An ADTRAN White Paper

Mitigating RFoG-OBI in DOCSIS 3.0/3.1 Networks

A statistically significant analysis of ADTRAN *LambdaShift*™

Executive Summary

Optical Beat Interference (OBI) can occur when multipleRF over Glass (RFoG) Optical Network Units (ONUs) transmit simultaneously on a shared fiber network. The phenomenon depends on multiple factors, which makes it extremely difficult to troubleshoot in the field as well as test in a laboratory. ADTRAN's patented *LambdaShift*™ OBI mitigation eliminates the chronic OBI conditions experienced in the field, while reducing the overall occurrence of OBI and the impact of unit-to-unit variations. In this white paper, we discuss why OBI occurs with RFoG, how our *LambdaShift*™ technology mitigates OBI, and analyze the impact that the DOCSIS 3.1 Orthogonal Frequency Division Multiple Access (OFDMA) upstream protocol has on upstream performance with RFoG. In particular, we focus on describing testing that is statistically significant and defining traffic that is representative of actual network traffic. We also provide test results that show the impact of OBI is minimal when *LambdaShift*™ RFoG ONUs are employed.

Key Takeaways

- **OBI Effect on DOCSIS/RFoG Networks**
  We explain OBI as a statistical phenomenon that is dependent on multiple factors. Those who test RFoG must be aware of these factors in order to design and conduct statistically significant tests when OBI is involved.

- **OBI Mitigation: How *LambdaShift*™ Works**
  ADTRAN La*mbdaShift*™ technology effectively mitigates most OBI problems by controlling and randomizing the transmit wavelength on each RFoG ONU.

- **Traffic Patterns and DOCSIS 3.1**
  DOCSIS 3.1 has more potential than DOCSIS 3.0 for concurrent upstream transmissions, and this increases the opportunities for OBI events; however, with representative traffic, most transmission slots go unused because there isn't any traffic to send. We model the actual applications that subscribers use—a Virtual Subscriber methodology—and extend that approach to define traffic levels anticipated 10 years from now.

- **Test Results with a DOCSIS 3.1 Cable Modem Termination System (CMTS)**
  Testing was performed using significant traffic as defined using the Virtual Subscriber methodology. Performance was measured by recording the number of uncorrectable codewords received at the CMTS. The results proved that mitigation of OBI was quite effective; the resulting impact of OBI was quite small (1-1.5 percent uncorrectable codewords for most test cases). Additional tests measured upstream traffic rates of 35-50 Mbps with similar results.

Testing RFoG and OBI

RFoG and OBI

Point-to-multipoint access networks accomplish multiple access by scheduling transmissions that are separated in time, frequency, wavelength, or by correlation code. Most native fiber PON systems employ Time Division Multiple Access (TDMA), while Data Over Cable Service Interface Specification (DOCSIS) typically employs both time and Frequency Division Multiple Access (FDMA) by scheduling (or provisioning) transmissions across different channels in DOCSIS 3.0 and across groups of Orthogonal Frequency Division Multiplexing (OFDM) carriers (channels and mini-slots) in DOCSIS 3.1.[1] When using DOCSIS over RFoG, FDMA can generate simultaneous transmissions from multiple RFoG ONUs (R-ONUs, a.k.a. micro-nodes). While these FDMA transmissions are separated in frequency and thus orthogonal to the RF signal, with RFoG, the RF signal is used to modulate the intensity of a laser and the resulting optical signal is recovered using a direct detection receiver. Since the laser wavelength is not precisely controlled in this process, it is possible that two R-ONUs can transmit simultaneously using lasers that are close enough together that the (non-linear) direct detection process generates optical intermodulation products that are within the optical bandwidth of the receiver. This phenomenon is known as OBI.

OBI is only generated when two factors are present:

- 1) two or more R-ONUs are simultaneously transmitting
- 2) the wavelengths of those transmissions are sufficiently close together that the resulting intermodulation products are within the bandwidth of the receiving photodiode

The combination of these two factors can make the presence of OBI appear random since the actual transmit wavelength (and line width) of a given R-ONU is typically not observable (through normal performance statistics collection) and varies with environmental conditions.[2]

The wavelength of a particular R-ONU transmission will vary within the allowed wavelength range due to unit-to-unit variations, environmental temperature, and short-term heating of the laser during transmission times. The laser will also exhibit a linewidth, which is influenced by the Lorentzian noise of the laser, the chirp of the laser and the modulator signal bandwidth. Several studies have modeled and measured OBI with directly modulated Distributed Feedback (DFB) lasers, indicating that wavelength separation from 0.1 to 0.8 nm is sufficient to avoid OBI, depending on how well the temperature of the laser is controlled.

1 *OBI also can result from simultaneous transmission of Out of Band (OOB) and DOCSIS transmissions from different R-ONUs on the same ODN. However, since the amount of OOB transmissions is small relative to DOCSIS transmissions, the presence of OBI is more likely due to multiple simultaneous DOCSIS transmissions.*

2 *While OBI can occur during normal operation, it can also appear as an artifact in fault conditions, such as when a defective R-ONU fails to completely turn off its laser.*

OBI mitigation techniques have been developed to reduce or eliminate the occurrence of OBI in RFoG deployments.

These techniques can be classified into two general categories:

- 1) Prevent wavelengths from overlapping in the receiver
- 2) Prevent transmission times from overlapping

Techniques for preventing the wavelengths from overlapping in the receiver include separating the wavelengths, either with static or dynamic wavelength assignment, and using one optical receiver per R-ONU. Overlapping transmission times are avoided through the use of scheduling at the Cable Modem Termination System (CMTS), either statically, where no R-ONUs are allowed to transmit simultaneously, or dynamically, where the R-ONUs pairs that cause OBI are learned so that the scheduler can avoid overlapped transmissions of just those R-ONUs which would cause OBI.

OBI's Impact on DOCSIS Performance

OBI causes intermittent noise in the received upstream signal when two R-ONUs transmit at the same time on overlapping wavelengths. With OBI, the noise in the receiver increases as the phase noise of the lasers is converted to intensity noise through the (non-linear) direct detection process. The amount of noise is determined by the optical spectra of the modulated lasers, the amount of overlap in their optical spectra and the relative states of polarization of the received signals. With a conventional R-ONU, the transmit wavelength is not carefully controlled and will vary due to unit-to-unit variations, ambient temperature changes and short-term heating of the laser during a burst transmission. Because the laser is not temperature controlled, the transmit wavelength during a single burst transmission increases as the laser heats up during the burst. This is observed as spectral broadening and drift on an Optical Spectrum Analyzer (OSA) and can result in an optical spectrum nearly 0.8 nm wide (-40 dB point) when averaged over the observation time of a typical OSA.

As mentioned earlier, OBI is observed as a random phenomenon, the statistics of which depend on several factors, including the relative state of polarization of the received signals, nominal wavelengths of the lasers, ambient temperature of the R-ONU and the transmit data traffic patterns (which affect overlapping transmissions and burst size). The nominal wavelengths of the lasers vary unit to unit, based on production variations. While OBI can come and go based on ambient temperature variation and varying traffic patterns, once a set of R-ONUs is deployed, the set of nominal wavelengths is fixed, which can make OBI more likely in some Passive Optical Networks (PONs) than others. Furthermore, if the nominal wavelengths and ambient temperatures of two R-ONUs on the same PON are similar, the probability of experiencing OBI can increase dramatically and occur over long time periods, creating a chronic condition.

The dependence of OBI on the nominal wavelengths of the selected R-ONUs also makes it difficult to reproduce field issues in the lab, where typically only a single PON's worth of R-ONUs are tested. It is therefore difficult to know whether a lab test is representative of what will be seen in the field or a sampling near one or the other extreme. Only with careful measurement of the wavelength of each R-ONU and comparing the statistics of the units under test to the broader population can confidence in laboratory test results be gained.

ADTRAN RFoG ONU with *LambdaShift*™ OBI Mitigation

ADTRAN's OBI-mitigated RFoG ONUs employ *LambdaShift*™, a patented method of autonomous dynamic wavelength assignment. Through temperature control of the laser, we eliminate the in-burst spectral broadening and a drift is eliminated, replacing it with an intentional randomization of the wavelength between transmissions. By using random wavelength assignment, no upstream controller or aggregator (and associated cabinet and power source) is required to mitigate OBI. Our approach to random wavelength assignment has two elements: 1) at power-on and after periods of inactivity, each R-ONU picks a random wavelength over a range of several nanometers. This ensures that there will be some statistical variation in the transmit wavelengths for the R-ONUs deployed on a single PON, even if all of the R-ONU lasers have the same nominal wavelength (which could occur if they are from the same production batch.) 2) After every transmission burst, each R-ONU's wavelength is varied over a small range. This step eliminates the chronic condition where two R-ONUs are unlucky to have nearly identical wavelengths and therefore chronic OBI problems.

The laser temperature control used in *LambdaShift*™ also minimizes the spectral broadening and shift seen in conventional R-ONUs as the laser is heated and cooled by the data transmission pattern (as well as the slower wavelength drift due to changes in the ambient temperature.) By narrowing the spectrum, we reduce the likelihood of an OBI event even if the total span of the wavelengths on a given PON isn't any more than that attained by conventional R-ONUs. This is illustrated in *Figure 1*, where we see that the optical spectrum of the *LambdaShift*™ R-ONU is substantially narrower (less than 20 percent of the spectral width at -30 dB) than a conventional R-ONU.

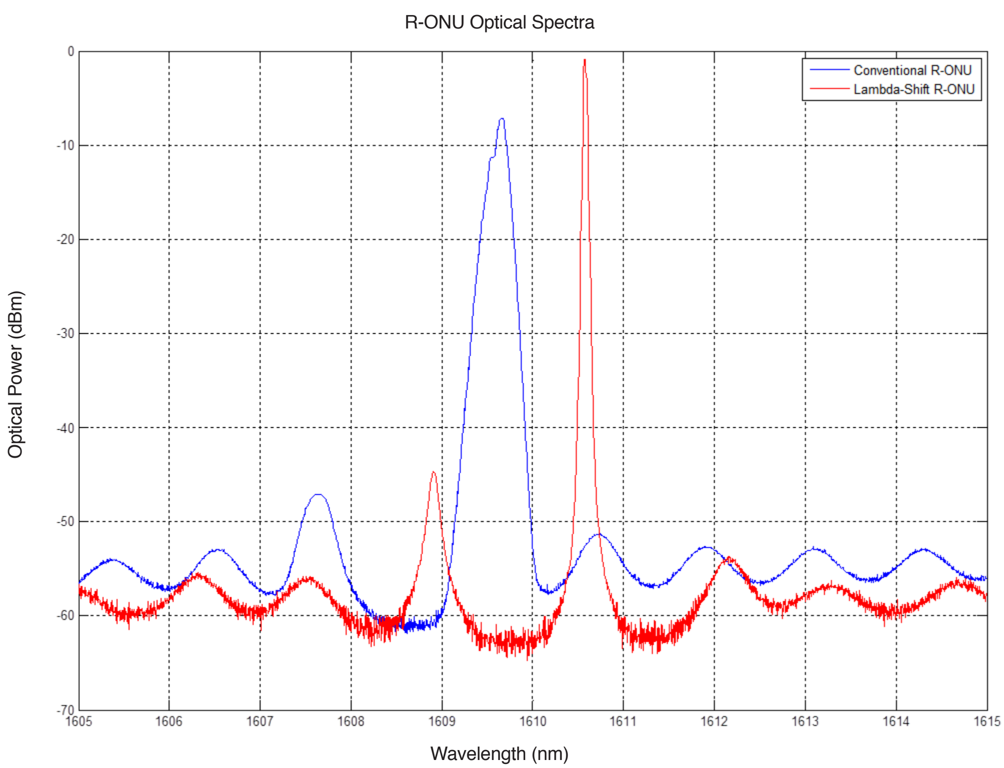

Figure 1: Optical Spectrum of Conventional and *LambdaShift™* R-ONUs

While *LambdaShift™* increases the randomness and spread of R-ONU transmit wavelengths, it doesn't completely eliminate the influence of unit-to-unit variation on test results. Typical deployments place 32 R-ONUs per PON. *Figure 2* shows the probability that all 32 *LambdaShift™* R-ONU transmit wavelengths will be within X nm for a given PON. This distribution includes the effects of both the unit-to-unit variations and the power-on random wavelength assignment. From the figure it can be seen than for virtually all cases, the wavelength spread will be between 5 and 10 nm. Over that range, the probability of wavelength overlap can vary by a factor of two, leading to variability in the test results. By observing the maximum wavelength spread for a given test, one can get a feel for whether the test is on the pessimistic (narrow wavelength spread) or optimistic (wide wavelength spread) side of the distribution. The wavelength spread will change after every power cycle of the R-ONUs, so by performing multiple test runs separated by a power on-off cycle, more statistically significant results can be collected.

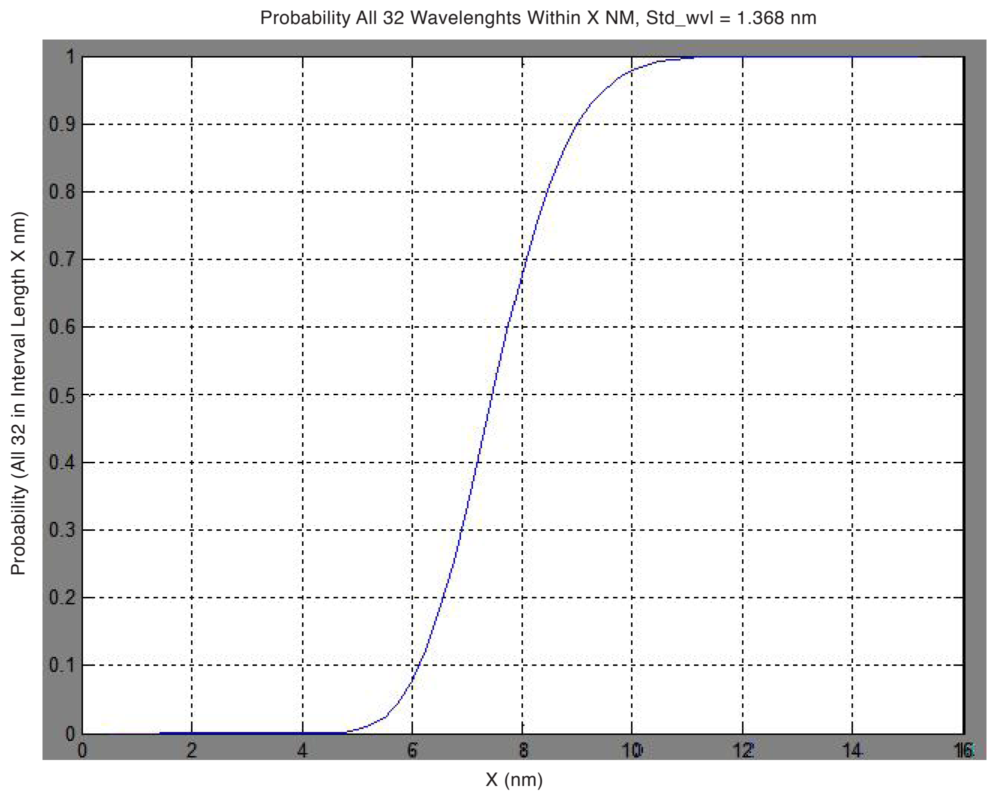

Figure 2: Probability That All Wavelengths From 32 *LambdaShift™* R-ONUs Are Within X nm

Traffic Patterns for Testing

While OBI does cause some packet loss, even with mitigation, it is important that any tests on OBI impact be done with as realistic traffic patterns as possible to gauge the performance impact that will be seen in deployment. This is often a challenge because test equipment is optimized for repeatable testing, not emulating real-world traffic. For instance, if OBI performance is tested with a traffic generator that sends out User Datagram Protocol (UDP) packets simultaneously on all ports, the Cable Modems (CMs) will all ask for transmission grants at the same instance and will be scheduled to transmit as simultaneously as possible. The result is a pathological worst-case test of OBI performance but is not reflective of how the units will operate in a deployment environment.

To have laboratory testing better reflect real-world deployment, ADTRAN has developed a simulation and testing environment built around the concept of a virtual subscriber. The companion white paper "Virtual Subscriber: Realistic Traffic Generation in the Lab Environment,"[3] describes the basic approach. We will note the highlights here.

The Virtual Subscriber is a set of application characteristics and subscriber behaviors implemented in software that, when taken together, generate traffic emulating the behavior of an actual subscriber as closely as possible. Characteristics and behaviors are defined for each of a number of application classes: currently, we are modeling Voice over IP (VoIP), real-time interactive video, Over-the-Top (OTT) video streaming, upstream video streaming, file sharing, non-video Web access, virtual reality, gaming and performance testing. Application traffic for all of the above classes is modeled in sessions, where each session represents a new instantiation of traffic (e.g., a video clip or program). Session initiation is modeled as a Poisson process with a defined Mean Time Between Sessions (MTBS) based on the application class and the overall subscriber load. The characteristics that define a session—duration, number of flows, flow rates, transport and application protocols, packet size, etc.—are also dependent on the application class and subscriber load. Since residential subscribers represent households containing multiple people and devices, it is common for multiple sessions to overlap in time.

Virtual Subscriber testing is useful for OBI testing as it generates a realistic distribution of transmissions over time and employs real Transmission Control Protocol (TCP) stacks to generate application packets and their acknowledgments. This results in a traffic pattern which may have a variable number of subscriber ports used during a given test, but even more variation when observed over shorter-term intervals. For instance, an OTT video flow generates packets in every one-second interval, but only in about seven percent of 10 ms intervals. As we get down to the timescales of DOCSIS transmission grants, the number of users trying to transmit in a given frame decreases further, which lowers the probability of OBI, even when there is wavelength overlap.

When applying the Virtual Subscriber methodology to a test environment, there are few choices that must be made. The first, is: What traffic load should be modeled? To answer that question, we must look at historical network speeds and loading information. Based on

3 *"Virtual Subscriber: Realistic Traffic Generation in the Lab Environment" whitepaper - Dr. Kevin W. Schneider, Ph.D, ADTRAN*

data published by various sources, the average peak load on fixed internet access is approximately 1 Mbps per subscriber downstream and approximately 50 kbps per subscriber upstream. Projecting that forward at a 25 percent compound annual growth rate (CAGR) (higher than reported over the last several years), we reach 20 Mbps per subscriber downstream and approximately 200 kbps upstream 10 years from now (2028). For testing RFoG performance, somewhere between today's load and the expected load 10 years out should be appropriate. We will go even a bit more aggressive and use 20 Mbps downstream and about 1.2 Mbps upstream in the following analysis. These values are shown in *Figure 3*, below.

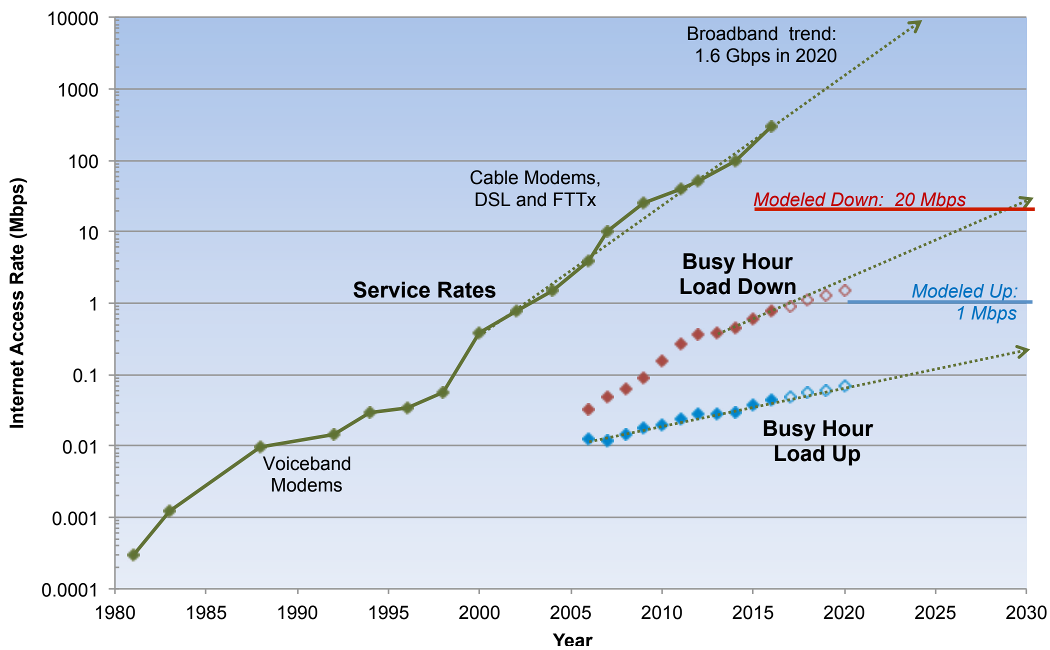

Figure 3: Historic, Projected and Modeled Average Peak Load Per Subscriber

Once we have selected the load, we need to find a Representative Set of Sessions (RSS) to use for a given test run. As mentioned earlier, the sessions are modeled with random initiation times, so many different sets of flows can be generated for a limited time test.

To find a representative set of sessions, we generated 10,000 sets of sessions for a 140-second test. We examined the sets for both overall session data load and the number of subscribers with data transmissions in the test window. A scatter plot of those two properties is shown in *Figure 4*. (For statistics buffs, the correlation coefficient between these two properties is 0.44) From these 10,000 trials, we picked two sets of sessions for use in our RFoG testing. They are shown with the orange circle (RSS 1) and square (RSS 2) in *Figure 4*.

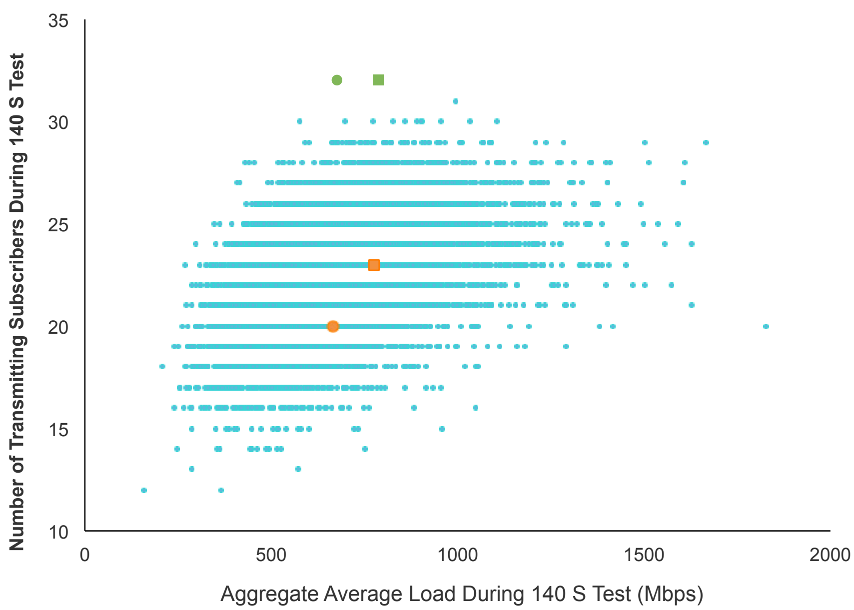

Figure 4: Scatter Plot of Average Load Per Subscriber and Total Number of Transmitting Subscribers in 10,000 Sets of Simulated Sessions

Figure 5 shows a normalized histogram of the average aggregate data transfer rate per test. Our tests were selected to have an average aggregate data transfer rate slightly higher than the median, coming in at 665 and 775 Mbps. (Mean was 677 Mbps).

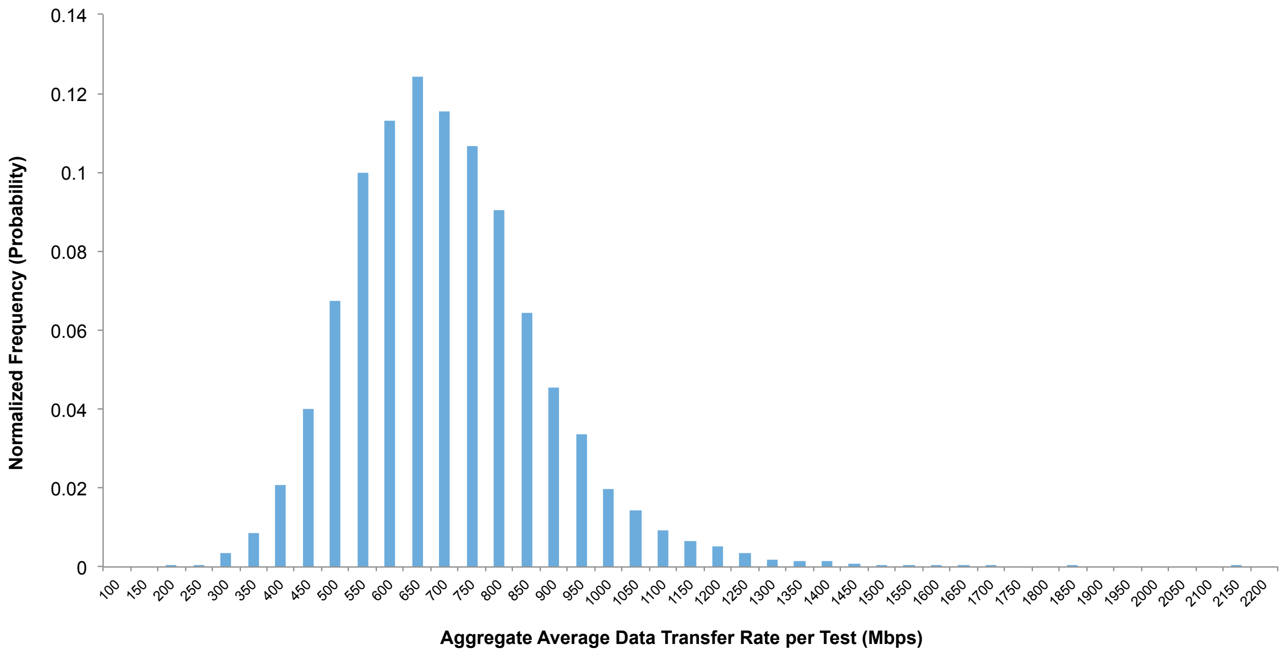

Figure 5: Normalized Histogram of the Average Aggregate Data Transfer Rate During the Test

Figure 6 shows a normalized histogram of the number of transmitting subscribers during the entirety of the test. With 20 and 23 transmitting subscribers, the sets we chose are near the median (and mean) of that distribution (mean was 22.5), but to better assess the sensitivity to the number of transmitting subscribers in the test, we augmented the sessions with additional TCP sessions for the 9-12 idle links. These test cases are shown as green circle (which we will refer to as Augmented Representative Set of Sessions (ARSS) 1) and square (ARSS 2) in *Figure 4*.

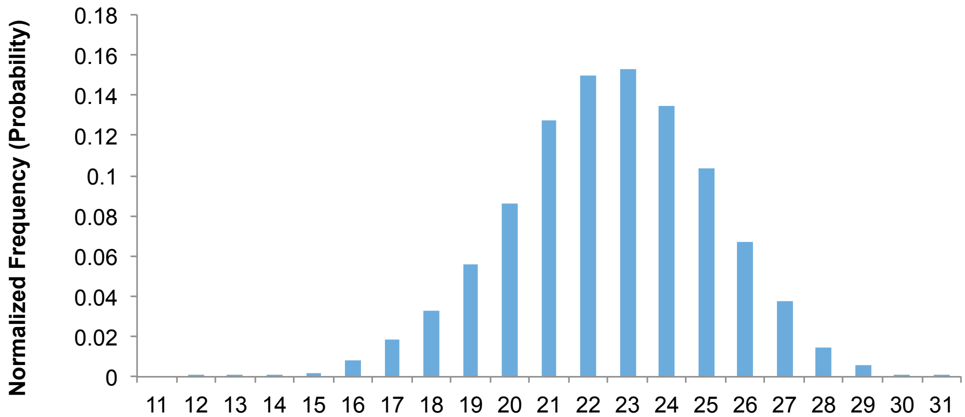

Figure 6: Normalized Histogram of the Number of Subscribers Transmitting During the Test (140 S Test, 10000 Trials)

RFoG, DOCSIS 3.0 and DOCSIS 3.1

In addition to the traffic pattern, the probability of simultaneous transmission depends on the scheduling opportunities in the DOCSIS upstream Media Access Control (MAC) address. With DOCSIS 3.0, there are typically no more than four 6.4 MHz upstream channels, each using Time-Division Multiple Access (TDMA) to schedule transmission. Therefore, there will be at most four (five if we include Set Top Box (STB) Out of Band (OOB) signaling) R-ONUs that are transmitting at a given time.

DOCSIS 3.1 uses Orthogonal Frequency-Division Multiple Access (OFDMA) for the upstream MAC, and the maximum number of simultaneous transmitters is not simple to compute. Upstream transmission grants are issued in units of mini-slots, which occupy 400 kHz of spectrum (eight tones for a 2k Fast Fourier Transform (FFT)) by a depth of N frames, where N can range from six to 36.

Depending on the modulation order and Forward Error Correction (FEC) parameters, the number of bits in a mini-slot can vary. The minimum size grant must contain at least one mini-slot, and also must contain at least 700 bits (420 information bits + 280 parity bits), as that is the minimum code-word length. The data capacity of the grants is further reduced by the presence of pilot tones within the grant. DOCSIS 3.1 specifies required and optional patterns of pilot tones. The optional patterns (patterns five to seven) have reduced overhead and use from four to seven tones-symbols for a single mini-slot grant. *Table 1* shows some exemplary values for minimum grant sizes based on the framing parameters chosen (2K FFT), and minimum overhead pilot patterns.

From the table, we can see that to make grants efficient for single packets, the frame length should be kept short. For 64-Quadrature Amplitude Modulation (QAM), and frame lengths of 16 symbols or less, grants to transmit a 64-byte packet will need to at least occupy two mini-slots.

Grant Size in Mini-slots	Modula-tion	Bits per Symbol per Tone	Tones per Mini-slot	Symbols per Frame	Pilot Symbols per Frame	Minimum Data Symbols per Grant	Total Bits per Mini-slot-Frame	Minimum Parity Bits	Minimum Infor-mation Bits per Grant	Minimum Informa-tion Bytes per Grant
3	64-QAM	6	8	6	10	134	804	280	524	65
2	256-QAM	8	8	6	8	88	704	280	424	53
2	512-QAM	9	8	6	8	88	792	280	584	73
1	64-QAM	6	8	16	4	124	744	280	488	61
1	256-QAM	8	8	16	4	124	992	280	744	93
1	512-QAM	9	8	16	4	124	1116	280	872	109
1	64-QAM	6	8	36	4	284	1704	560	1144	143

Table 1: Data Capacity of Minimum Sized Grant vs. Frame Length and Constellation Size

The number of mini-slots per symbol depends on the amount of spectrum allocated to upstream transmission. For 24 MHz of spectrum and a 16 symbol frame length; there could be as many as 30 simultaneous transmissions. With a six-symbol frame length, this drops to 20, but it is still quite an increase from four or five. While this is theoretically possible, it will only occur if there is data waiting to be sent by 20-30 CMs and the Cable Modem Termination System (CMTS) issues grants of single packets to all waiting CMs in a single OFDMA frame. More likely, an upstream transmission grant will occupy more mini-slots, some CMs won't have data to send on a given frame and the number of simultaneous transmissions will remain small. This can be confirmed both with analysis and testing.

The Number of Simultaneous Transmitters in User Data

With the Virtual Subscriber Test system, the number of simultaneous transmitters in an interval can be measured in the absence of the CMTS, CMs and RFoG. This provides an upper bound on the number of simultaneous transmitters in the DOCSIS 3.1 environment. *Figure 7* shows the normalized histogram of the number of simultaneous transmitters in a 420 μs window generated by the Virtual

Subscriber system with 20Mbps x 1Mbps (Down/Up) average load per subscriber. We picked 420 µs because a six symbol OFDMA frame (2k FFT) produces approximately a 420 µs frame length. This data does not include the effects of CMTS scheduling, which can group transmissions for a given CM/R-ONU into larger grants and thus reduce the number of transmitters per frame. Even so, over 90 percent of frames will have five or less transmitters.

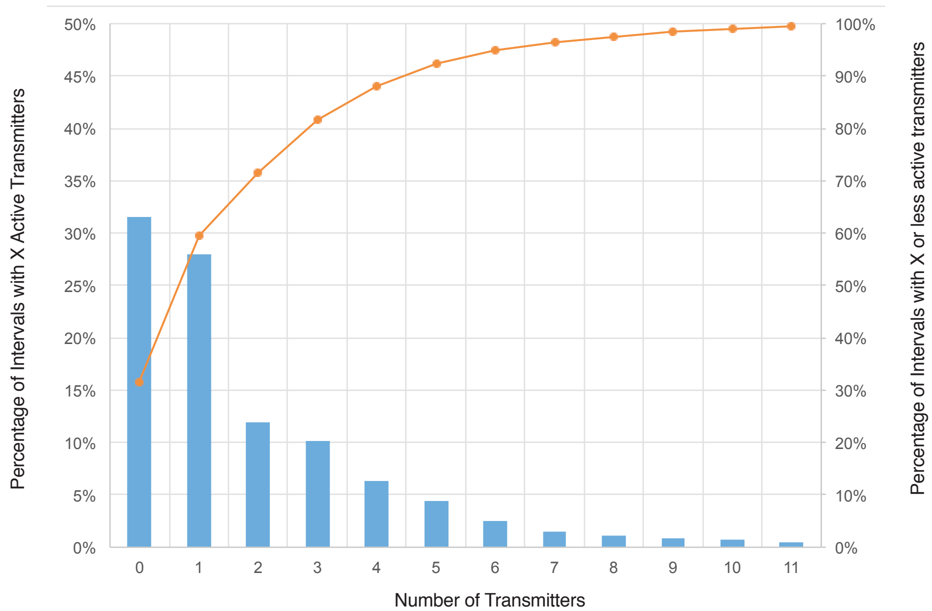

Figure 7: Histogram and Cumulative Percentage for the Number of Transmitters in a 420 µs Interval

This data shows that even if DOCSIS 3.1 has the capability to schedule 20-30 simultaneous transmissions, the data is unlikely to be available to use all of those opportunities. Based on this, we should expect to see slightly more simultaneous transmitters from DOCSIS 3.1 than with DOCSIS 3.0, but these events will occur with low probability.

OBI Impact On Data Transmission

OBI causes the phase noise of the laser to show up as noise in the RF modulating signal. [1] This increased noise floor lowers the Carrier-to-Noise Ratio (CNR) at the receiver, causing errors in the received data. Since DOCSIS uses Forward Error Correction (FEC), some percentage of these errors can be corrected by the FEC (RS in DOCSIS 3.0 and LDPC in DOCSIS 3.1). The Low-Density Parity-Check (LDPC) coding in DOCSIS 3.1 offers about 2.5 dB more noise immunity for 64-QAM than the Reed-Solomon Forward Error Correction (RS-FEC) in DOCSIS 3.0[4]. This additional noise immunity can allow the FEC in DOCSIS 3.1 to correct more OBI-caused errors than the FEC can in DOCSIS 3.0. Any errors that cannot be corrected by the FEC show up in the count of Uncorrected Codewords (UC), and an estimate of the impact of OBI in DOCSIS performance can be made by computing the percentage of total codewords that are uncorrected.

4 *This number is based on difference in the required CNR for downstream 64-QAM in the DOCSIS 3.1 and 3.0 PHY specs. While there is a specification for upstream performance in D3.1, there is no spec for upstream performance in the D3.0 PHY specification.*

Test Results

Testing the impact of OBI with *LambdaShift*™ R-ONUs was performed using the virtual subscriber test configurations with a DOCSIS 3.1 CMTS and CMs. The testing was performed at an independent third-party test lab. The CMTS was configured to use 2k FFT, 64-QAM, 24 MHz of spectrum, a frame size of 16 symbols, Cyclic Prefix (CP) of 256 (2.5 us) and Roll-off Period (RP) of 128. (The units were also successfully operated with the minimum CP values.) One set of 32 ONUs was tested. Three different power-on cycles were tested, and the session sets representing the orange and green dots in *Figure 4* were used. In *Figure 8*, the round markers correspond to the RSS/ARSS with lower aggregate data rates (points on the left in *Figure 4*) and the square markers correspond to the RSS/ARSS with the higher aggregate data rates (points on the right in *Figure 4*.)

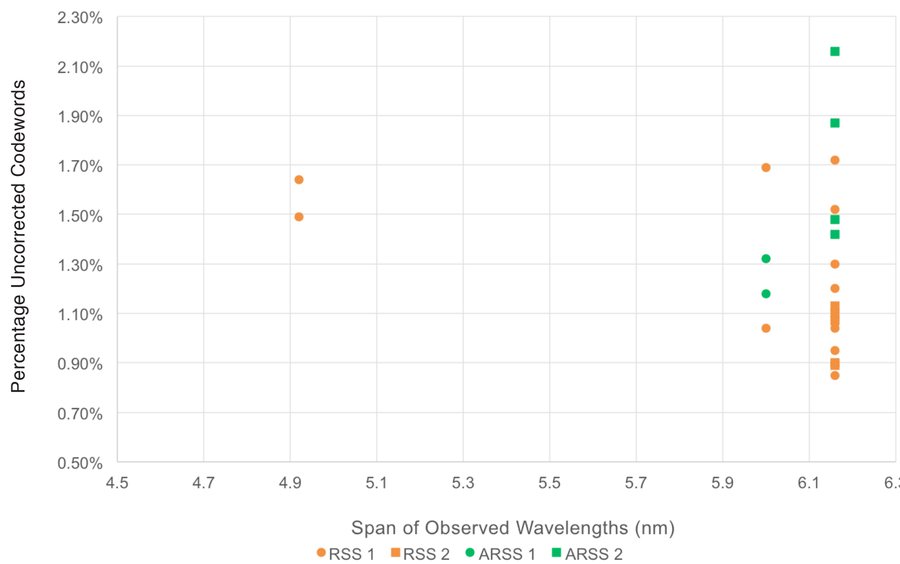

Figure 8: Percentage Uncorrected Codewords vs. Span of Observed Wavelengths for Four Different RSS Tests of *LambdaShift*™ R-ONUs with DOCSIS 3.1

While the results of each test run vary, the trend of the uncorrected codeword rate dropping as the wavelength span increases is apparent. Comparing the span of observed wavelengths on this plot with that of *Figure 2*, we see that the test conditions were all in the worst 10 percent of all expected wavelength spans, so we expect that UC rates in the field, even with the same high levels of traffic will be, on average, lower than the test results.

Virtual Subscriber testing was also conducted with an additional subscriber performing an upstream speed test (iperf). Upstream speed test rates of 35 to 50 Mbps were observed, but the UC rates remained between 1 and 1.5 percent.

Conclusions

Obtaining repeatable and representative test results for OBI is difficult, and it is easy to unintentionally generate test results that differ substantially from what is seen in actual deployment. Testing with traffic patterns that are representative of actual traffic is an important component of good OBI test processes. The test results shown here indicate that ADTRAN La*mbdaShift*™ R-ONUs are able to be used with DOCSIS 3.1 OFDMA upstream and provide good performance.

[1] T. H. Wood and N. K. Shankaranarayanan, "Operation of a Passive Optical Network with Subcarrier Multiplexing in the Presence of Optical Beat Interference, " *Journal of Lightwave Technology*, vol. 11, no. 10, October 1993.

Acronym	Description
ARSS	Augmented Representative Set of Sessions
CAGR	Compounded Annual Growth Rate
CM	Cable Modem
CMTS	Cable Modem Termination System
CNR	Carrier to Noise Ratio
CP	Cyclic Prefix
dB	Decibel
DFB	Distributed FeedBack
DOCSIS	Data Over Cable Service Interface Specification
FDMA	Frequency Division Multiple Access
FEC	Forward Error Correction
FFT	Fast Fourier Transform
kbps	kilobits per second
LDPC	Low-Density Parity Check
MAC	Media Access Control
Mbps	Megabits per second
MTBS	Mean Time Between Sessions
nm	nanometer
OBI	Optical Beat Interference
ODN	Optical Distribution Network
OFDMA	Orthogonal Frequency Division Multiple Access

Acronym	Description
ONU	Optical Network Unit
OOB	Out of Band
OSA	Optical Spectrum Analyzer
OTT	Over The Top
PHY	Physical Layer of OSI Model
PON	Passive Optical Network
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
RFoG	Radio Frequency over Glass
R-ONU	RFoG Optical Network Unit
RP	Roll-off Period
RS-FEC	Reed-Solomon Forward Error Correction
RSS	Representative Set of Sessions
STB	Set Top Box
TCP	Transmission Control Protocol
TDMA	Time Division Multiple Access
UC	Uncorrected Codewords
UDP	User Datagram Protocol
VoIP	Voice over Internet Protocol
µs	microsecond

ADTRAN, Inc.
901 Explorer Boulevard
Huntsville, AL 35806
256 963 8000

General Information
800 9ADTRAN
www.adtran.com/contactus

Canada Headquarters— Toronto, Ontario
+1 877 923 8726
+1 905 625 2515
sales.canada@adtran.com

Canada—Montreal, Quebec
+1 877 923 8726
+1 514 940 2888
sales.canada@adtran.com

Mexico and Central America
+1 256 963 3321
+1 52 55 5280 0265 Mexico
sales.cala@adtran.com

South America
+1 256 963 3185
sales.brazil@adtran.com
sales.latam@adtran.com

TL19.1270

ADTRAN is an ISO 9001, ISO 14001, and a TL 9000 certified supplier.