant operator, the town board itself was responsible for resolving problems in the network which its members were not equipped to do. This led to them paying for contract work with other companies who were not trained to operate TVWS equipment [51]. Ultimately, due to TVWS being an emerging technology and the town being in a rural area, there were few options for support for the town board to turn to.

As residents became frustrated with the TVWS network, many of them dropped the service for other options that were starting to become available [49]. By the end of 2020, there were only around 10 clients still subscribing to the TVWS network. With the grant conditions fulfilled and other internet options becoming available to residents in the town, the town stopped maintaining the network altogether [54].

7 Key Findings and Recommendations

In this section we summarize key results and findings, and make recommendations for rural deployments to maximize the utility and longevity of a network. While our insights are drawn based on a deployed TVWS network, the findings apply more broadly to community-owned fixed wireless networks.

7.1 Design Topology to Minimize the Angle between Client and Base Station Antennas

Based off our results, minimizing the antenna angle between client and serving base station within 30 degrees is more important than minimizing their distance. As the angle increases, the bit-error-rate worsens. The distance between client and base station should still

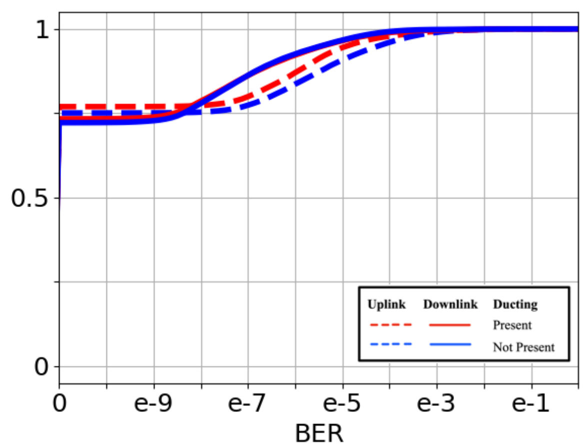

(a) BER

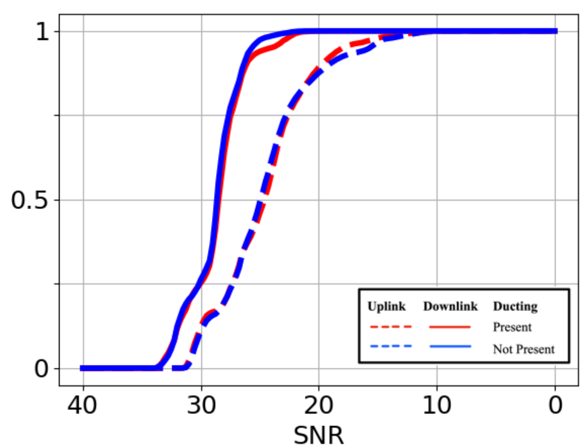

(b) SNR

Figure 18: Tropospheric Ducting CDFs. During the months of June through October, we compared periods with potential ducting effects, meaning DX Info Centre "tropo index" above 0 [21], to periods without them We saw no effect on the network performance.

be within a few kilometers, but this distance does not need to be minimized. Alternatively, network designers could consider the tradeoffs of using an omnidirectional antenna for the base station, or using larger beamwidth client equipment. These would decrease the distance covered, but better ensure a good antenna angle for clients.

7.2 Trade Off between Cross-Talk and Maintenance Efficiency

Collocating base stations saves on infrastructure investment, and prior research recommends minimizing backhaul in rural areas [32]. However, we show that cross-talk, i.e. the downlink signals on one

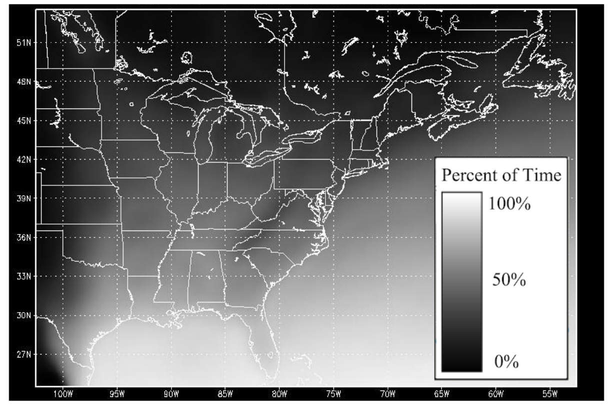

Figure 19: Percentage of time with refracting effects along the east coast of the United States. We see more ducting in warmer and more humid areas. Data comes from DX Info Centre [21].

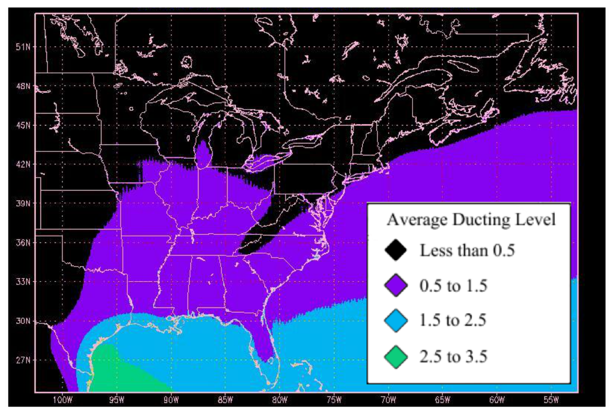

Figure 20: Average ducting intensity along the east coast of the United States. We see higher averages in warmer and more humid areas. Data comes from DX Info Centre [21].

base station negatively affecting the uplink performance on another, can impact the performance of multi-sector deployments. To avoid performance issues with collacted base stations, antennas should be carefully separated and shielded. Even with such precautions, cross-talk could still be an issue. Thus an alternative approach could be to synchronize uplink and downlink activity across collacted sectors, so no antenna is transmitting while another one is receiving.

7.3 TVWS is a Good Option for Cold Environments

Based on our analysis, we saw that there was better performance in cold temperatures and that snowfall had no effect on performance. Therefore, we conclude that TVWS networks may be well suited for colder environments. In addition, although we saw negative effects due to rain and wind, the TVWS network worked well in the temperate climate of the northeast United States. We expect that TVWS networks may face more challenges in warmer coastal areas that experience higher temperatures, wind speeds, and heavier precipitation. These effects could be compounded by potentially more tropospheric ducting in those areas as well. One caveat to this finding is that there may be additional challenges with equipment maintenance and set up in colder climates. If engineers are able to reliably maintain the network equipment, cold environments will work for TVWS networks.

7.4 Emerging Technologies Have Fewer Available Experts

Existing research has shown that network maintenance can be especially difficult in rural areas [20, 55]. Thus it is especially important for rural towns to consider whether they will have the sustained expertise and maintenance support for a technology that might be unusual and hard to maintain for a typical network engineer. It is evident that with emerging technologies in rural communities, the project's success is dependent on fewer people. However, it is important for the long-term success of projects that multiple people can fill in for any critical role. When considering the viability of a project in a rural area, we recommend taking steps from the start so that the project has no single person without whom the project would fail.

7.5 Accurate Understanding of Community Buy-In

Although there was much support for the TVWS network in this rural town, there were also residents who were skeptical of its value. This led to a deteriorating network when the new town board was elected, as they did not view the TVWS network as a priority for the town. The town board who initially brought the grant to the town and started the network construction process was not there when the network was actually supporting users which led to unpredictability and a lack of reliability. The town initially overestimated the support for the network, perhaps due to biased sampling. Care should be taken to ensure an accurate understanding of community support. Any surveys should be scientifically valid and represent the entire community who will be supporting a project.

8 Discussion and Conclusion

This paper presents a longitudinal and multi-dimensional study of the performance and adoption of a community Television White Space network in a rural mountainous town. Our analysis is based on four months of transport layer and three years of link-layer data collected from the network, data from a nearby weather station, data for regional tropospheric ducting, and the minutes posted from five years of town board meetings. There are several limitations of our data, that are worth reiterating before we conclude. First, we used in-situ user traffic from a real-world and live network. While this is useful to understand TVWS performance in the field, we had no control over factors such as client concurrency, network contention, client location, and offered load, which impact performance. We have ensured that our analysis excludes confounding factors to elucidate effects based on isolated criteria. Second, the duration of data across different data types varied: e.g., while we had nearly three years worth of data link layer metrics, we were only able to collect transport layer traces for four months during the winter. This precluded us from analyzing the effects of weather seasonality of transport layer performance. Third, for our analysis of network topology and user placement (Section 5.2), we were missing location data for several clients and terrain data, including antenna height, for all links. Terrain data, specifically, would have further elucidated TVWS applicability in mountain areas, such as that of our partner community. Fourth, our weather analysis (Section 5.5) was based on data from a weather station 10 miles (16 km) away, and was performed under the assumption that the recorded weather conditions hold throughout the region. For mountainous terrains, this may not always be the case; however, we had no way of checking what the deviation of weather conditions was between the weather station and the location of the TVWS network. Furthermore, we could not sufficiently analyze some practitioners' observations for TVWS network performance, such as poorer performance with freezing rain, as our weather data seldom indicated the presence of such conditions. Finally, our study of social factors that influenced the network evolution and termination (Section 6) was dependent on the quality of publicly available board meeting minutes and statements. We did not interview individuals associated with the network for this study.

Despite these data limitations, our analysis sheds unique insights into the performance and applicability of TVWS networks for rural broadband access. We studied the effects of external factors (such as climate conditions and network architecture) on the network and studied social factors that lead to the success or failure of community-owned communication infrastructure. With regard to network infrastructure design, we found that base station collocation, the distance and angle between base station and client antennas, all play important role in the network performance. Considering the weather impacts on a TVWS networks, we found that higher wind speed, precipitation and temperature, negatively impacted the network's usability. Thus, we concluded that TVWS is a good fit for cold and dry climates, and less so for warm-coastal areas that have heavy wind and precipitation. Finally, we found that although the network generally supported users' traffic, there were also periods of poor connectivity and difficulties maintaining and fixing equipment. Based off our analysis, we better understand the potential and limitations of TVWS and of networks, and have made multiple recommendations to improve future TVWS deployments and future community network projects.

Finally, we note the impact of our findings beyond TVWS networks. Although our conclusions about the impact of external effects are limited to TVWS, our approach and evaluation framework can be utilized for future studies of other types of networks. Furthermore, our findings and recommendations relating to managing and deploying rural community networks can be broadly applied beyond TVWS.

Acknowledgments

This work was funded under NSF awards CMMI-1831547 and CNS-1845858. We would also like to thank our community partners Fred Engelmann and Cynthia Hyde for their help in facilitating this work.

References

[1] 2015. New NY Broadband Program. https://broadband.ny.gov/new-ny-broadband-program

[2] 2024. About BroadbandUSA. https://broadbandusa.ntia.doc.gov/about

[3] 2024. Internet for All: Funding by State/Teritory. https://www.internetforall.gov/funding-recipients

[4] 2024. New York State Mesonet. http://www.nysmesonet.org/

[5] Andreas Achtzehn, Marina Petrova, and Petri Mähönen. 2011. Deployment of a cellular network in the TVWS: A case study in a challenging environment. In *Proceedings of the 3rd ACM workshop on Cognitive radio networks*. 7–12.

[6] Tuheen Ahmmed, Afsoon Alidadi, Zichao Zhang, Aizaz U Chaudhry, and Halim Yanikomeroglu. 2022. The digital divide in Canada and the role of LEO satellites in bridging the gap. *IEEE Communications Magazine* 60, 6 (2022), 24–30.

[7] Joseph Amajama and M Eshiet. 2016. Impact of weather components on (UHF) radio signal. *International Journal of Engineering Research and General Science* 4, 3 (2016), 481–494.

[8] Lahlou Mohammed Amine, El Bouzrazi Adil, and Pr. Hayar Aawatif. 2017. Toward enhancing connectivity through TVWS in outdoor rural isolated environment. In *2017 Sensors Networks Smart and Emerging Technologies (SENSET)*. IEEE, 1–4.

[9] Tarmo Anttalainen and Ville Jaaskelainen. 2014. *Introduction to Communication Networks.* Artech House, Boston, 214.

[10] Tom Arbuthnot. 2018. What are Thresholds for Good and Poor Network Packet Loss, Jitter and Round Trip Time for Unified Communications? https://tomtalks.blog/what-are-thresholds-for-good-and-poor-network-packet-loss-jitter-and-round-trip-time-for-unified-communications/

[11] Susan Baker. 2017. Budget Workshop. https://thurmanny.gov/board-minutes/

[12] Timothy X Brown, Ermanno Pietrosemoli, Marco Zennaro, Antoine Bagula, Hope Mauwa, and Sindiso M Nleya. 2014. A survey of TV white space measurements. In *International Conference on e-Infrastructure and e-Services for Developing Countries*. Springer, 164–172.

[13] European Commission. 2024. Connecting Europe Facility - CEF Digital. https://digital-strategy.ec.europa.eu/en/activities/cef-digital

[14] Babu R Dawadi, Danda B Rawat, Shashidhar R Joshi, and Martina M Keitsch. 2019. Recommendations for energy efficient SoDIP6 network deployment at the early stage rural ICT expansion of Nepal. In *2019 International Conference on Computing, Networking and Communications (ICNC)*. IEEE, 714–718.

[15] Madhukar Deshmukh, Kishor Patif, Flemming Frederiksen, Knud Skouby, and Ramjee Prasad. 2013. Wireless broadband network on TVWS for rural areas: An Indian perspective. In *2013 16th International Symposium on Wireless Personal Multimedia Communications (WPMC)*. IEEE, 1–5.

[16] Marisa Elena Duarte, Morgan Vigil-Hayes, Ellen Zegura, Elizabeth Belding, Ivone Masara, and Jennifer Case Nevarez. 2021. As a squash plant grows: social textures of sparse internet connectivity in rural and tribal communities. *ACM Transactions on Computer-Human Interaction (TOCHI)* 28, 3 (2021), 1–16.

[17] FCC. 2019. Amendment of Part 15 of the Commission's Rules for Unlicensed White Space Devices. https://www.fcc.gov/document/amendment-part-15-rules-unlicensed-white-spaces-devices

[18] FCC. 2024. Inquiry Concerning the Deployment of Advanced Telecommunications Capability to All Americans in a Reasonable and Timely Fashion. https://www.fcc.gov/document/fcc-increases-broadband-speed-benchmark-0

[19] Senka Hadzic, Amreesh Phokeer, and David Johnson. 2016. TownshipNet: A localized hybrid TVWS-WiFi and cloud services network. In *2016 IEEE International Symposium on Technology and Society (ISTAS)*. IEEE, 1–6.

[20] Shaddi Hasan, Yahel Ben-David, Colin Scott, Eric Brewer, and Scott Shenker. 2013. Enhancing rural connectivity with software defined networks. In *Proceedings of the 3rd ACM Symposium on Computing for Development.* 1–2.

[21] William Hepburn. 2024. DX Info Centre. https://www.dxinfocentre.com/

[22] Cynthia R. Hyde. 2014. Minutes of the Fiscal and Regular Meeting of the Town Board of the Town of Thurman Held July 15, 2014. https://thurmanny.gov/board-minutes/

[23] David L Johnson, Elizabeth M Belding, Kevin Almeroth, and Gertjan van Stam. 2010. Internet usage and performance analysis of a rural wireless network in Macha, Zambia. In *Proceedings of the 4th ACM Workshop on Networked Systems for Developing Regions.* 1–6.

[24] Kyu-Min Kang and Byung Jang Jeong. 2014. TV band device for TV white space field trial. In *2014 IEEE International Conference on Consumer Electronics (ICCE)*. IEEE, 450–451.

[25] Alexander Karlsson, Osama Al-Saadeh, Anton Gusarov, Renuka Venkata Ramani Challa, Sibel Tombaz, and Ki Won Sung. 2016. Energy-efficient 5G deployment in rural areas. In *2016 IEEE 12th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*. IEEE, 1–7.

[26] Mohsin Khalil, Junaid Qadir, Oluwakayode Onireti, Muhammad Ali Imran, and Shahzad Younis. 2017. Feasibility, architecture and cost considerations of using TVWS for rural Internet access in 5G. In *2017 20th Conference on Innovations in Clouds, Internet and Networks (ICIN)*. IEEE, 23–30.

[27] Mohsin Khalil, Zaheer Shamsi, Akmal Shabbir, and Abdul Samad. 2019. A comparative study of rural networking solutions for global internet access. In *2019 International Conference on Information Science and Communication Technology (ICISCT)*. IEEE, 1–5.

[28] V Krizanovic, D Zagar, and K Grgic. 2011. Techno-economic analyses of wireline and wireless broadband access networks deployment in Croatian rural areas. In *Proceedings of the 11th International Conference on Telecommunications*. IEEE, 265–272.

[29] YH Lee and YS Meng. 2007. Tropical weather effects on foliage propagation. In *The Second European Conference on Antennas and Propagation, EuCAP 2007*. IET, 1–5.

[30] Frank Loh, Fabian Poignée, Florian Wamser, Ferdinand Leidinger, and Tobias Hoßfeld. 2021. Uplink vs. downlink: Machine learning-based quality prediction for http adaptive video streaming. *Sensors* 21, 12 (2021), 4172.

[31] Jari Luomala and Ismo Hakala. 2015. Effects of temperature and humidity on radio signal strength in outdoor wireless sensor networks. In *2015 Federated Conference on Computer Science and Information Systems (FedCSIS)*. IEEE, 1247–1255.

[32] Carleen F Maitland. 2018. Now you see it, now you don't: Digital connectivity in marginalized communities. *Computer* 51, 6 (2018), 62–71.

[33] Nicoletta Metri. 2018. Piloting "White Space" to connect the underserved of rural Tanzania. https://www.internetsociety.org/blog/2018/04/piloting-use-tv-white-space-community-networks-rural-tanzania/

[34] Libor Michalek, Marek Dvorsky, Ondrej Grunt, Jan Skapa, and Roman Sebesta. 2015. Analysis of signal attenuation in UHF band. *Advances in Electrical and Electronic Engineering* 13, 4 (2015), 338–343.

[35] C Mikeka, M Thodi, JSP Mlatho, J Pinifolo, D Kondwani, L Momba, M Zennaro, and A Moret. 2014. Malawi television white spaces (TVWS) pilot network performance analysis. *Journal of Wireless Networking and Communications* 4, 1 (2014), 26–32.

[36] NTIA. 2024. BEAD: Broadband Equity Access and Deployment Program. https://broadbandusa.ntia.doc.gov/funding-programs/broadband-equity-access-and-deployment-bead-program

[37] Meteorological Service of Canada. 2013. Manual of Surface Weather Observations. https://globalcryospherewatch.org/bestpractices/docs/MANOBS_weather_2013.pdf

[38] Miquel Oliver and Sudip Majumder. 2019. Motivation for TV white space: An explorative study on Africa for achieving the rural broadband gap. (2019).

[39] I Opawoye, Nasir Faruk, Olayiwola W Bello, and AA Ayeni. 2015. Recent trends on TV white space deployments in Africa. *Nigerian Journal of Technology* 34, 3 (2015), 556–563.

[40] Rabin K Patra, Sergiu Nedevschi, Sonesh Surana, Anmol Sheth, Lakshminarayanan Subramanian, and Eric A Brewer. 2007. WiLDNet: design and implementation of high performance WiFi based long distance networks.. In *NSDI*, Vol. 1. 1.

[41] Francesca Righetti, Carlo Vallati, Daniela Comola, and Giuseppe Anastasi. 2019. Performance measurements of IEEE 802.15. 4g wireless networks. In *2019 IEEE 20th International Symposium on" A World of Wireless, Mobile and Multimedia Networks"(WoWMoM)*. IEEE, 1–6.

[42] Matthias Sander-Frigau, Tianyi Zhang, Chen-Ye Lim, Hongwei Zhang, Ahmed E Kamal, Arun K Somani, Stefan Hey, and Patrick Schnable. 2021. A measurement study of tvws wireless channels in crop farms. In *2021 IEEE 18th International Conference on Mobile Ad Hoc and Smart Systems (MASS)*. IEEE, 344–354.

[43] National Weather Service. 2024. Estimating Wind Speed. https://www.weather.gov/pqr/wind

[44] Ljiljana Simić, Marina Petrova, and Petri Mähönen. 2012. Wi-Fi, but not on steroids: Performance analysis of a Wi-Fi-like network operating in TVWS under realistic conditions. In *2012 IEEE International Conference on Communications (ICC)*. IEEE, 1533–1538.

[45] P.L. Slingsby. 1991. Modelling tropospheric ducting effects on VHF/UHF propagation. *IEEE Transactions on Broadcasting* 37, 2 (1991), 25–34. https://doi.org/10.1109/11.86959

[46] Jeanie Sprague. 2016. Minutes of the Fiscal and Regular Meeting of the Town Board of the Town of Thurman Held May 10, 2016. https://thurmanny.gov/board-minutes/

[47] Susan E. Staples. 2017. Minutes of the August 08, 2017 Regular Meeting at The Town Hall. https://thurmanny.gov/board-minutes/

[48] Susan E. Staples. 2018. Minutes of the Organizational Meeting of the Town Board of the Town of Thurman, NY Held January 04, 2018 at the Town Hall. https://thurmanny.gov/board-minutes/

[49] Susan E. Staples. 2018. Minutes of the Regular Board Meeting of the Town Board of the Town of Thurman, NY Held August 08, 2018 at the Town Hall. https://thurmanny.gov/board-minutes/

[50] Susan E. Staples. 2018. Minutes of the Regular Board Meeting of the Town Board of the Town of Thurman, NY Held March 14, 2018 at the Town Hall. https://thurmanny.gov/board-minutes/

[51] Susan E. Staples. 2018. Minutes of the Regular Board Meeting of the Town Board of the Town of Thurman, NY Held May 09, 2018 at the Town Hall. https://thurmanny.gov/board-minutes/

[52] Susan E. Staples. 2019. Minutes of the Regular Board Meeting of the Town Board of the Town of Thurman, NY Held January 09, 2019 at the Town Hall.

https://thurmanny.gov/board-minutes/
[53] Susan E. Staples. 2020. Minutes of the Board Meeting of the Town Board of the Town of Thurman, NY Held December 09, 2020 at the Town Hall. https://thurmanny.gov/board-minutes/
[54] Susan E. Staples. 2020. Minutes of the Emergency Board Meeting of the Town Board of the Town of Thurman, NY Held December 29, 2020 at the Town Hall. https://thurmanny.gov/board-minutes/
[55] Sonesh Surana, Rabin K Patra, Sergiu Nedevschi, Manuel Ramos, Lakshminarayanan Subramanian, Yahel Ben-David, and Eric A Brewer. 2008. Beyond Pilots: Keeping Rural Wireless Networks Alive.. In *NSDI*, Vol. 8. 119–132.
[56] Carlson Wireless Technologies. 2012. *RuralConnect® Generation II Brochure*.
[57] KE Ukhurebor, SA Olayinka, W Nwankwo, and C Alhasan. 2019. Evaluation of the effects of some weather variables on UHF and VHF receivers within Benin City, South-South region of Nigeria. In *Journal of Physics: Conference Series*, Vol. 1299. IOP Publishing, 012052.
[58] UNESCAP. 2020. Universal access and service fund. https://www.unescap.org/resources/ap-policy-brief-universal-access-and-service-fund
[59] International Telecommunications Union. 2020. The Last-mile Internet Connectivity Solutions Guide: Sustainable Connectivity Options for Unconnected Sites. https://www.itu.int/en/ITU-D/Technology/Pages/LMC/LMC-Home.aspx
[60] International Telecommunications Union. 2023. From electricity grid to broadband Internet: Sustainable and innovative power solutions for rural connectivity. https://www.itu.int/hub/publication/d-tnd-09-2023-01/
[61] International Telecommunications Union. 2023. Measuring digital development: Facts and Figures 2023. https://www.itu.int/en/ITU-D/Statistics/Pages/facts/default.aspx
[62] USAASA. 2024. About USAF. http://www.usaasa.org.za/usaif/index.html
[63] Jaap Van de Beek, Janne Riihijarvi, Andreas Achtzehn, and Petri Mahonen. 2011. UHF white space in Europe—a quantitative study into the potential of the 470–790 MHz band. In *2011 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*. IEEE, 1–9.
[64] Morgan Vigil, Matthew Rantanen, and Elizabeth Belding. 2015. A first look at tribal web traffic. In *Proceedings of the 24th International Conference on World Wide Web*. 1155–1165.
[65] Morgan Vigil-Hayes, Md Nazmul Hossain, Alexander K Elliott, Elizabeth M Belding, and Ellen Zegura. 2022. Lorax: Repurposing LoRa as a low data rate messaging system to extend internet boundaries. In *Proceedings of the 5th ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies*. 195–213.
[66] Xiao-Long Wu, Wei-Min Li, Fang Liu, and Hua Yu. 2012. Packet size distribution of typical Internet applications. In *2012 International Conference on Wavelet Active Media Technology and Information Processing (ICWAMTIP)*. 276–281. https://doi.org/10.1109/ICWAMTIP.2012.6413493
[67] Wenjie Zhang, Jingmin Yang, Guanglin Zhang, Liwei Yang, and Chai Kiat Yeo. 2018. TV white space and its applications in future wireless networks and communications: A survey. *IET Communications* 12, 20 (2018), 2521–2532.
[68] Yaguang Zhang, David J Love, James V Krogmeier, Christopher R Anderson, Robert W Heath, and Dennis R Buckmaster. 2021. Challenges and opportunities of future rural wireless communications. *IEEE Communications Magazine* 59, 12 (2021), 16–22.