



Impact of Crosstalk from Non-Standard Links on IEEE 802.3bz

VERSION 1.0
OCTOBER 2017

George Zimmerman, CME Consulting, Inc.
Amrik Bains, Principal Engineer, Cisco Systems
Ken Naumann, Technical Lead, Cisco Systems

Introduction

The performance of 2.5Gb/s and 5Gb/s 802.3bz links, like all Ethernet BASE-T links since 1000BASE-T, is determined by the alien crosstalk generated by other transmissions in the user's network, generally from adjacent links. IEEE 802.3bz™ 2.5G/5GBASE-T and NBASE-T™ links use the same signaling and transmit power levels, making them compatible in terms of alien crosstalk performance. The development and deployment of these BASE-T networks relies on careful planning of the transmit power levels and frequencies so that multiple rates and technologies may coexist with predictable performance. By using standards-based transmit power levels and spectral densities, along with length based transmit power back off levels, suitable link performance in the face of alien crosstalk can be managed as a function of the cabling system and installation.

Systems utilizing nonstandard signaling at 2.5Gb/s and 5Gb/s recently appeared and have been confirmed by multiple parties to exploit transmit power levels greater than those in standardized systems. These increased transmit power levels increase the noise levels seen by other links and may degrade link performance for standardized transmission technologies in the same network. The resulting interference from the proprietary modes may reduce the predictability and performance of both IEEE 802.3bz and NBASE-T systems, and also of a wide range of standardized technologies, including IEEE 802.3an 10GBASE-T, and even, in some cases, 1000BASE-T links.

This white paper first summarizes the impact of the increased noise level on standardized systems, and then briefly reviews the history of managing alien crosstalk in BASE-T Ethernet, beginning with IEEE 802.3an-2006 10GBASE-T, and continued in NBASE-T and IEEE 802.3bz-2016 2.5G and 5G BASE-T technologies. Finally, it presents the observed problem of higher power using the non-standard implementation vs. standard IEEE 802.3bz 2.5Gb/s and 5Gb/s on three test setups including 6-around-1 bundled cables.

Managing Alien Crosstalk in BASE-T Ethernet

Increased power levels from proprietary systems are significant because they can impact the relative levels of alien crosstalk into other links. Ever since 1000BASE-T, Ethernet PHY performance has required management of alien crosstalk. The 1000BASE-T standard recognized the potential of alien crosstalk in bundled and hybrid cables (see IEEE Std 802.3-2015 at 40.A.1.2) by making recommendations for allowed coupling between cables, but fell short of making alien crosstalk coupling a specified

This NBASE-T AllianceSM white paper describes observations of non-standard operating modes in products advertising IEEE 802.3bz performance. It describes the importance and history of alien crosstalk management in BASE-T Ethernet, including NBASE-T/IEEE 802.3bz, summarizes measurements and then shows the impact of interference, in the form of alien crosstalk, from these non-standard devices on IEEE 802.3bz-compatible links, such as NBASE-T. It recommends use of devices with standard-compliant transmit power and power spectral densities to maintain good performance in BASE-T networks.

requirement on link segments supporting 1000BASE-T. Similarly, alien crosstalk performance of Category 5e and Category 6 cabling (ISO/IEC Class D and Class E channels, respectively) is not specified in cabling standards. 10GBASE-T introduced alien crosstalk requirements, and, in tandem, cabling standards bodies introduced Category 6A cabling (and ISO/IEC Class Ea channels) which specified alien crosstalk levels. NBASE-T technology extended the methodology from 10GBASE-T to manage alien crosstalk on existing Category 5e and Category 6 (Cat5e/Cat6) cabling where alien crosstalk performance is not specified by cabling standards. Transmit power levels and power spectral densities were carefully chosen so as not to impair the performance of the existing deployed Ethernet systems. This was adopted by, and standardized in, IEEE 802.3bz-2016. Recent observations confirm, however, that some vendors have added their own 2.5Gb/s and 5Gb/s modes and jacked the power up, potentially ruining the fun for every other link in the bundle.

Standardized BASE-T systems specify alien crosstalk performance of cabling by analyzing the noise from potential disturbing signals in the cabling of the network. When new or nonstandard signals are introduced into an installation, the performance of the standardized systems will suffer if the new signals use higher power levels or have higher power spectral density. First 10GBASE-T, in IEEE 802.3an, and later 2.5G/5GBASE-T estimated the signal-to-noise ratio (a measure of receiver performance) based on the potential interference. The next two subsections describe in more detail how standardized BASE-T systems do this, and why introducing different transmit power levels and power spectral densities might be problematic.

Alien Crosstalk Margin Computation in 10GBASE-T

The 802.3an-2006 10GBASE-T standard was the first to consider alien crosstalk as a primary limiting noise source¹, providing a metric, the Alien Crosstalk Margin Computation (ACMC), for specifying a link segment (see IEEE Std 802.3-2015 at 55.7.3.3).

The ACMC procedure considered both alien near-end crosstalk (PSANEXT) and alien far-end crosstalk (PSAFEXT) from many sources. The specified procedure “ensures the total combined PSAFEXT and PSANEXT coupled into a duplex channel is limited in order to maintain the minimum signal-to-noise ratio”, by implementing an optimum mean-square decision feedback equalizer SNR calculation, also known as a “Salz” SNR calculation after its original author². The well-known theoretical method requires assumptions of the signal and noise power spectral densities. Since 10GBASE-T was the first standardized signal where alien crosstalk was an issue, the ACMC made simplifying assumptions, including that 10GBASE-T signals compliant to the standard represent the worst-case noise a receiver is expected to see. This means that the transmit power levels of both the disturbed and disturbing signals follow the specifications in IEEE Std 802.3-2015 Clause 55, including the transmit power, transmit PSD mask and the power back off schedule³. If a proprietary signal with a higher transmit power spectral density level, occupying a substantial portion of the frequency range of 10GBASE-T were to be introduced, it could become the worst-case interferer, increasing the level of interference, jeopardizing performance of other, standards-compliant links. Even when the alien crosstalk margin computation is not used to qualify the link, the alien crosstalk coupling specifications for Cat6A cabling, and for IEEE Std 802.3-2015 Clause 55.7.2.4, were derived using similar assumptions, and link performance can be similarly jeopardized by proprietary links.

Alien Crosstalk Limited Signal-to-Noise-Ratio (ALSNR) Criterion for NBASE-T and IEEE 802.3bz

Alien crosstalk specifications for 2.5G and 5GBASE-T link segments are based on the same methodology as the ACMC and are called the ALSNR criterion. The ALSNR criterion is based on the same SNR calculations as the ACMC, but analyzes the more complex situation with multiple types of disturbers. The ALSNR criterion specified in IEEE Std 802.3bz-2016 126.7.3 (and applicable to NBASE-T, as well) considers the impact from any

disturbing mixture of 1000BASE-T, 2.5GBASE-T, 5GBASE-T or 10GBASE-T signals on a disturbed 2.5G or 5GBASE-T link. As in the 10GBASE-T ACMC, standards-compliant disturbing signals are assumed, and a “Salz” SNR calculation is performed for each mixture of disturbers, found in Equation 126-35 of IEEE Std 802.3bz-2016. The noise spectrum into the receiver is computed from the PSD templates defined in Table 126-22 of IEEE Std 802.3bz-2016, with the standardized transmit PSD masks, and power back off levels found in the relevant IEEE Std 802.3 clauses for each transmission technology. While NBASE-T technology uses these same transmit PSD templates and power levels, if a transmitter was to transmit at a higher level, for example x dB above the standard, its contribution to the overall noise would increase similarly. When that disturber is the dominant disturber type, the degradation in the ALSNR margin approaches that same x dB, lowering link performance for systems complying with the standard’s power levels.

Summary of Measurements of IEEE 802.3bz and Proprietary Modes

Measurements performed by independent parties have confirmed non-IEEE-compliant proprietary transmit levels and power spectral densities (PSDs) for 2.5G and 5GBASE-T products [Cibula, McClellan, Naumann]. These measurements were conducted on systems which were clearly labeled as supporting IEEE 802.3bz, without mention of a proprietary 2.5G or 5GBASE-T mode. Both measurements showed that when the systems were connected to link partners with a similar proprietary physical layer device (e.g., another port on the same multi-port switch), increased transmit power was observed for both 2.5G and 5G modes. Transmit power spectral density shapes were approximately unchanged, except for the increased level.

Long Link Testing

Figure 1 shows a 5GBASE-T standards-compliant transmit PSD and the 5Gb/s transmit PSD of the proprietary link relative to the IEEE 802.3bz upper limit for the 5GBASE-T PSD on long (100m) links. Figure 2 shows the same measurements for 2.5Gb/s operation. The comparison indicates between 2 and 3 dB higher transmit power on long links.

Additional measurements showed that when the systems were connected to an NBASE-T or IEEE 802.3bz compliant link partner

1. While the 1000BASE-T standard recognized the potential of alien crosstalk in bundled and hybrid cables (see IEEE Std 802.3-2015 at 40.A.1.2), alien crosstalk was not a requirement on link segments for 1000BASE-T.

2. J. Salz, “Optimum Mean-Square Decision Feedback Equalization”, Bell System Technical Journal, Vol. 52 No. 8, October 1973 pp. 1341-1373.

3. Because 10GBASE-T was considered both the victim signal and the disturber in the ACMC, the shape of the transmit power spectral density (PSD) of the disturbed and disturbing signals were identical, and therefore the power spectral density itself is not present in the equations for the “Salz” computation, Equations (55-44) and (55-51) in IEEE Std 802.3-2015.

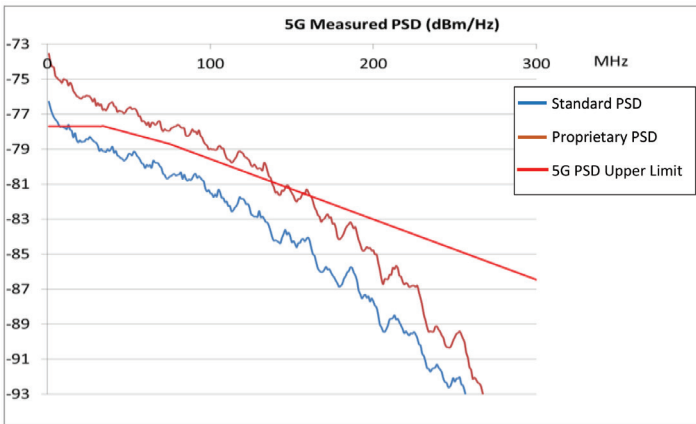


Figure 1: Measured 5Gb/s Standard-compliant vs. Proprietary non-standard compliant Transmit PSDs and IEEE 802.3bz PSD Mask [McClellan]

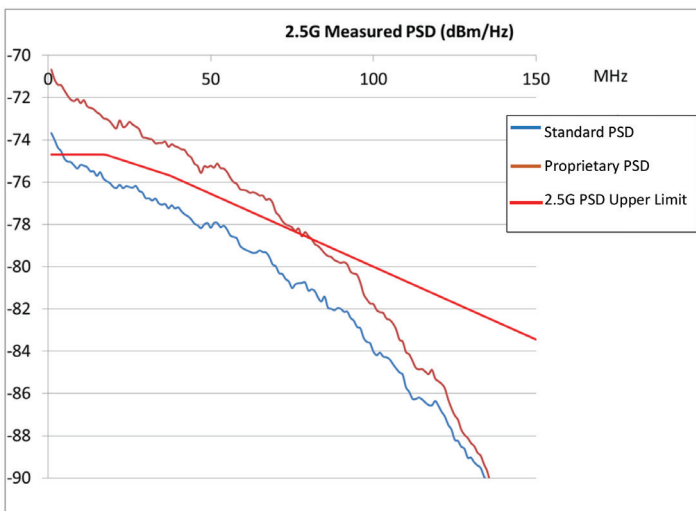




Figure 2: Measured 2.5Gb/s Standard-compliant vs. Proprietary non-standard compliant Transmit PSDs and IEEE 802.3bz PSD Mask [McClellan]

(e.g., a design using transceiver of another vendor), the resulting link showed transmit power levels and power spectral densities (PSDs) prescribed by both IEEE 802.3bz and NBASE-T for both 2.5G and 5GBASE-T modes. Figure 3 and Figure 4 show the transmit PSDs observed when the link partner (the connected unit) was of the same vendor (Red) curve, and a different vendor (Green), at 5 Gb/s and 2.5Gb/s respectively. The difference between the green and red curves for the different PHY connections shows the same 2 to 3 dB higher transmit power on long links seen in Figure 1 and Figure 2.



Higher Tx power is again observed in back-to-back vs. interop configurations. Total non-IEEE transmit power is estimated to be ~2 to 3 dBm higher.

Figure 3: Transmit PSDs at the 5Gb/s rate, with both similar and different PHYs on the link partner [Cibula]



Higher Tx power is again observed in back-to-back vs. interop configurations. Total non-IEEE transmit power is estimated to be ~2 dBm higher.

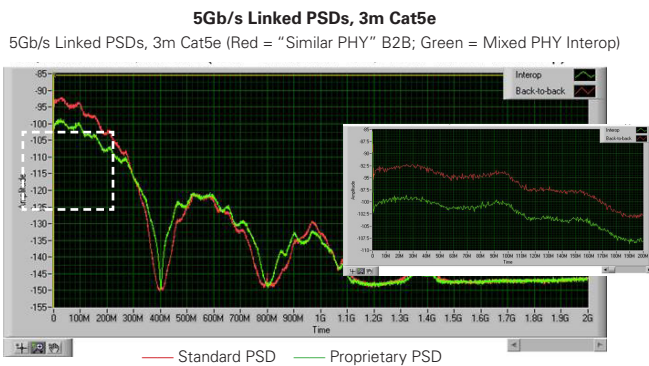
Figure 4: Transmit PSDs at the 2.5Gb/s rate, with both similar and different PHYs on the link partner [Cibula]

Short Link Testing

In addition to specifying the maximum transmit PSD level on long links, 2.5G/5GBASE-T and 10GBASE-T standards manage interference by specifying reductions in the allowed transmit power on short links. This “power back off” is an important method of managing far-end crosstalk’s impact on long links disturbed by shorter links. When the systems were tested on shorter links with similar link partners, the power back off levels observed were less than those specified in IEEE 802.3bz, meaning that they transmitted at still higher levels than prescribed in the standard. See Figure 5 for an example showing 5-7 dB higher transmit power on short links than NBASE-T/IEEE 802.3bz compliant PHYs at 5Gb/s rates. As in previous figures, the red curve shows the higher transmit level when the units were linked to a partner with the same proprietary mode [e.g., another port

of the same device] relative to the lower green curve when they were linked with a unit compliant to IEEE 802.3bz transmit power specifications. In this case, the green (compliant) PSD is from 5 to 7 dB lower than the red PSD for 5Gb/s transmission. This shows that the suspect units *can* operate at a lower transmit PSD but elect not to do so when establishing links in the proprietary mode.

Short link power levels for 2.5Gb/s were not reported, but management registers indicate 2dB greater power-back-off at short reach than would be mandated by the standard, potentially normalizing power levels when the proprietary links are on short links [McClellan, slide 2].



Note higher Tx power in the back-to-back (same PHY, assumed non-IEEE mode) configuration compared to the interop (different PHYs, assumed IEEE-compliant mode) configuration. Total non-IEEE transmit power is estimated to be ~6dBm higher.

Figure 5: Decreased power back off level on short links for 5Gb/s transmission with similar (proprietary) PHYs [Cibula]

As demonstrated by these measurements, the system under test showed higher transmit powers at both long and short links for 5Gb/s transmission. Transmit power was approximately 2.5dB higher on long links for 5Gb/s, and 2 to 3 dB higher for 2.5Gb/s. On short links, because of differences in power back off, transmit power for 5 Gb/s was 5 to 7 dB higher than expected.

Impact of IEEE 802.3bz and Non-standard 2.5Gb/s and 5Gb/s Operation in the Same Cable Bundle

Testing of live systems was performed to assess the impact of the increased crosstalk on a standards-compliant victim. The tests used a standard cabling configuration known as 6-around-1 where six cables surround the victim under test. The six cables are then activated with different disturbing signals to assess the impact. The impact may be observed in NBASE-T systems employing

* Note: The number of non-IEEE compliant links was limited to 4 instead of 6 due to only 8 ports of the non-IEEE compliant system available. System performance would be at least as bad if 2 additional non-IEEE links were added.

NBASE-T Downshift by watching for a lower transmission speed when the interference levels are too high to support 2.5Gb/s or 5Gb/s traffic.

The impact of the increased power level employed by a non-IEEE 802.3bz compliant system is readily demonstrated by comparing 6-around-1 behavior using IEEE 802.3bz compliant disturbers the non-IEEE 802.3bz compliant disturbers in a 2m (patch cord) + 100m (horizontal) + 2m (patch cord) Cat5e channel.

In the IEEE 802.3bz case, all 6 disturbers are running IEEE 802.3bz compliant signaling. The result for this case is all disturbers and the victim link at 2.5Gb/s. Diagrams 1 and 2 show a 2.5Gb/s idle waveform and a zoom to estimate symbol time. 2.5Gb/s is the expected rate for this cabling setup with standards-compliant disturbers.

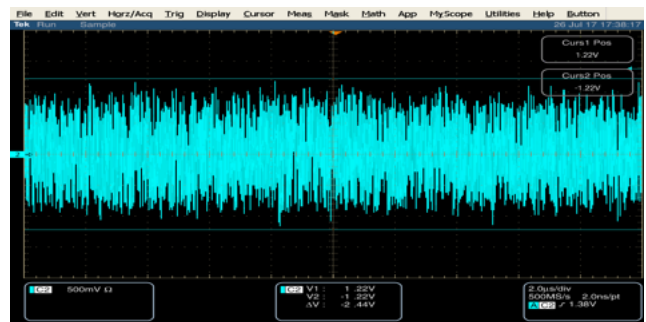


Figure 6: Amplitude of Victim Pair A IEEE802.3bz 2.5Gb/s 2m (p) + 100m (h) + 2m (p) amplitude. All disturbers IEEE802.3bz linked at 2.5Gb/s.




Figure 7: 6-around-1 Symbol time of Victim Pair A IEEE 802.3bz 2.5Gb/s 2m (p) + 100m (h) + 2m (p). All disturbers IEEE 802.3bz linked at 2.5Gb/s.

In the second setup, we employed just 4 (out of possible 6) disturber pairs with non-IEEE compliant links and the other 2 disturber pairs with IEEE 802.3bz compliant links*. What should be emphasized here is the IEEE 802.3bz compliant victim has

now down-shifted to 1Gb/s – a downgrade from the prior IEEE 802.3bz compliant case. It is our estimate that the reason for the lower link speed is due to the cross-talk exhibited by the non-compliant system preventing the victim receiver to converge at 2.5Gb/s. Figure 8 shows that this resulted in a 1Gb/s idle waveform, and Figure 9 shows a zoom to estimate symbol time.

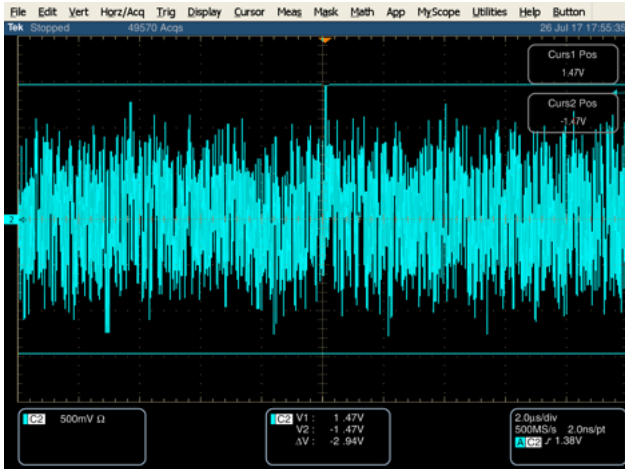


Figure 8: 6-around-1 Amplitude of Victim (Pair A) in IEEE 802.3.40 1000Base-T 6-around-1 channel 2m (p) + 100m (h) + 2m (p) amplitude. 6 disturbers comprised of 2 IEEE 802.3bz 2.5Gb/s links + 4 proprietary non-IEEE 802.3bz compliant links. Here the victim has downshifted from 2.5Gb/s to 1Gb/s after application of 4 non-compliant aggressor links.

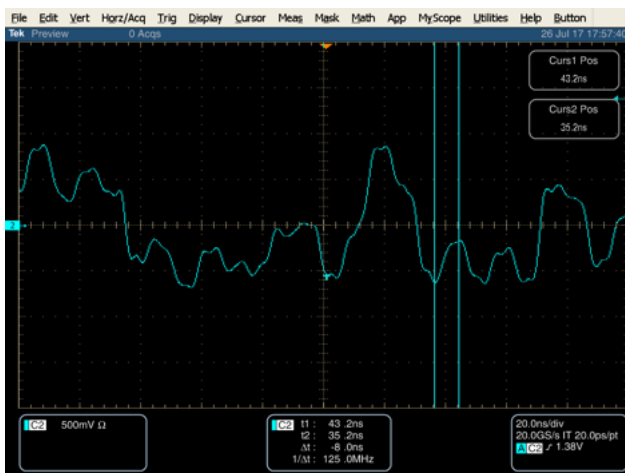


Figure 9: Symbol time of Victim (Pair A) IEEE 802.3.40 1000Base-T in 6-around-1 channel 2m (p) + 100m (h) + 2m (p) amplitude. 6 disturbers comprised of 2 IEEE 802.3bz 2.5Gb/s links + 4 proprietary non-IEEE 802.3bz compliant links. Here the victim has downshifted from 2.5Gb/s to 1Gb/s after application of 4 non-compliant aggressor links.

A snapshot comparison between the IEEE 802.3bz link and non-IEEE compliant link is shown in Figure 10 below. The non-IEEE compliant link shows almost 30% higher amplitude than the IEEE 802.3bz compliant link, (captures taken using same 2m patch + 100m horizontal + 2m patch Cat5e channels). This is consistent with the 2dB power difference reported by the other measurements.




Figure 10: Amplitude of non-IEEE compliant link is much higher (3.5Vpk-pk) than the IEEE 802.3bz compliant link (2.75V pk-pk). High amplitude causes noise contamination on neighboring links.

Pictures of the setup follow:



Figure 11: Compliant and non-Compliant switch set up.




Figure 12: Cat5e 6-around-1 cable rack.

Analysis of the Impact on Existing Standardized Systems

Consistent with the laboratory results, an analysis of the impact of the increased power levels was performed. The mathematics of the ALSNR criterion and the ACMC are clear, and they model the optimum performance of the signal processing underlying 1000BASE-T, 2.5GBASE-T, 5GBASE-T and 10GBASE-T standardized systems. The ALSNR criterion predicts that, in the best case, if many signals are introduced with an increased power level and similar spectral content, such as those measured and reported above, they will degrade the SNR margins of standardized links. Calculations reported in [Cibula] confirm the ALSNR degradation is approximately equal to the number of decibels that the proprietary signal exceeds the standardized power levels. Degradations computed for a particular topology and use case may be slightly less than the increased power level due to the specific alien crosstalk frequency responses in the case examined. This is especially true when the limiting cases for the ALSNR criterion include some amount of standardized 1000BASE-T and/or 10GBASE-T disturbers in the mixture. In all cases, use of the standardized power and PSD levels for 2.5G or 5G signals would result in improved ALSNR.

The full impact of 2.5G proprietary signals is not yet known, because the low frequency content makes alien FEXT a particularly important component of the noise, and the short-link levels have not yet been fully evaluated. However, when long 2.5G proprietary links are coupled into standards-compliant 2.5GBASE-T victims, alien crosstalk limited margins degrade. The loss in margin can be as much as 2 to 3 dB⁴ and is directly proportional to the contribution of the increased transmit power level of the proprietary disturbers in the total alien crosstalk noise. This degradation impacts 1000BASE-T and 2.5GBASE-T primarily, because of the frequency content, but would also impact 5GBASE-T systems to a lesser extent. 10GBASE-T impact should be substantially less (< 0.5dB), because the frequency content of the proprietary 2.5G signals occupies only a fourth of the 400 MHz 10GBASE-T band.

Because the relative increase in short-link power measured is higher than that of long links, the impact of the 5G proprietary signals can vary from up to 2 dB to up to 6 dB, depending on the link configurations. This should primarily impact 5GBASE-T and 10GBASE-T links.

5GBASE-T link margin may be degraded by as much as 2 to 3 dB when long disturbing links are present. The straightforward analysis, that a 2 to 3dB increase in the noise level can cause up to between 2 and 3dB loss in ALSNR, is confirmed by analysis on specific long links in [Cibula], as shown in Figure 13. Because 5GBASE-T occupies approximately half the bandwidth of the 10GBASE-T signal, 10GBASE-T margins can be expected to be impacted by approximately 1-1.5dB when confronted with a mixture of 10GBASE-T and proprietary 5G alien aggressors. However, because of the difference in power back off, the impact of a short reach proprietary signal coupling into a short reach 5GBASE-T or 10GBASE-T signal may be substantially more severe. At short reach, the 6 dB power difference presents a much more significant disturber into a standards-compliant 5GBASE-T or 10GBASE-T⁵ signal utilizing the standards-compliant power back off. For the specific short-into-short disturbing configuration, the ALSNR margin of a standards compliant 5GBASE-T victim degrades by as much as the full 6 dB of the transmit power difference, and a 10GBASE-T compliant link may be impacted by as much as 3 dB. These degradations will vary substantially with the link configurations, but are sufficient in magnitude to result in failed links for 5GBASE-T and 10GBASE-T.

The impact of the disturbances from these non-standard power levels will be situationally dependent, with the magnitude of degradation dependent on the topology and coupling of the links, as well as the number of proprietary links used. Additionally, the impact will be lessened, but still present, in installations utilizing Cat6A or even shielded cabling.

Calculated ALSNR

The calculated ALSNR is the minimum across all 4 pairs of victim

Operating Mode	Pair A ALSNR (dB)	Pair B ALSNR (dB)	Pair C ALSNR (dB)	Pair D ALSNR (dB)
5Gb/s victim, 6x IEEE aggressors	25.03	22.88	22.28	23.15
5Gb/s victim, 6x non-IEEE aggressors	22.78	20.34	19.66	20.66
2.5Gb/s victim, 6x IEEE aggressors	30.36	28.28	27.69	28.36
2.5Gb/s victim, 6x non-IEEE aggressors	29.38	26.90	26.23	27.01

The non-compliant aggressor PHY can reduce the operating margin of a compliant victim PHY by as much as 3 dB

Figure 13: ALSNR degradation findings from [Cibula] for specific long links.

4. The actual computed degradation may be slightly less due to different mixes of disturbers and the actual frequency weighting of the alien crosstalk levels.

5. Because the 2.5GBASE-T signal overlaps 1000BASE-T, which lack power back-off across nearly its full bandwidth, the non-backed off proprietary 5Gb/s signal is not expected to substantially degrade 2.5Gb/s performance over that already seen in 1000BASE-T alien crosstalk.

Conclusion

The performance of BASE-T links at gigabit and higher rates is dependent on a cooperative, standardized management of the interference environment, especially alien crosstalk. IEEE 802.3 Ethernet standards have carefully considered transmit power levels and power spectral densities (PSDs) to limit the levels of interference present and provide good performance. Systems should maintain compliance with the transmit power levels and PSDs in IEEE 802.3 standards to be good neighbors and provide good performance for all. Independent measurements disclose proprietary transmission modes for 2.5Gb/s and 5Gb/s with 2 to 6 dB higher transmit power spectral densities than standardized IEEE 802.3bz and NBASE-T signal levels. These increased transmit signal levels violate the foundation of IEEE 802.3 specifications for alien crosstalk levels for 1000BASE-T, 2.5GBASE-T, 5GBASE-T and 10GBASE-T PHYs. The resulting performance degradation to 2.5G, 5G and 10GBASE-T links will vary by link topology and speed used, but can be expected to be similar in magnitude to the disturbing power increase, from 2 to as high as 6 dB, resulting in frame loss, flapping links, failed links and reduction in link rate and/or reach when these proprietary systems are introduced into a standards-based network topology.

References

[Cibula] Hossein Sederat (Aquantia) and Peter Cibula (Independent (for Aquantia)), "The Effects of Non-Compliant MultiGBASE-T PHYs on Network Ecosystem"; NBASE-T TWG contribution, 22 June 2017.

[McClellan] Brett McClellan, Marvell, "MGBASE-T Coexistence Issue"; NBASE-T TWG contribution, May 2017.



3855 SW 153RD DRIVE
BEAVERTON, OR 97003 USA
WWW.NBASET.ORG

The NBASE-T Alliance is an industry-wide cooperative effort of more than 45 companies focused on enabling the development and deployment of products that support 2.5G and 5GBASE-T Ethernet.

To support early product development, the alliance developed and released the NBASE-T specification, which is compatible with the IEEE 802.3bz standard, and includes an added Downshift feature. The alliance focuses on publishing optimizations to the specification, facilitating interoperability and educating the market about the multiple applications of the NBASE-T technology.

© 2017 NBASE-T Alliance, Inc. All rights reserved. Unauthorized use strictly prohibited.

NBASE-TTM, NBASE-T AllianceSM, and all NBASE-T logos shown in this document are registered or unregistered trademarks, service marks, and/or certification marks of NBASE-T Alliance, Inc. or its licensors in the United States and other countries. All rights reserved. Unauthorized use strictly prohibited. Other names, marks, and logos shown in this document are trademarks and/or service marks of their respective owners. Specifications and content subject to change without notice.