White Paper

Delivering High-Quality Video in Your IPTV Deployment

This white paper introduces Broadcom's PhyR™ retransmission system. This innovation achieves the low residual BER required for the deployment of IPTV across the xDSL network, while also achieving higher rates and extended service reach at no extra cost. PhyR overcomes the limitations of the high bit error rates in the DSL standard, giving telecom service providers and carriers a cost-effective, scalable solution that is available today and ready to roll out across existing Broadcom CO and CPE hardware via a simple firmware upgrade.

April 2009

Video Deployments over DSL

As telecom service providers around the world begin to deploy Internet Protocol Television (IPTV) and interactive video applications over their network, they have to provide competitive video quality and user experiences to acquire or at least retain market share.

In this context, DSL is often mixed with other access technologies and, from an operational viewpoint, should be deployed exactly as Fiber or Ethernet, with the same or better performance and without any additional system complexity or infrastructure overhead.

ADSL has been a huge success for fast Internet access, but the job was rather easy from a service perspective; rates are still fairly low (only a few Mbps) and Bit Error rate (BER) requirements are not too stringent, as the TCP-IP retransmission protocol effectively hides transmission errors at these rates.

With the evolution towards IPTV, much lower BER figures are required. Typically, multicast video streams are no longer protected by native retransmission protocols, no more than one error per movie is allowed (equivalent with a BER lower than 10^{-10} for a 120' movies at 20 Mbps), and application-level retransmission protocols are expecting Ethernet-level BER to scale economically.

As explained below, the evolution of DSL technologies towards improvements of rate, reach, and error-free operations is of key importance in addressing both technical and business goals for successful deployments of video over a DSL network.

The requirements are definitively changing, and DSL has to adapt or become irrelevant.

Service Rate—Move to ADSL2+ and VDSL2

To deploy video over the copper access network and ensure that enough bandwidth is available to provide services ranging from a single standard-definition video stream to a combination of multiple standard-definition and high definition video streams, higher bit rates are expected. To meet this need, telecom service providers and carriers worldwide are naturally turning now to the higher-speed DSL technologies that are widely available, such as ADSL2+ or VDSL2.

Service Reach—Reduce Transport Overhead

Service reach is the “guaranteed” amount of DSL data rate that is available at a specific loop length that complies with the service requirements.

DSL data rates and serviceable reach are closely related; higher DSL data rates typically result in shorter serviceable reach and a much higher sensitivity to attainable customers as a function of the distance. A small service reach variation translates into big differences in serviceable customers. Service reach is, therefore, of key importance for an access provider’s business case and time-to-profit; the greater the number of residential homes serviced from an existing CO/RT, the lower the investment into new fiber network extensions and Remote Terminal equipment, and the higher and faster the revenue potential.

While most DSL solutions today have physical layer rate performances close to theoretical limits, BER and robustness requirements for IPTV have a downside effect on service reach. This document explains that the data transport and coding overhead required achieve video’s much stricter BER, and related robustness against impulsive noise becomes prohibitive with today’s implementations, where the overhead is typically above 30%.

This is clearly unacceptable and has to be addressed urgently, as it has a huge impact on service reach and, therefore, the business case.

Error-Free Operations—Use Broadcom PhyR™

The focus in recent years has been on improving the robustness of DSL connections against nonstationary or impulsive noise (leading to residual errors and video artifacts), driven by new video requirements. A considerable effort has been invested in standardization, but with the exception of a new slightly improved interleaving scheme (extensions of some framing parameters and interleaving memory leading to even higher coding overhead), no conclusive progress has been made until now.

However, as is described below, the standard Reed Solomon (RS) + interleaving approach is believed to have reached its technical and acceptability limits, given the associated implementation and rate overhead costs. Hence, there is a need for these requirements to be adapted with a new approach that is radically different from what has been proposed until now.

To address these issues, Broadcom® has developed an innovative data retransmission system implemented above the DSL physical layer. The new data retransmission system is called PhyR™ (pronounced “Fire”) and is described in this document. PhyR has been tested in the field and is proven to provide ten times more protection against impulsive noise compared to a traditional RS + interleaving scheme and this with much less overhead and constraints.

Moreover, due to Broadcom’s unique soft-modem architecture, the PhyR retransmission system can be rolled out across any currently deployed or future Broadcom DSL CO and CPE devices by a simple software upgrade.

Limitations of the Traditional RS + Interleaver Scheme

DSL technologies, including ADSL1/2/2+ and VDSL2, define a Forward Error Correction (FEC) scheme based on a combination of RS coding and convolutional interleaving to provide extra coding gain in first instance and, by extension, protection against impulsive noise (SHINE or REIN).

When ADSL2+ was defined a few years ago, the importance of the impulse noise issue was not identified, and impulse noise protection was made worse as there was no initial provision in the new standards to scale Impulsive Noise Protection (INP) with the increasing line rates.

The most recent amendments to the standard aimed to improve this are interleaving memory and interleaving depth were increased from 16 KB to 24 KB and from 64 to 511 respectively, and the RS decoding speed pushed upward. Those improvements appear largely insufficient today, where DSL struggles to achieve 1 ms impulse noise protection, despite field data demonstrating that at least 5 ms are required.

“They steal your bandwidth” (and reduce your service reach)

The RS + Interleaver scheme suffers from a high overhead burden because for every errored byte, two additional overhead bytes must be transmitted to allow for a successful FEC correction. For example, assuming that an overhead of 10% correctable is tolerable, correcting a burst of 1 ms of impulsive noise (INP = 4 DMT symbols) requires an interleaver depth (or delay) of minimum 10 ms:

INP (ms) = ½ OH . delay (ms)

The ITU amendments mentioned above do not change anything in that fundamental limitation and only intend to be as close as possible to the formula without introducing additional framing limitations. Figure 1 clearly illustrates the capacity losses on an ADSL2+ system that complies with latest ITU amendments for increasing INP when delay is constrained to 8 ms. When INP is set to 4 and at 16 Mbps (4 kilobits per symbol), capacity loss is about 20%!

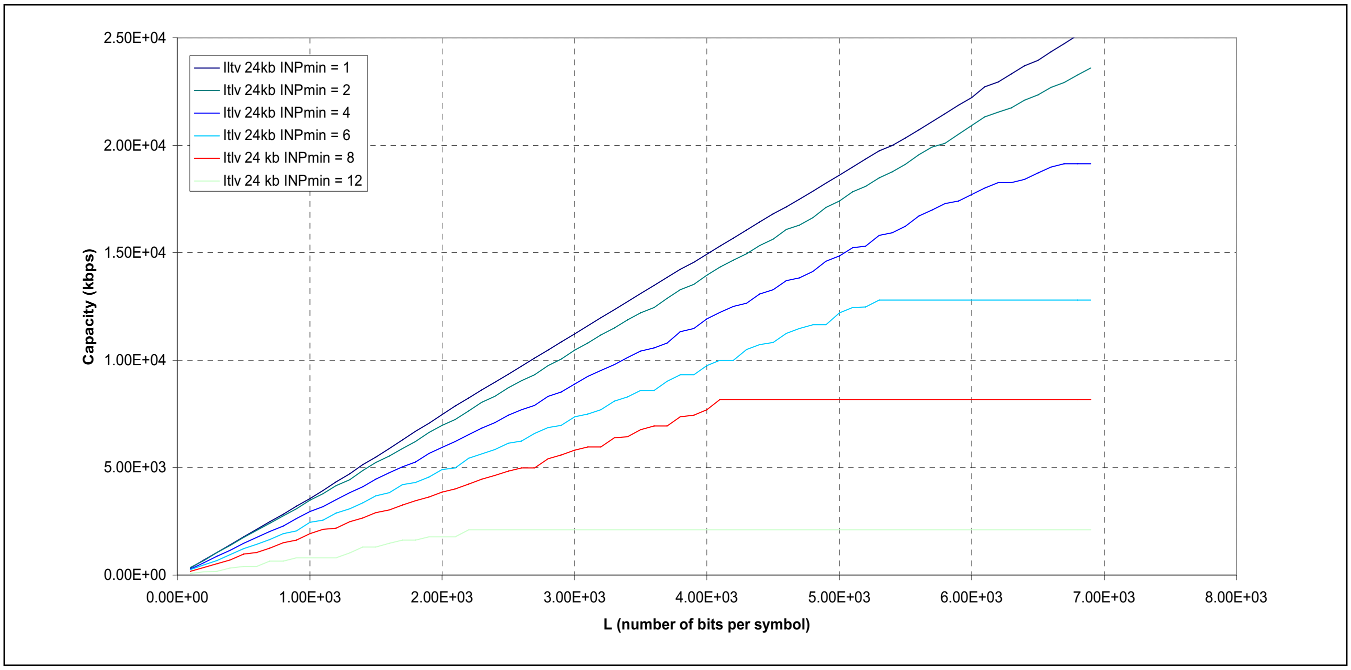

Figure 1: Rate Limitation of an RS + Interleaver Scheme

"They cheat with your margin" (Coding Gain Overbooked)

When we claimed above that the RS coding provided extra coding gain, we were fooling ourselves together with the whole DSL industry!

Indeed, the RS coding gain assumes that the RS decoder is correcting the random stationary errors present when the line runs at a low-noise margin. In those margin conditions, the likelihood is very high that there is no room anymore to correct extra errors from noise impulse popping up on the line and the other way around.

Figure 2 on page 5 shows that with a 2 dB noise margin, a (128,112) RS code with 24 KB interleaving memory will fail to correct an impulse it should correct in 33% of cases. In practice, at least 3 dB of margin is needed to correct more than 99.5% of the expected impulses.

RS FEC capabilities are, therefore, double-booked. In practice, an additional 2–3 dB margin is required to recover the expected BER and INP figures, at the cost of huge capacity loss.

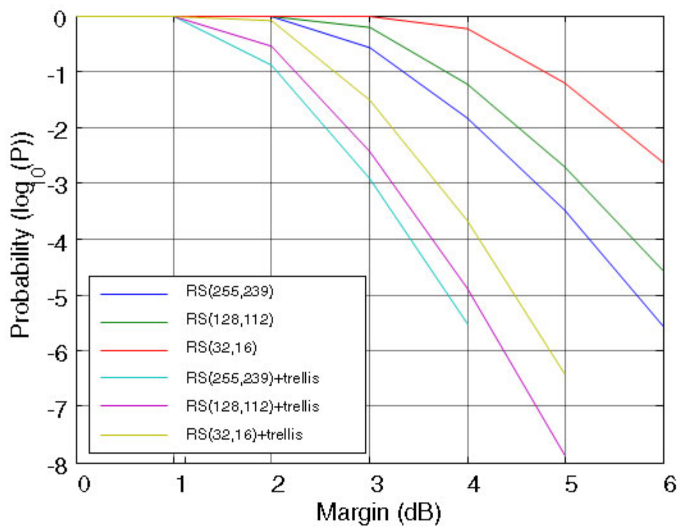

Figure 2: Probability of Not Correcting an Impulse Because of Stationary Noise (Interleaver = 24 KB)

"They sell you systems they can't guarantee at 10^{-10} BER"

The next issue is even more fundamental for DSL systems. Previously, they have been designed with a BER target of 10^{-7} BER, which is far from the 10^{-10} required to provide high quality video services and compare to Ethernet or Fiber.

As explained before, this 10^{-7} BER is highly insufficient for high-quality IPTV distribution. With the typical error patterns observed in DSL modems, a BER of 10^{-7} at 30 Mbps translates into one error every 13 seconds!

Increasing the noise margin certainly helps overcome this issue. However, not guaranteeing the BER reduces the service reach, as, at these levels, many other vicious error sources also play a role (clipping or cable unbalance, for instance). Therefore, no chipset vendor is ready to commit, with today's RS + interleaver scheme, to such low BER values under acceptable deployment conditions.

"They can make it even worse" (Error Propagation)

The next problem is the well-known issue of error propagation caused by the deinterleaver.

The processing that scatters the errors to make them correctable by the RS code also makes the situation much worse when the RS code is stressed beyond its correction capability. Whenever a burst exceeds the RS code correction capability, the impulse affects data over the full depth (or delay) of the deinterleaver, as illustrated in Figure 3 on page 6.

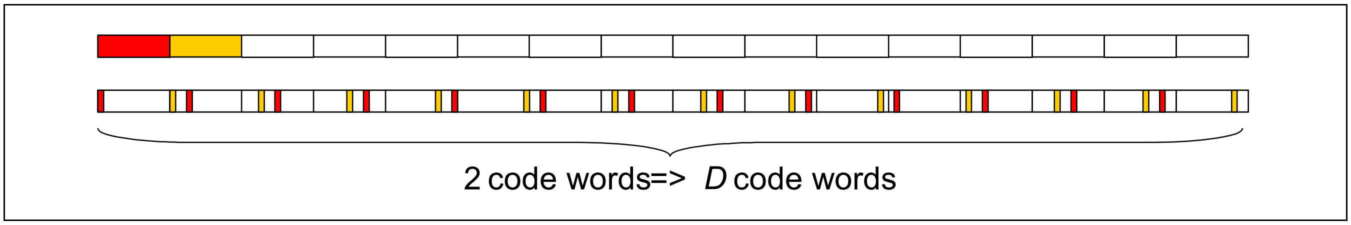

Figure 3: Interleaver Spreading Errors on Multiple Symbols

This means that for a system set with INP = 2 (500 µs) and delay = 8 ms, an impulse slightly above 500 µs can destroy 8 ms worth of data. Remember as well that we are talking about impulse of up to 5 ms in the field, and that +500 µs is very common.

In such a case, the error is likely to cause so many packet losses that it will not go unnoticed, making the work of any MPEG forward error correcting code or application-level retransmission extremely difficult.

"They ask you to pay to do their job" (Provisioning Nightmare)

All the above-mentioned issues make it clear that access providers have to pay a heavy operational price to get only minimal and best effort protection.

Due to the huge price of these high INP settings (reduced service reach), they can only be applied on an individual basis, assuming a setting can be found to cope with both service and INP requirements.

This is a provisioning nightmare: service providers have to develop or buy intelligent network analyzer tools that play with limited INP (RS/interleaver), noise margin, minimum rate, and delay settings and try to limit the rate losses to get a bare minimum INP and/or BER protection.

PhyR™—DSL Physical Layer Retransmission

What is PhyR?

PhyR is a new INP technology based on physical layer retransmission that enables carriers and OEMs to dramatically improve voice, data, and video services to their customers. This new technology is incorporated in Broadcom's industry-leading ADSL2+/VDSL2 firmware and provides major improvements in residual BER as well as resistance against impulsive noise without additional provisioning overhead.

How Does PhyR Work?

As illustrated in Figure 4 on page 7, a retransmission buffer (T) stores the transmitted data, packed in a retransmit unit. When a data unit is received, its Frame Check Sequence (FCS) is checked, and a first retransmission request is immediately launched if it is found to be corrupted. Even if corrupted, the data units are pushed into the receive buffer (R). If the retransmitted data unit arrives while the corrupted one is still present in the receive buffer, the corrupted data is replaced. If the retransmitted data unit does not arrive on time, the corrupted data will be further processed by the receiver data path.

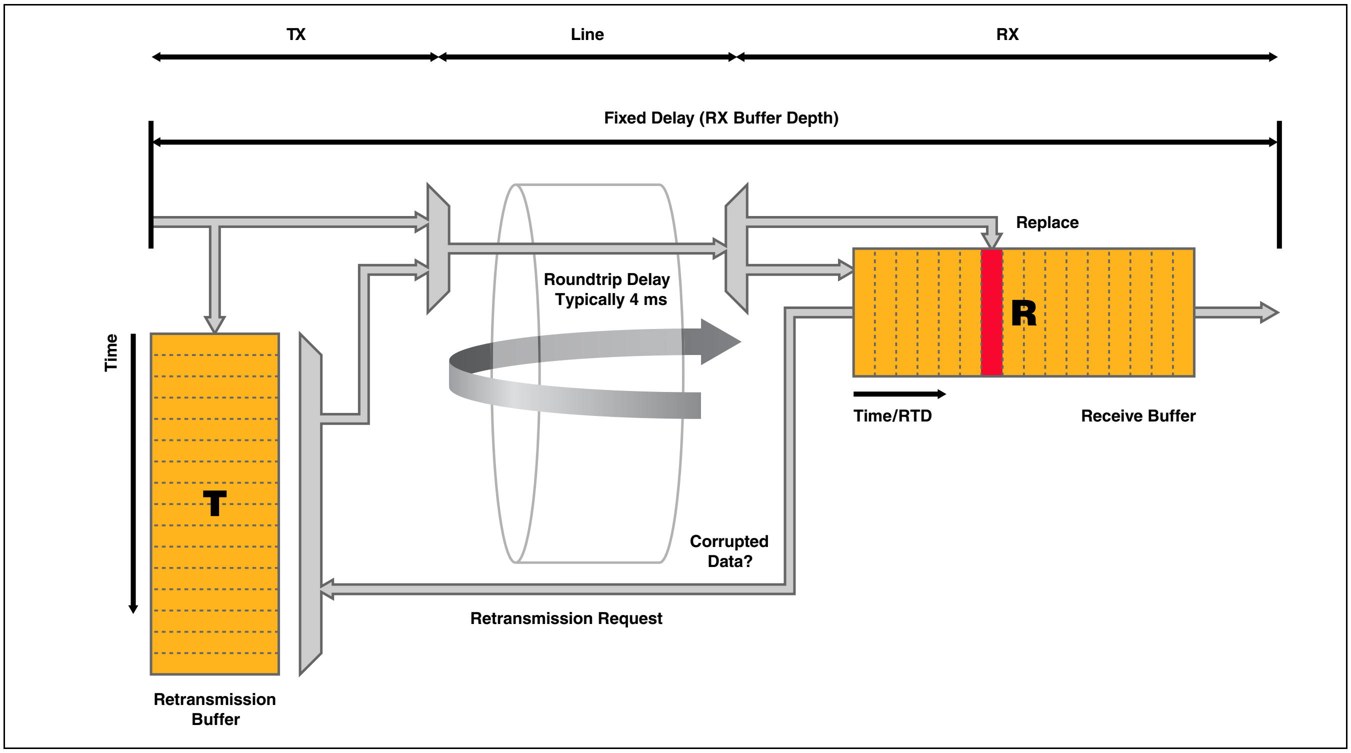

Figure 4: Physical Layer Retransmission Principle

It should also be noted that:

- The end-to-end delay is kept constant, even when retransmission takes place. This makes the end-to-end behavior very similar to a standard interleaving scheme.
- Retransmission does not prevent the use of FEC at the receiver side to correct a data unit prior to retransmission and provide additional coding gain (i.e., complementary).
- The roundtrip delay is quite small, typically 3 to 4 ms. As 2 ms maximum per direction is mandated by the standard today, a link running retransmission operates at comparable (and typically much lower for decent INP) delay than an interleaver-based system.
- By using the PhyR retransmission scheme, an effective INP as high as 16 with a maximum delay of 4 ms at 24 Mbps can be achieved with the same buffer size as today's ADSL2+ interleaver without compromising the effective data rate (the same gain applies to VDSL2).

Let's now review the benefits of a retransmission technique compared to the legacy approaches.

PhyR—Pay Per Fault

The first advantage of a retransmission technique is that the real overhead is only present when errors are there. If errors are present only 1% of the time, the associated cost is only present 1% of the time and not 100%, as with an RS + interleaver approach. There is some steady-state overhead associated with the retransmission of a transport data unit's ID or CRC, but both are negligible.

Hence, as illustrated in Figure 5, for high rates, a retransmission scheme will typically provide 10x the INP protection offered by an interleaver (or a much higher rate for a given INP). Furthermore, the rate is not directly affected by the INPmin settings.

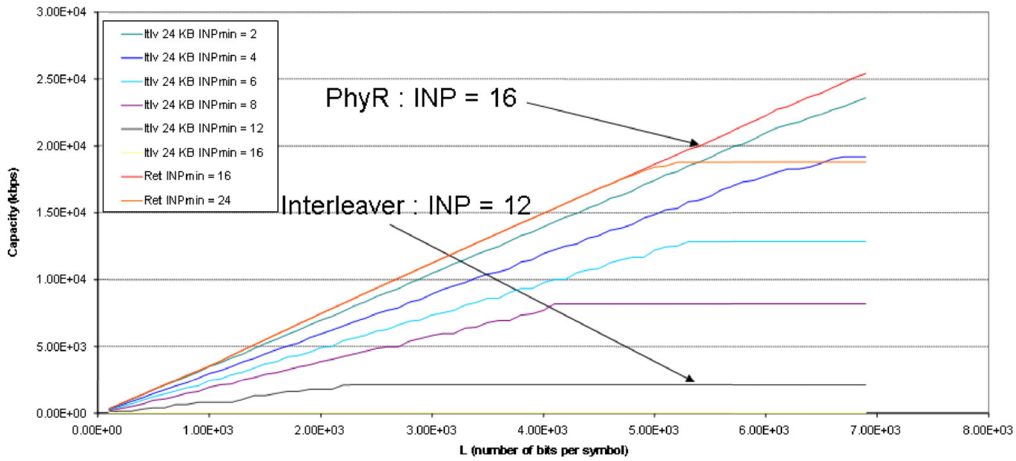

Figure 5: Capacity of PhyR Compared to RS + Interleaving

As retransmission is also complementary to RS encoding, the RS overhead can now be selected solely to optimize the coding gain. Thus, RS encoding becomes a benefit rather than an overhead.

PhyR—Much Higher INP than a Standard Interleaver

While providing higher rates, PhyR is also able to cope with a much wider variety of impulse types. Figure 6 highlights the correction capability of both interleaver and retransmission schemes for different impulse lengths and impulse inter-arrival time (IA, impulse periodicity), under the fair assumption that the overhead introduced to cope with the impulses is the same in both cases (RS overhead for the inter-leaver, bandwidth lost for retransmission on the other side).

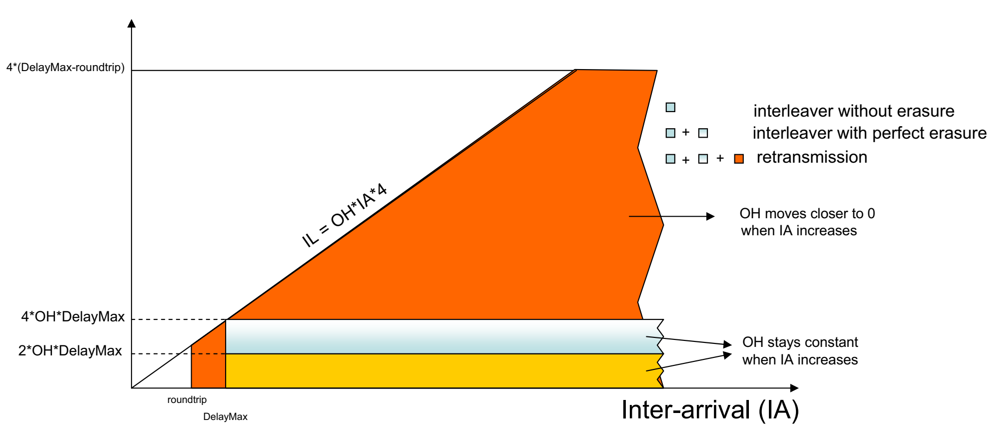

Figure 6: INP of PhyR Compared to RS + Interleaving Schemes

Figure 6 on page 8 demonstrates that retransmission copes with all the impulses that the interleaver is able to correct, even assuming ideal erasure decoding (not possible in practice), but it can also correct impulses that are totally beyond the capabilities of the interleaver. This approach is mainly limited by the maximum overhead allocated to retransmission in order to provide a minimum guaranteed payload rate.

PhyR—No Margin Cheating

Less obvious, but also essential, is that PhyR stops the "organized cheating" of overbooking coding gain and impulse protection.

Figure 7 shows the probability to correct an impulse at 0 dB margin in function of the overhead allowed for retransmission. Allowing for 1% of bandwidth lost by retransmission, it shows that 99.99% of the impulses are corrected at a 0 dB noise margin. This is much better than the interleaver at 3 dB noise margin!

PhyR-enabled modems, therefore, provide much less sensitivity to additional noise in low margin conditions and are the only ones capable of remaining in Showtime when Repetitive Impulse Noise (REIN) is applied during these conditions, thereby dramatically improving modem robustness against extreme stress in the field.

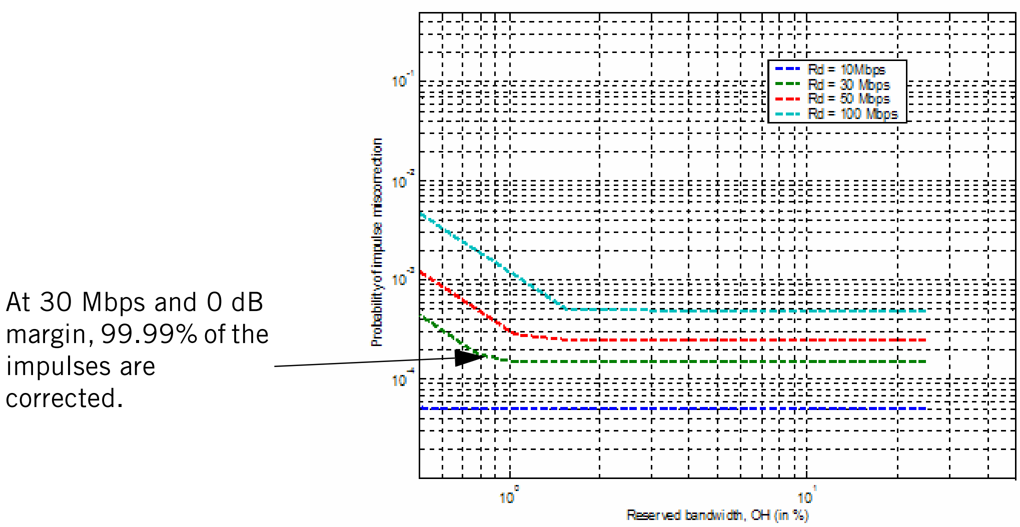

Figure 7: PhyR INP Protection at 0 dB Margin

PhyR Offers 10^{-10} BER

PhyR also provides an answer for the lower BER required by IPTV services. Indeed, retransmission even provides some coding gain!

Figure 8 compares the residual BER at a 0 dB margin of an interleaver scheme versus a retransmission scheme. While the interleaver results in one error every 10 seconds at 40 Mbps, retransmission results in only one error every 100 hours! Note as well that in such a case, the data rate lost because of retransmission is only about 6 Kbps—completely negligible in comparison to 40 Mbps.

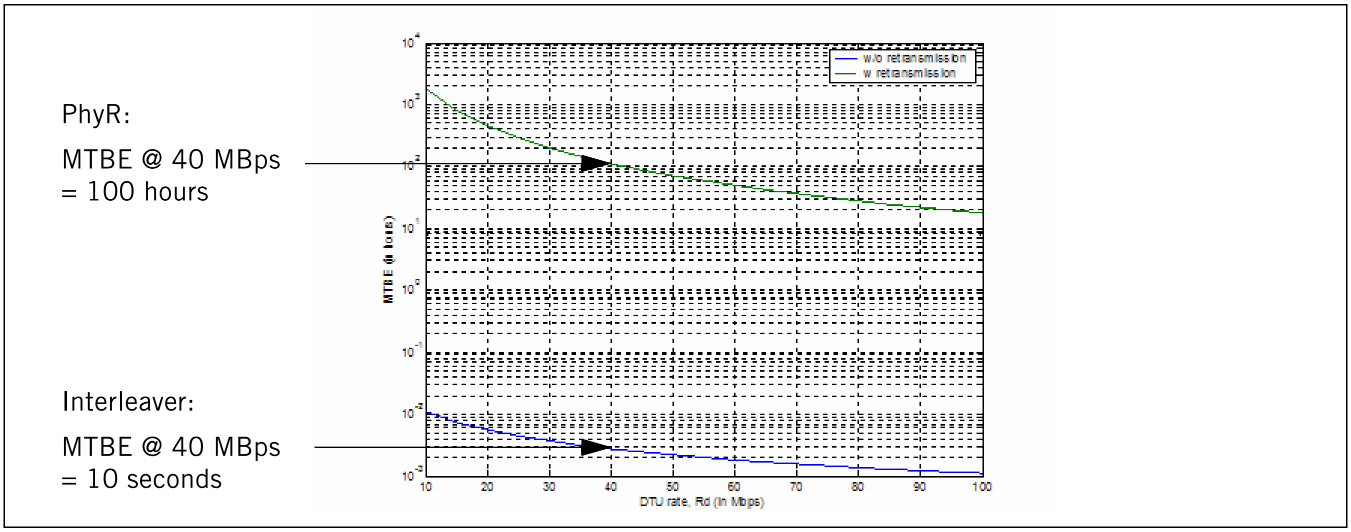

Figure 8: Mean Time Between Error (MTBE) at 0 dB Margin

PhyR— 'Set and Forget' Provisioning

A major benefit of having these advantages without any fixed penalty is that they can be blindly provisioned to everyone! There is no need to tune the configuration on a per line basis, so everyone can receive a low delay and high rate service, with only the few bad lines effectively spending some controlled part of the data rate bandwidth into retransmissions. This removes all need for the provisioning nightmare of the current implementation or even dual latency: gamers, VoIP, and IPTV requirements are met with a single, low-latency path.

Figure 9 on page 11 shows the relative delay of retransmission compared to interleaver: typical values used in the field today are 16 ms as maximum interleaving delay against typically 5 ms for retransmission.

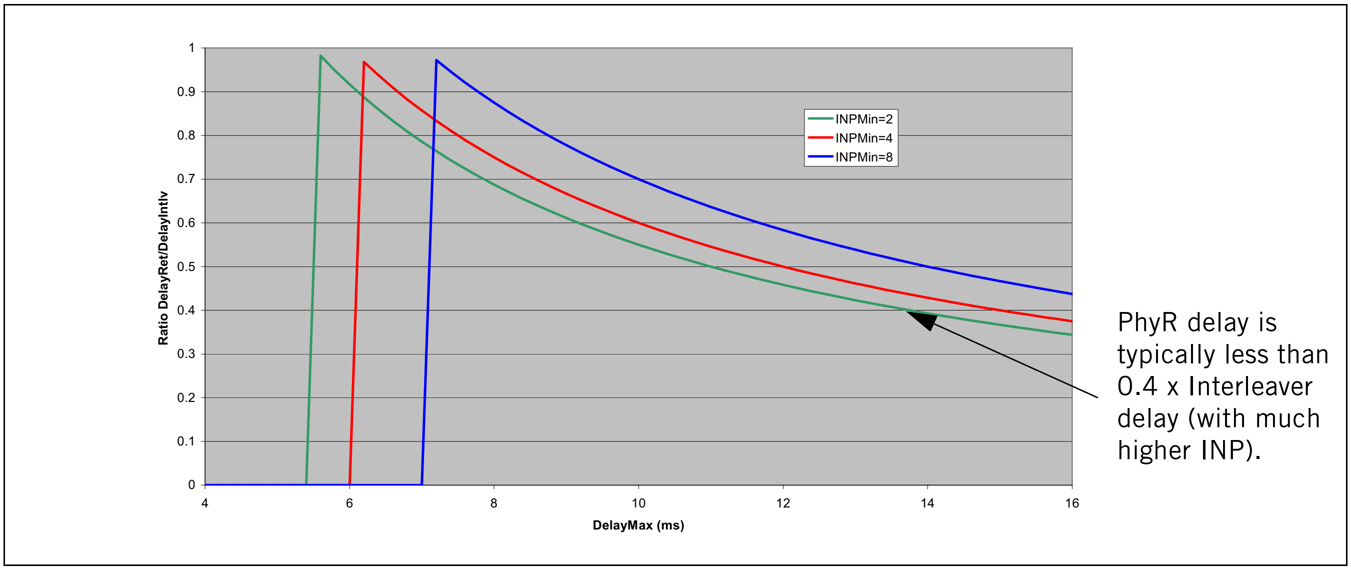

Figure 9: PhyR vs. Interleaver Delay

PhyR—DSL just as Ethernet or Fiber

Another key advantage of implementing a retransmission protocol at the physical layer is that no additional and dedicated complexity is required at upper layers in the system (e.g., on a network processor or the backbone network) to address the poorer BER characteristic of DSL. This makes DSL provisioning similar to other physical layers used in access provider network (e.g., Ethernet or GPON) and helps access providers' operation and planning.

PhyR is mostly transparent and significantly reduces the burden on a network that uses retransmission on the application layer (e.g., Microsoft® IPTV), thereby improving the efficiency and scalability of the network. In essence, physical and application layer retransmission schemes are complementary and work constructively together to increase triple-play service protection against packet losses.

PhyR—Some Field Results

Broadcom has run extensive laboratory tests to back up all PhyR claims presented in this document. These tests have demonstrated not only uninterrupted video streaming over a link corrupted by strong periodic impulses up to 40% of the time, but also error-free links in many operators' labs, using their reference worst case noise and loop configuration and adding the most severe impulse noise type they recorded in their network.

Beyond these lab results, we were also able to validate PhyR in the field with operators, thanks to collaboration with Free® (operating roughly two-million lines in France), who allowed us to deploy PhyR technology by upgrading all their CO and CPE equipment and gave us access to detailed modem statistics for analysis purposes. Some of these results are presented below.

Figure 10 shows the cumulative line distribution in function of the number of frame errors per hour. This was recorded at the same time and place on about 5000 lines with default interleaver configuration for one line on two and with PhyR enabled for the others.

It is obvious and remarkable that PhyR-enabled lines have consistently fewer CRCs by an order of magnitude (10x)!

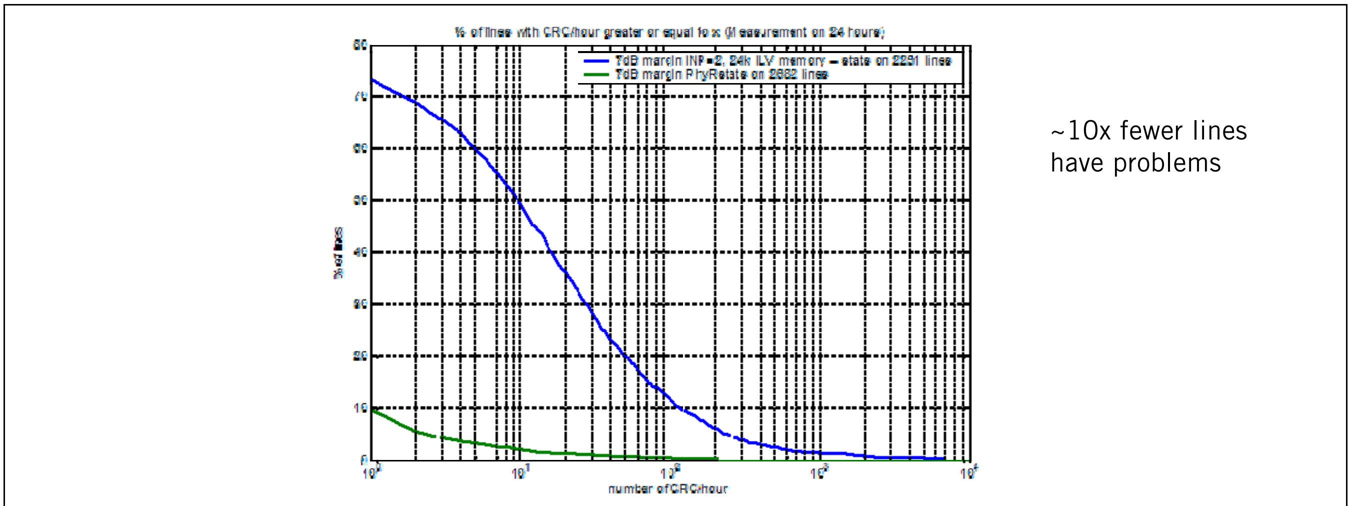

Figure 10: PhyR vs. Interleaver CRC

The comparison of the rates achieved on those same lines, as illustrated in Figure 11, also confirms our previous claims: this huge improvement in the robustness of the system is accompanied by a rate improvement!

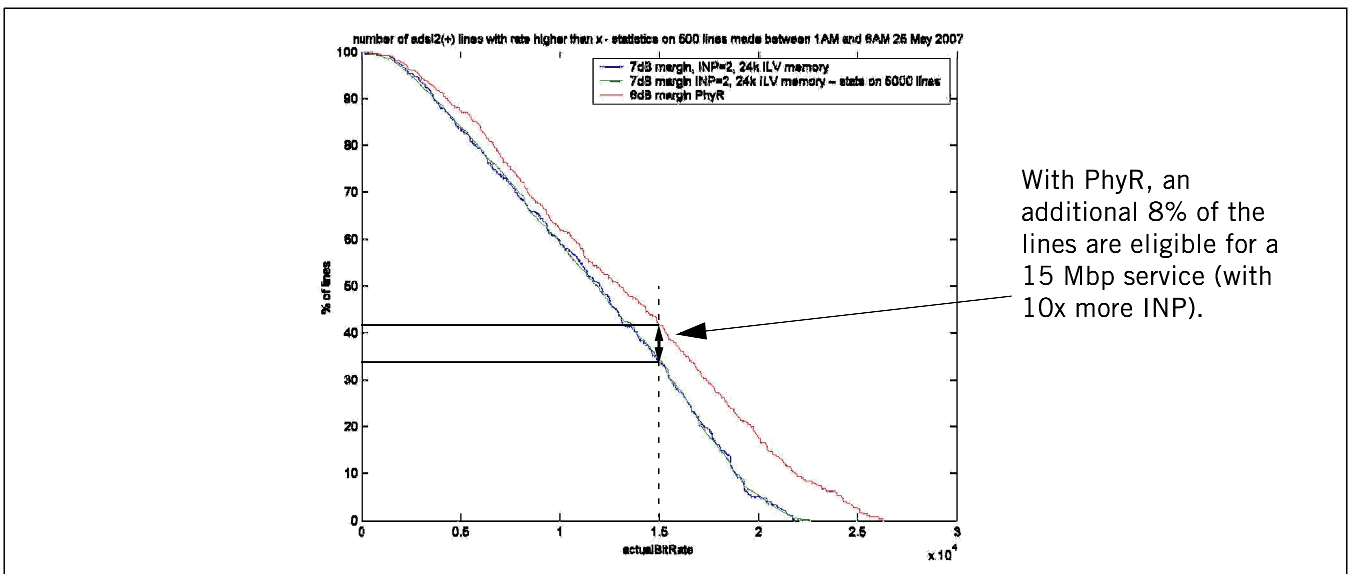

Figure 11: Service Reach Improvement

This rate improvement comes from, on the one hand, the lower margin used (this solely explains the improvement below 12 Mbps), and on the other hand, from the poor efficiency of the interleaver at high rates (above 12 Mbps). For example, Figure 11 on page 12 shows that 8% more customers would be eligible for a 15 Mbps service compared to a standard implementation (provisioned only with INP = 2, much less than similar lines running PhyR, see below), and gain is even higher at higher rates.

Figure 12 further demonstrates this advantage. Today's scheme cannot address impulse noise issues in the field above 2…4 (i.e., all lines with higher INP will create packet errors), while PhyR keeps up with INP > 100 for some lines with the current implementation.

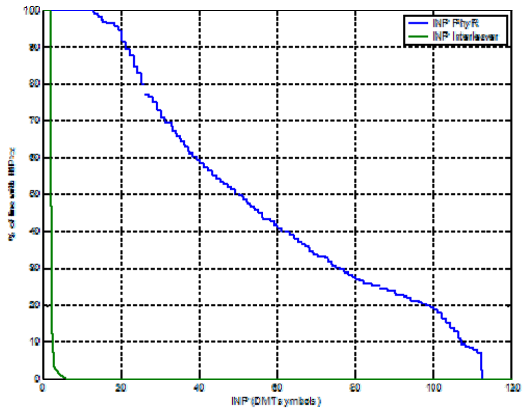

Figure 12: Rate vs. INP Settings

Figure 13 demonstrates, finally, that in almost all cases, the amount of bandwidth "stolen" to retransmit data remains statistically negligible.

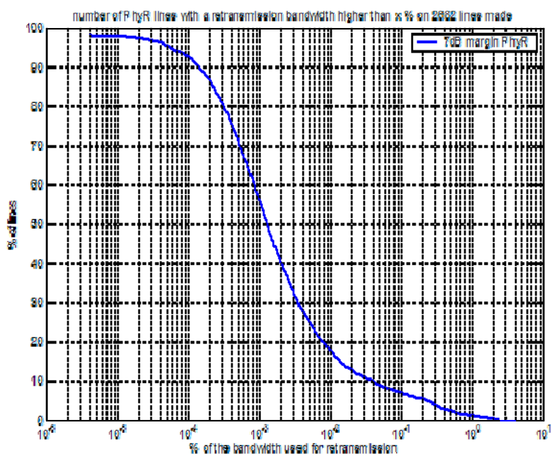

95% of the lines are using less than 0.1% of the bandwidth for retransmissions:

=> negligible!!

Figure 13: Bandwidth Used by Retransmission

Conclusions

Broadcom's PhyR retransmission scheme provides a number of key advantages:

- As much as a ten times (10x) higher impulse-noise resilience without impacting link delay and error propagations.
- Significantly lower residual BER or packet loss even in REIN conditions. This repositions DSL for the delivery of IPTV, negating the need for dedicated PHY-level protection schemes to achieve BER comparable to Ethernet or Fiber.
- Extended service reach: higher rate, longer reach, lower delay.
- Much simpler network provisioning ('set and forget', no per user tuning). Dual-latency and Dynamic Rate Repartioning (DRR) complexity are avoided by providing a fast path with very high INP capabilities.
- The scheme is transparent to both network and upper layer applications and significantly reduces the burden on networks using higher-layer retransmission schemes to improve network efficiency (Microsoft IPTV, TCP/IP).
- Easy to implement on any existing Broadcom-powered ADSL2+ or VDSL2 central office and customer premise equipment via a simple firmware upgrade (deployed easily through remote upgrade functions).
- Same silicon/memory cost as classical interleaving.
- Available today from Broadcom!

Connecting everything®

BROADCOM.

BROADCOM CORPORATION
5300 California Avenue
Irvine, CA 92617

Phone: 949-926-5000
Fax: 949-926-5203
E-mail: info@broadcom.com
Web: www.broadcom.com

XDSL-WP101-R April 2009