DOCSIS® Waveforms Designed for Improved Upstream RF Leak Detection and Localization

A Technical Paper prepared for SCTE by

Richard A Primerano
Principal Engineer
Comcast
Richard_Primerano@comcast.com

Benny Lewandowski
Engineering Architect IV
Comcast
Benny_Lewandowski@comcast.com

Table of Contents

List of Figures

1. Introduction

RF signal leakage occurs in the coaxial distribution system due to cable damage, loose connections, and other issues. FCC regulations limit allowable leakage field strengths and require cable providers to routinely survey the plant for RF signal leakage. In standard-split plants, the primary spectrum of interest falls in the 108 MHz to 137 MHz VHF aeronautical band, and leaks are detected by having headend equipment generate a test signal that is easily detected by specially designed meters. In high-split and full duplex (FDX) plants, critical bands now fall in the upstream, and the cable modems must generate the probe signal. The upstream OFDMA signal is noise-like, making it difficult to detect directly. The OFDMA frame contains a known preamble pilot pattern that can be detected with a traditional matched filter approach, but since the pilot is a small portion of the overall frame, an equally small portion of the leak energy can be captured for detection.

In this paper, we present a new upstream leak detection technique built around a specially designed DOCSIS ODFMA upstream data profile (OUDP) testing burst. By sending two identical pulses in a row, with known spacing between them, we can build a *correlation-based detector* that captures the energy of the full leak pulse. As a result, we have demonstrated a sensitivity increase of greater than 20 dB over the preamble detection approach.

The output of the correlation-based detector has both magnitude (indicating the strength of the leak) and a phase (indicating the relative velocity of the detector with respect to the leak) components. We demonstrate in the paper how this Doppler velocity measurement can be used to localize leaks in space using particle filtering techniques borrowed from the radar and image processing fields.

2. Background

2.1. Using the OUDP Signal for Upstream Leak Detection

With the introduction of high-split and FDX plants, critical FCC-regulated bands that were previously in the downstream now exist in the upstream. This raises new leakage detection challenges because of the busty nature of upstream signals and the noise-like nature of OFDMA signals. The DOCSIS specification requires both DOCSIS 3.1 and DOCSIS 4.0 cable modems to support generation of the OUDP signal. The signal is used for measuring channel performance and performing sounding. In addition, the industry is settling on using this pulse to support upstream leak detection. The current technique relies on building a matched filter on a portion of the OUDP pulse that is known [3,4]. Next, we look more closely at the structure of the OUDP pulse to understand the level of achievable processing gain using this matched filter approach.

2.1.1. The Anatomy of the OUDP Pulse

The OUDP pulse passes through the same upstream transmission block as regular data transmissions. This block, shown in Figure 1, contains error control coding (LDPC encoder), scrambling, and data interleaving. This cascade of operations converts the incoming data to a noise-like signal. Unless the incoming data, scrambler state, and randomizer state are all known, we will not know what signal is being transmitted over the wire, preventing us from building a matched filter to detect it.

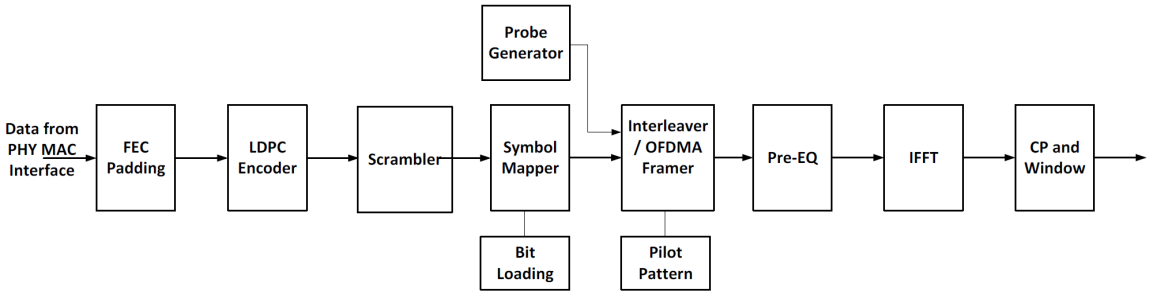

Figure 1 - DOCSIS 3.1 Cable modem upstream transmitter block diagram [2]

Note that the upstream signal path also contains a *Pilot Pattern* block. This block inserts a known data pattern into the minislot just before transmission. In normal operation, this pattern is used for upstream channel equalization. Some pilot patterns defined by the DOCSIS PHY specification [2] are shown in Figure 2. Those shown in the figure are for the 4k subcarrier OFDMA configuration. While we can't build a matched filter on the payload data subcarriers (white squares in the figure), we can build one around the pilots (red/green squares). Pilot pattern 11 has the highest density of pilot subcarriers compared to other patterns.

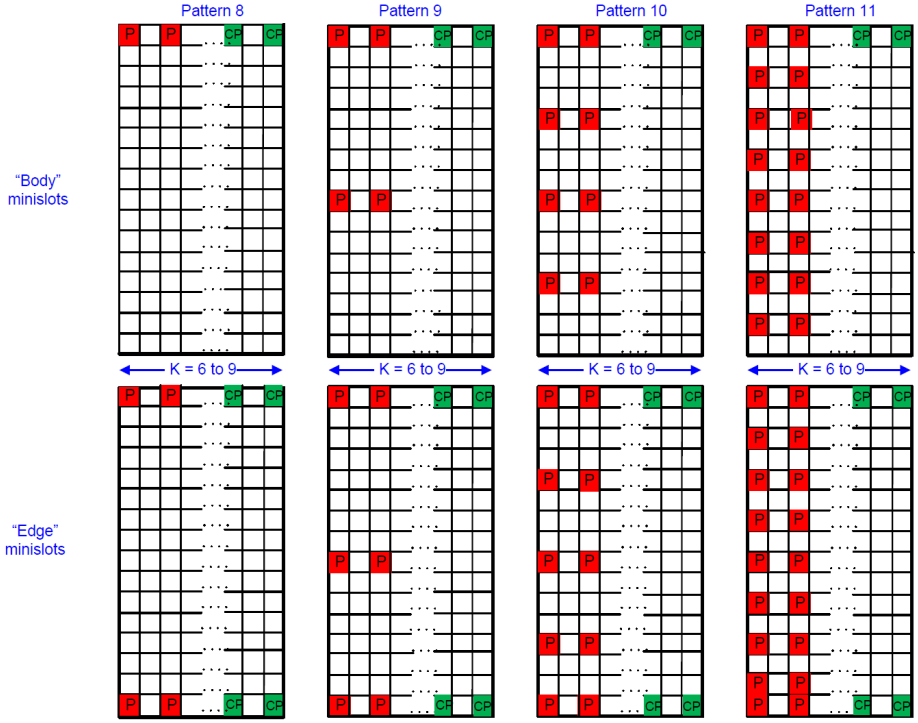

Figure 2 - Pilot pattern 11 has the highest density of pilot symbols [2]

2.2. Detection via Matched Filtering

With a deterministic portion of the OUDP signal identified, a matched filter can be built to detect it. The matched filter has an impulse response that is a conjugated version of the signal to be detected. Mathematically, the (discrete time) matched filter is given by

$$y[n] = \sum_{k=0}^{T} h[n-k]x[k]$$

where $x[k]$ is the leak signal, $h[k]$ is the matched filter impulse response, and T is the support of the matched filter. The filter output $y[n]$ is then sampled and if the signal exceeds some threshold, a leak detection is declared. The matched filter is optimal in that it maximizes the signal-to-noise ratio (SNR) of the output signal in the presence of a leak. The SNR of the matched filter can be expressed as

$$SNR = \frac{\sum_{k=0}^{T} x[k]}{kBT}$$

which is the ratio of the energy in the signal to be detected to thermal noise energy. In the denominator, k is the Boltzmann constant, B is the detector bandwidth, and T is the temperature. Note that the SNR of this filter can be increased either by increasing signal energy or by reducing detector bandwidth. In the case of DOCSIS leak detection, bandwidth reduction will not help because any reduction in bandwidth gives an equal reduction in signal energy. Furthermore, the fixed pilot pattern in the OUDP pulse places a fixed value on the numerator of the above equation.

3. The OUDP Pulse Train

The matched filter detector relies on knowing some portion of the transmitted signal. Other detectors can be built to leverage known statistical properties of the signal. Consider a signal like the one in Figure 3, representing a burst of multiple OUDP pulses, evenly spaced. The waveform is defined by three parameters: OUDP pulse width (T_p), burst repetition period (T_r), and number of transmitted bursts (n).

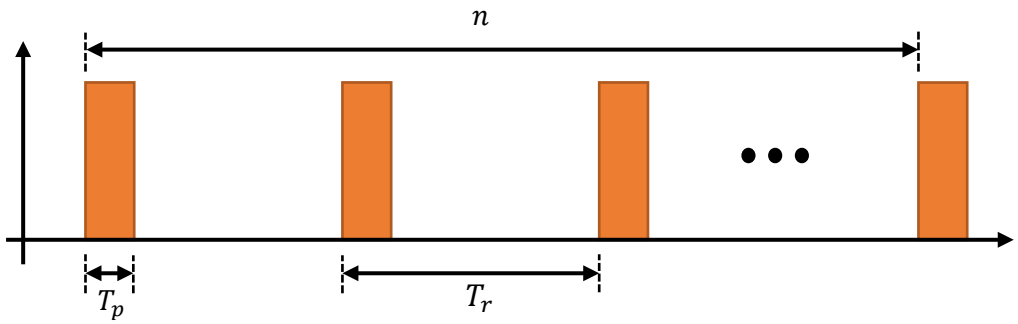

Figure 3 – A multi-pulse OUDP burst

3.1. Detection using Autocorrelation

A repeated waveform like the one in Figure 3 can be detected through *autocorrelation*. In this process, the signal is compared to a delayed copy of itself. If the spacing between pulses is known, the autocorrelation can be computed at exactly that lag, resulting in a low-complexity detector. The block diagram in Figure

4 shows the implementation of the detector. The incoming signal $x[n]$ is multiplied by a delayed and conjugated version of itself. The delay is equal to the pulse repetition interval T_r in Figure 3. The resulting signal is averaged over a period of T_p, the OUDP pulse width in Figure 3.

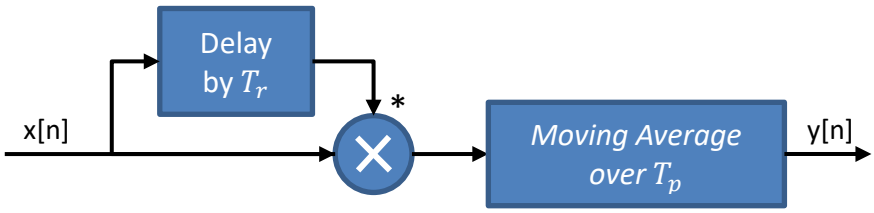

Figure 4 - An autocorrelation-based detector for detecting multi-pulse bursts

3.2. Demonstrating the Detector in a Cabled Test

To demonstrate the approach, we generated an OUDP multi-pulse burst by capturing a single OUDP burst from a cable modem and repeating it five times with known spacing. The pulse, along with its parameters, is shown in Figure 5.

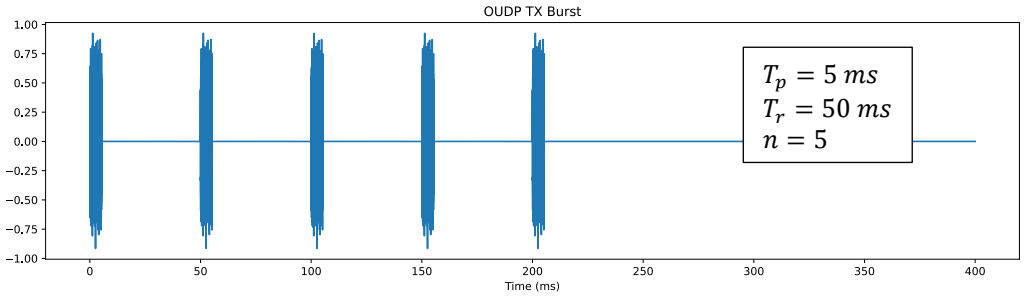

Figure 5 - A synthesized 5-pulse OUDP burst

The burst was transmitted from one software-defined radio (SDR), over cable to another SDR and recorded. Finally, the detector in Figure 4 was implemented in Python and used to process the recorded signal. The output of the detector is shown in Figure 6. Note that the received signal is not time-aligned to the transmitted one, resulting in a time shift. As expected, the detector output contains four triangular-shaped pulses for the five received pulses.

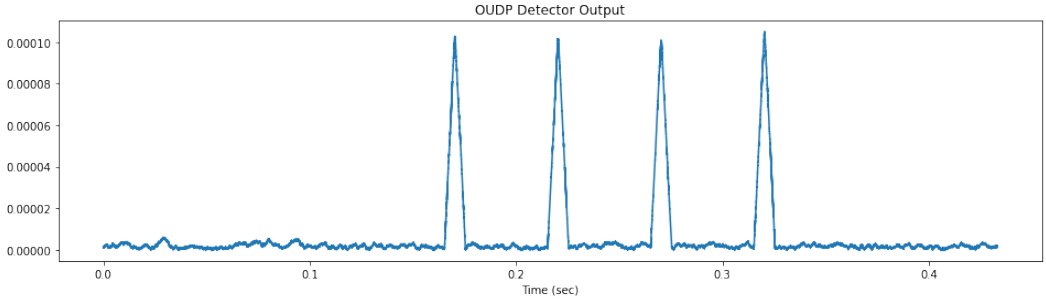

Figure 6 - Output from correlation-based detector (magnitude only)

3.3. Over-the-Air Testing

Using the same five-pulse burst shown above, we conducted indoor over-the-air (OTA) testing. The test was conducted in the 433 MHz portion of the 70-centimeter amateur radio band, but the same principles apply for VHF aeronautical (108 MHz to 137 MHz) and lower LTE (600 MHz) bands. Transmitter and receiver were 50 meters apart and the transmit power was set to a representative DOCSIS OUDP pulse.

Because the test was done in an amateur radio band, we expect that at times there will be significant interference. One weakness of the correlation-based detector is its susceptibility to correlated interference signals. Any interference that has the same correlation properties as our OUDP burst will pass through the detector. Figure 7 shows the output of the detector when to leak signal was being transmitted. This baseline capture shows the effect of correlated interference on the detector. This could have been from licensed amateur radio communications or from non-licensed Part 15 devices[1] operating in the vicinity of 433 MHz.

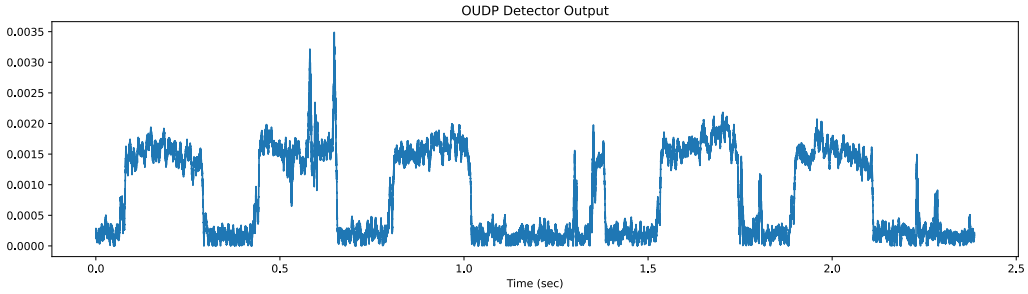

Figure 7 - Baseline capture showing strong interference

Next, we performed the same capture while the OUDP leak signal was being transmitted continuously on a loop. The capture in Figure 8 shows the detector output in this case. Four pulse groups are clearly visible among the interference energy. During quiet periods, we observe post-detector signals with an SNR in excess of 20 dB. Later in the paper we discuss techniques for rejecting interference.

[1] A "Part 15 device" is a low-power, non-licensed transmitter that complies with the regulations in Part 15 of the FCC rules. Part 15 transmitters typically use low power, often less than a milliwatt. Part 15 transmitters are non-licensed because their operators are not required to obtain a license from the FCC to use them. Some examples of Part 15 devices include automotive, garage, and consumer remote controls. Per the FCC Rules, Part 15 devices must not cause harmful interference, and must accept any interference caused by licensed OTA services. Some Part 15 devices operate in the vicinity of 433 MHz, which is in the 70-centimeter amateur radio band, a licensed over-the-air radiocommunication service.

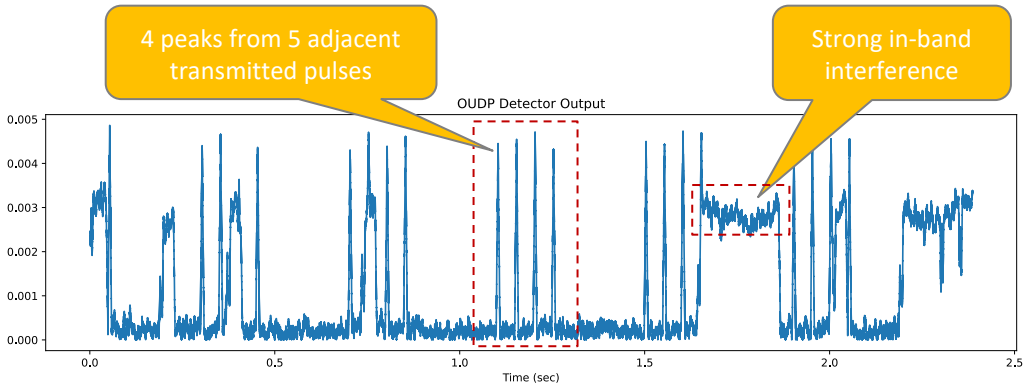

Figure 8 - Correlator output showing OUDP peaks and strong interference

3.4. Incorporating Doppler Processing for Leak Localization

The detector output in Figure 8 shows magnitude only. If the receiver is mounted in a moving vehicle, we can use the phase of the detector output to measure the relative velocity of the leak with respect to the vehicle. The relationship between velocity and phase is given by:

$$v = \frac{\theta c}{T_r f_c}$$

where θ is the phase of the detected peak (in radians), c is the speed of light, T_r is the pulse repetition interval, and f_c is the center frequency of the OUDP pulse.

If a leak is observed over a span of several minutes as a vehicle moves through an area, we will build a time-history of Doppler velocity estimates between the vehicle and leak. Furthermore, if we know the position and velocity of the vehicle for each measurement, we can apply tracking techniques to help localize the leak. This is important because detected leaks must generally be repaired. Figure 9 shows a simulation of a particle filter being used to localize a leak (gold star at x=700 m, y=250 m) using velocity measurements captured from a vehicle-mounted leak sensor (blue star starting at x=400 m, y=50 m). The arena is 1 km wide by 500 meters high and the four frames shown below would represent a collection spanning several minutes.

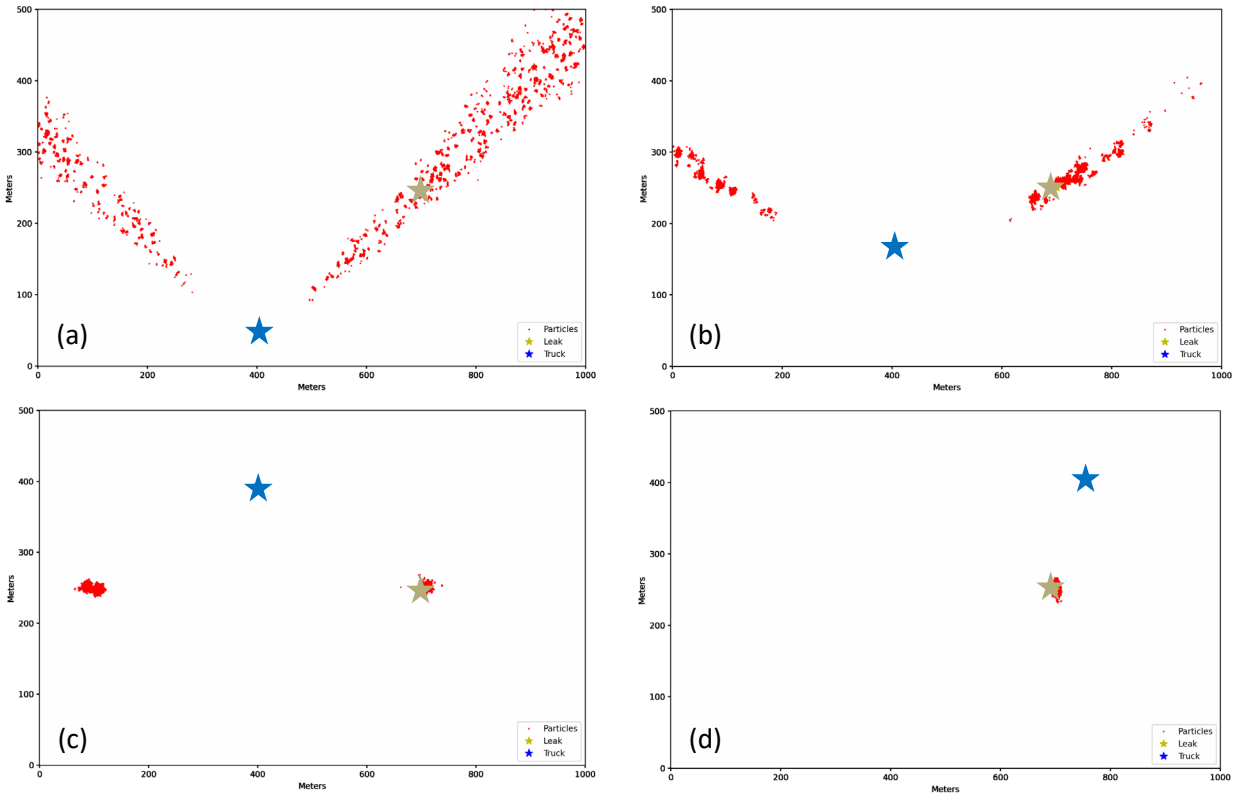

Figure 9 - A particle filter converging on the true location of a leak over time

The red particles represent assumed locations for the leak. The particle filter begins by randomly distributing them in the arena. As the vehicle moves forward (a), it measures a velocity relative to the leak indicating the leak is in front of the vehicle at a bearing of approximately 45 degrees. As the vehicle continues to move (b), it collects additional measurements indicating that the leak is almost stationary relative to the vehicle, i.e., the leak is almost at a right angle to the vehicle's motion. Further in time (c), the vehicle observes a change in sign of the velocity, indicating the leak is moving away from the vehicle. The set of measurements taken between (a) and (c) are all in a single direction. This is not enough data for the particle filter to determine if the leak is to the left or right. As the vehicle changes bearing, the velocity relative to the leak shows the leak is moving toward the vehicle. This immediately allows the particle filter to determine which of the candidate locations from (c) is the true leak location. Although we have described the process above as the tracker making decisions about leak location at discrete points in time, the actual operation happens continually. The particle filter is solving a maximum-likelihood problem iteratively such that the particles coalesce around the point(s) that best agree with the velocity measurements the filter has seen over time.

4. Choosing OUDP Pulse Train Parameters

The three pulse-burst parameters in Figure 3 can be freely adjusted by the designer to optimize system performance. Some considerations when selecting pulse parameters are discussed next.

Pulse Width (T_p)

Pulse width is the main parameter that sets detection sensitivity. The relationship between pulse width (T_p), signal bandwidth (BW), and processing gain (PG) is approximately $PG = 10 \log (T_p * BW)$. Using typical values of $T_p = 5\ ms$ and $BW = 1.6\ MHz$, we get $PG = 10 \log(8000) = 39\ dB$.

Pulse Repetition Interval (T_r)

When both the transmitter and receiver are stationary, the phase of the autocorrelator's output peak is zero. If there is relative motion between the two, the phase will be non-zero. The phase of the peak encodes the distance (in wavelengths) the receiver has traveled between pulse. The value of T_r sets the maximum unambiguous velocity (v_{max}) that can be measured, $v_{max} = \lambda/2T_r$. The receiver can travel no more than one half of a wavelength between pulses for unambiguous measurement. Using typical values of $T_r = 50\ ms$ and $\lambda = 2.3\ meters$ ($f = 130\ MHz$), we get $v_{max} = 23\ m/s = 51\ mph$.

Number of Pulses (n)

The minimum number of pulses that can be transmitted is two, resulting in a single peak at the detector output. Transmitting *n* pulses results in $n - 1$ detection events. If the relative motion between transmitter and receiver is constant over the duration of these pulses, the detection events can be coherently averaged to form a single effective detection. In this case, the overall processing gain becomes $PG = 10 \log ((n - 1) * T_p * BW)$.

4.1. Choosing Parameters based on Environment

The autocorrelation-based detection is *blind*. It requires no knowledge of the transmitted signal and is sensitive only to the repetition within the signal. One drawback of this approach is that any signal exhibiting repetition of length T_r will pass through the detection and appear as interference. We may tune the three OUDP leakage burst parameters based on background interference characteristics to minimize their impact on detector performance.

Waveform parameter selection can also be used to tune system performance based on the environment. Examples include:

- **A dense urban environment:** *many cable modems in RF range, close to the receiver's path of travel*
  - A narrow pulse width is used. This results in low sensitivity, but a larger number of devices can be scheduled in a fixed time interval.
  - Since homes are near the road, RF path loss is low and the sensitivity loss is acceptable.
- **A sparse rural environment:** *few cable modems in RF range, far to the receiver's path of travel*
  - A wide pulse width is used. This results in high sensitivity, but fewer devices can be scheduled in a fixed time interval.
  - Since homes are far from the road, the high sensitivity pulse overcomes RF path loss.

5. Conclusion

In this paper, we have presented a new leak detection approach built on the same OUDP waveforms currently defined in the DOCSIS specification. By sending multiple pulses at known spacing, we can employ an autocorrelation-based detector that has high processing gain to detect weak signals. This approach has a drawback: specific forms of interference have the potential to adversely affect

performance. By integrating Doppler processing and adjusting pulse parameters based on environment, we expect some of these issues can be overcome.

Abbreviations

BW	Signal Bandwidth
dB	decibel
DOCSIS	Data-Over-Cable Service Interface Specifications
FCC	Federal Communications Commission
km	kilometer
m	meter
v_{max}	maximum unambiguous velocity
MHz	megahertz
Ms	millisecond
N	number of pulses
OFDMA	orthogonal frequency-division multiple access
OTA	over-the-air
OUDP	OFDMA upstream data profile
PG	processing gain
RF	radio frequency
SDR	software-defined radio
SNR	signal-to-noise ratio
T_p	pulse width
T_r	pulse repetition interval

Bibliography & References

[1] *Data-Over-Cable Service Interface Specifications DOCSIS® 3.1, MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I25-230419*; Cable Television Laboratories, Inc.
[2] *Data-Over-Cable Service Interface Specifications DOCSIS® 3.1, Physical Layer Specification, CM-SP-PHYv3.1-I20-230419*; Cable Television Laboratories, Inc.
[3] *Leakage Detection in a High-Split World, Industry Progress Toward a Viable Solution*; R. Coldren et al., SCTE Fall Technical Forum, October 2021
[4] *Leakage Detection in Full Duplex DOCSIS*; B. Lewandowski et al., SCTE Fall Technical Forum, September 2022

