

Low-Power Modes for ADSL2 and ADSL2+

George Ginis

Broadband Communications Group, Texas Instruments

1. Introduction

This document gives a description of low-power modes for ADSL2 and ADSL2+. These include the power management states known as L2 and L3 and the power cutback feature. The ways in which these states can be used are explained, and their various restrictions are outlined.

It is observed that the L2 mode (and the L3 mode when coupled with fast initialization) offer power consumption savings when averaging over long periods of time, but they contain several risks mainly due to the creation of time-varying crosstalk.

The power cutback feature provides an alternative mechanism for reducing power consumption without introducing time-varying crosstalk. This is achieved by applying an amount of power cutback (selected either manually or automatically) at the beginning of a connection.

2. Description of the L2 and L3 power management modes

The ADSL2 [1] and ADSL2+ [2] standards define power management states as described in Table 1. The allowed transitions between these states are graphically shown in Figure 1. More detail on each of these states and their transitions is given in the following paragraphs.

Table 1: Power management states (from [1])

| State | Name | Description |
|-------|-----------|--|
| L0 | Full On | The ADSL link is fully functional. |
| L2 | Low Power | The ADSL link is active, but a low-power signal conveying background data is sent from the ATU-C to the ATU-R. A normal data carrying signal is transmitted from the ATU-R to the ATU-C. |
| L3 | Idle | There is no signal transmitted on the line. The ATU may be powered or unpowered in L3. |

Note:
 -ATU is ADSL Transceiver Unit.
 -ATU-C is ATU at the central office end.
 -ATU-R is ATU at the remote terminal end.

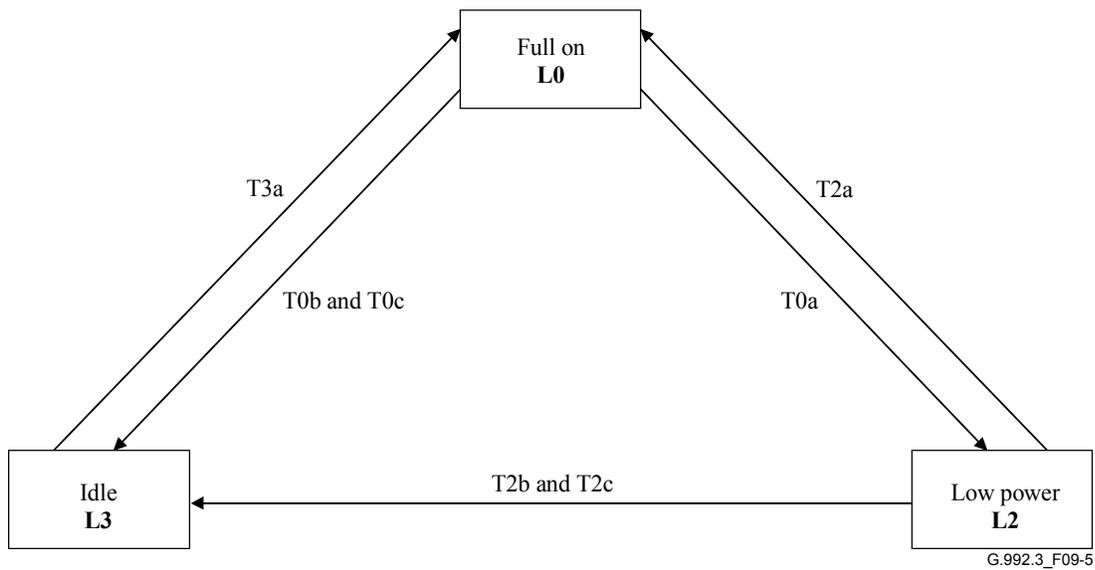


Figure 1: ADSL link power management states and transitions (from [1])

2.1. Full on L0 state

The L0 state is the mode of normal operation. The full data rate for a given environment can only be achieved when the ADSL link is in the L0 state.

2.2. Low-power L2 state

The L2 state is a low-power mode that is to be used only when there is little to no need for transmission of information from the ATU-C to the ATU-R. During this time, the downstream data rate can be reduced, which allows the average transmitted power to be scaled down. The effect of lowering the transmitted power is to potentially reduce the power consumed by the ADSL chipset, as well as lower its thermal dissipation.

Note that the power savings apply only for the ATU-C. The reason for this is that central office chipsets generally have tighter constraints on power consumption and thermal dissipation, especially when deployed in the outside plant (e.g. remote DSLAM).

Although the currently defined L2 state achieves power savings by reducing the average transmitted power, an alternative proposal on which there was no agreement in standards committees was to achieve similar power savings by reducing the peak-to-average ratio of the transmitted signals. For example, changing the constellation size of all tones to 4-QAM would give a reduction of the peak-to-average ratio. This leads to

power consumption savings without resorting to a reduction of the average transmitted power.

It is important to note that bit-swapping operations are not allowed during the L2 state. If a bit-swapping operation must be performed, then a transition to L0 state must first take place.

2.3. Idle L3 state

The L3 state is the mode when no signal is transmitted on the line, and thus no transmission of information is possible. When a link is in the L3 state, power savings are possible for the ATU-C and ATU-R chipset. The power savings depend on whether the corresponding chipsets (or ports of the chipset) are placed on a “monitor” state or on a “disabled” state, where the latter provides the larger savings.

2.4. Transitions from L0 to L2

If the ATU-C detects that there is little to no downstream traffic (e.g. by checking that the FIFO transmit buffers are empty or almost empty), then it may initiate a transition from L0 to L2. The ATU-C and the ATU-R coordinate the transition by exchanging a number of messages. The ATU-C sends to the ATU-R the allowed range for the downstream power cutback and for the downstream data rate. The ATU-R can either reject the request or select parameters within the ranges requested by the ATU-C. Note that the ATU-R response includes new bits and gains tables.

2.5. Transitions from L0 to L3

The transition from full on to idle can be initiated by either the ATU-C or the ATU-R. This transition can be coordinated between the two sides following the so-called “orderly shutdown procedure.” Also, there is provision for a “disorderly shutdown procedure,” which can be used when the ATU-R suddenly loses power.

2.6. Transitions from L2 to L0

This transition can be initiated by either the ATU-C or the ATU-R and involves a fast procedure to coordinate the two sides. The intention is to return to full on mode as soon as possible in order to resume normal transmission of information (e.g. when the transmit FIFO buffers of the ATU-C start filling up). Note that upon returning to L0, the ADSL link must operate with the same parameters as those in the previous L0 state.

2.7. Transitions from L2 to L3

This transition can be initiated by either the ATU-C or the ATU-R. Upon encountering loss of power, the ATU-R can initiate a “disorderly shutdown procedure.” Also, the ATU-C can initiate the transition using the “orderly shutdown procedure.”

2.8. Transitions from L3 to L0

The transition from L3 to L0 can be started by either the ATU-C or the ATU-R and of course must follow the initialization procedures defined in the ADSL2 and ADSL2+ standards.

2.9. L2 trim

While in L2 state, the ATU-C can initiate a “power-trim” procedure. This allows the ATU-C to further reduce the transmitted power by adjusting the amount of power cutback (for example, when the downstream reported margin is seen to be too high). The ATU-C and the ATU-R coordinate the trim operation by exchanging a number of messages. The ATU-C sends to the ATU-R the allowed range for the downstream power cutback and for the downstream data rate in the L2 state. The ATU-R can then either reject the request or select parameters within the ranges requested by the ATU-C.

3. Management control of L2 and L3 modes

This section provides information on how the low-power modes can be controlled through the management interface defined for ADSL2 and ADSL2+ in [3].

3.1. Power management state enabled (PMMode)

This configuration register defines the allowed line states for the ADSL link. Using this parameter, the L2 or L3 states can be independently enabled or disabled. Disabling a state means that a transition to that state is not allowed.

3.2. Power management state forced (PMSF)

This register allows the external forcing of a state on a line. The following transitions can be forced:

- a) Idle L3 to full on L0 state.
- b) Full on L0 to low-power L2 state. (Useful only for testing.)
- c) Full on L0 or low-power L2 to idle L3.

3.3. Minimum L0 time interval between L2 exit and next L2 entry (L0-TIME)

This parameter defines the minimum time allowed between an exit from the L2 state and the next entry into the L2 state. The maximum value is 255 seconds. It can be used to control how often the link can enter the L2 state.

3.4. Minimum L2 time interval between L2 entry and first L2 trim (L2-TIME)

This parameter defines the minimum time allowed between an entry into the L2 state and the first power trim in the L2 state and between consecutive L2 trims. The maximum value is 255 seconds. It can be used to control how often power trim operations can take place.

3.5. Maximum aggregate transmit power reduction per L2 request or L2 trim (L2-ATPR)

This register defines the maximum aggregate transmit power reduction that can be performed at transition from L0 to L2 state or during L2 trim. Its range is from 0dB to 31dB. It can be used to control the change in the aggregate transmit power during a transition to L2 or during an L2 trim.

3.6. Total maximum aggregate transmit power reduction in L2 (L2-ATPRT)

This parameter defines the total maximum aggregate transmit power reduction that is allowed in L2 state. It is a constraint on the sum of all reductions including the L0 to L2 transition and the L2 trims. It ranges from 0dB to 31dB. It can be used to control the total change in the aggregate transmit power between the L0 and the L2 states.

3.7. Illustration of the management parameters for line state

An example is given here to show how the transmitted PSD of a line may change over time for specific values of the parameters defined above. The example is graphically depicted in Figure 2.

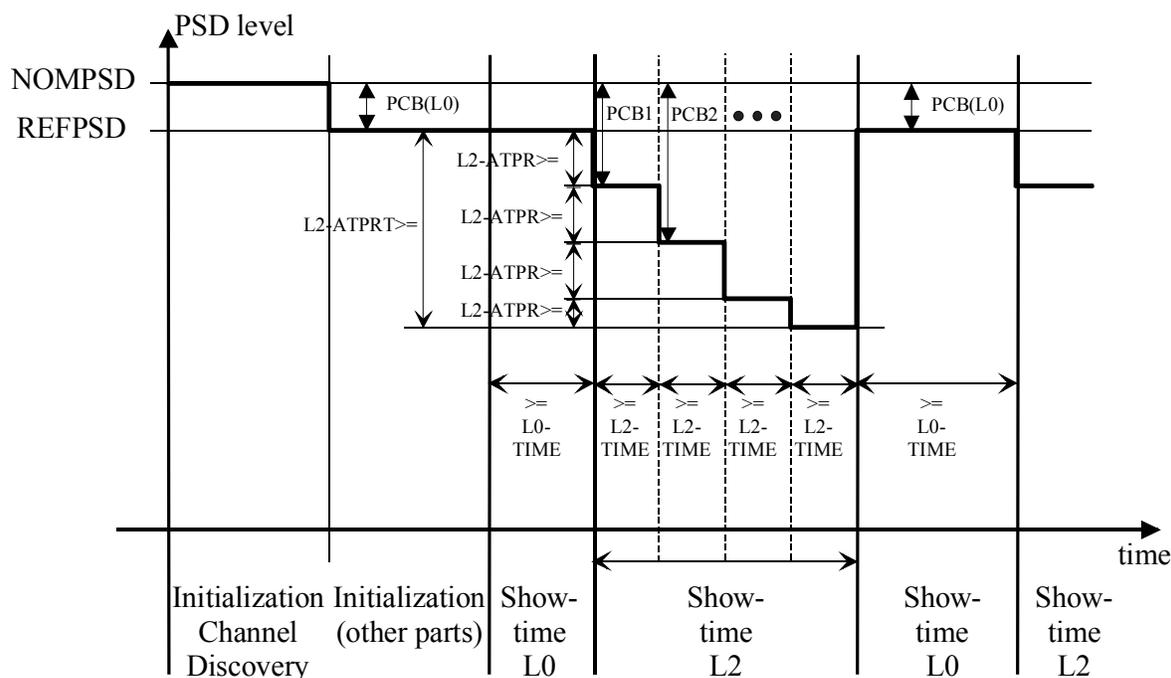


Figure 2: Illustration of management parameters for line state (from [1])

The horizontal axis depicts time, while the vertical axis depicts the PSD level transmitted on the line. The line starts initialization at time 0. During the “channel discovery” phase of initialization, the PSD level is equal to NOMPSD (nominal transmit PSD level). After this phase concludes, initialization continues with the PSD level equal

to REFPSD (reference transmit PSD level). REFPSD equals NOMPSD minus the amount of power cutback that is negotiated during “channel discovery” between the ATU-C and the ATU-R. This is here denoted as PCB(L0).

At the end of initialization, the link enters showtime (L0 state) with the PSD equal to REFPSD. Some time later, the link transitions to L2 state. Upon entering L2 state, the transmitted PSD changes by PCB1 relative to NOMPSD. The difference between PCB1 and PCB(L0) can be no larger than L2-ATPR. The next change in the transmitted PSD comes from an L2 trim. This can happen only after L2-TIME (or more) seconds have passed. After the L2 trim, the transmitted PSD is lower than NOMPSD by PCB2. Again, the difference between PCB2 and PCB1 can be no larger than L2-ATPR. Several L2 trim operations can follow, provided that the PSD adjustment is no larger than L2-ATPR and that they do not occur within less than L2-TIME seconds of each other. Also, the difference between the PSD level and the REFPSD must never become larger than L2-ATPRT. When the link exits from L2 to L0, the PSD level changes back to REFPSD. The link must remain in L0 for at least L0-TIME seconds before re-entering L2.

4. Potential issues with L2 and L3 modes

The primary application for the low-power modes is to reduce the power consumption and thermal dissipation of central office ADSL chipsets. When modem users become inactive, then the ADSL links can transition from L0 to L2 state and thus reduce their average transmitted power. If inactivity persists for a very long time, then the ADSL link may even transition to the L3 state. When data download is again requested, then the transition from L2 to L0 must take place very fast so that it is barely perceivable for the end-user. The transition from L3 to L0 needs to take more time. (Even the optional fast initialization specifies that the transition can take as many as three seconds.) When averaging over long periods of time (in the order of days), it is expected that the adoption of low-power modes may have a reduction effect on the power consumption and on the thermal dissipation. (The exact amount depends on the implementation.)

On the other hand, low-power modes create a number of complications both at the link and at the network level. These potential issues are listed next.

4.1. Robustness of fast initialization

The use of the L3 mode for the purpose of power savings seems to make sense only when combined with the optional fast initialization procedure, so that the transition time from L3 to L0 can be minimized.

However, fast initialization is for now an unproven concept. Shortening the initialization time from more than 20 seconds down to three seconds requires a number of initialization shortcuts which can affect the robustness of the connection. The following shortcuts are listed in Table 8-41 of [1]:

- a) No hybrid fine-tuning for ATU-R or ATU-C.
- b) Faster upstream channel estimation, less precise timing.
- c) Faster downstream channel estimation and equalizer training.
- d) Less accurate SNR estimation for downstream and upstream.
- e) Faster and simpler bit allocation algorithms for downstream and upstream.

These shortcuts are thought to have small impact if the line conditions remain the same as during the previous L0 state. However, if the line conditions change while the link is in L3 state, then it is likely that fast initialization will not achieve a robust connection and that a full initialization will be required.

4.2. Robustness of L2 state

It is important to note that bit-swapping operations cannot be performed while the link is in the L2 state. If a bit-swap operation must take place, then a transition to L0 must take place. This has two implications:

- a) Transitions between the L0 and L2 states can take place not only due to traffic demands but also due to the need to maintain the connection through bit-swapping. The next subsection explains that such transitions can have an adverse effect on the network due to the introduction of time-varying crosstalk.
- b) It becomes more difficult to recover from abrupt changes in the line conditions, since the modems must first exit from L2 mode to L0 mode and then start performing bit-swapping operations to restore the connection margin.

4.3. Time-varying crosstalk

The most serious concern about low-power modes as defined in ADSL2 and ADSL2+ comes from the fact that they can introduce time-varying crosstalk in the DSL network, which may in the worst-case lead to network instability.

This effect can be illustrated by a simple example. Assume that link L is initialized on a line that neighbors other lines with links that are in L2 or L3 state. During initialization, link L would experience low levels of noise and thus connect at a high rate. If sometime later one or more of the neighboring lines exit from L2 (or L3) state back to L0 state, then their transmitted power will suddenly change by as much as tens of dB. This can result in a large increase of the crosstalk experienced by link L. This change will be very abrupt, and it is likely that link L will not be able to recover by performing bit-swapping operations, in which case it will have to drop the connection and reinitialize.

One could note that such transitions from L2/L3 to L0 states have an impact that is comparable to the impact of switching a modem from off to on. However, the difference lies in how often such events take place. L2/L3 to L0 transitions depend on user traffic and are expected to happen much more frequently than modem turning on events. Additionally, L2 to L0 transitions can be triggered by the need to perform bit-swapping, which further increases the occurrence of time-varying crosstalk in the network. Such crosstalk variation occurring with a relatively high frequency over the whole DSL network is a cause for concern.

4.3.1. Time-varying crosstalk and L2 mode for VDSL2

The issue of time-varying crosstalk resulting from L2 mode has recently attracted the attention of DSL standards committees during the discussions for VDSL2. The short lines, for which VDSL2 is targeted, suffer from strong Far-End-Crosstalk (FEXT). In this case, the impact of time-varying crosstalk is expected to be very significant.

For this reason, it was proposed to completely abandon L2 mode for VDSL2 [4,5,6,7,8,9]. This proposal gained the support of a number of service operators, including SBC, Bell Canada, BT, Qwest, France Telecom, Telecom Italia and TeliaSonera. The ITU Study Group 15 agreed in the October 2004 meeting that an L2 mode shall not be specified for VDSL2 [10].

4.3.2. Time-varying crosstalk and L2/L3 mode for ADSL2/ADSL2+ mixed deployment

The issue of time-varying crosstalk for ADSL2/ADSL2+ systems becomes very serious in the cases of so-called “mixed” deployment scenarios, where services sharing the same cable-bundle are deployed from both a central office and a remote terminal. This subsection illustrates this.

The mixed loop topology of Figure 3 is assumed. ATU-C/R 1 are modems corresponding to ADSL2 services deployed from a Central Office (CO). ATU-C/R 2 are modems corresponding to ADSL2 or ADSL2+ services deployed from a Remote Terminal (RT). At first, this case may be considered ideal for the application of L2 mode for the downstream transmission from the RT. The RT has stricter power and thermal limitations compared to the CO, so statistical power savings would be very beneficial for the system.

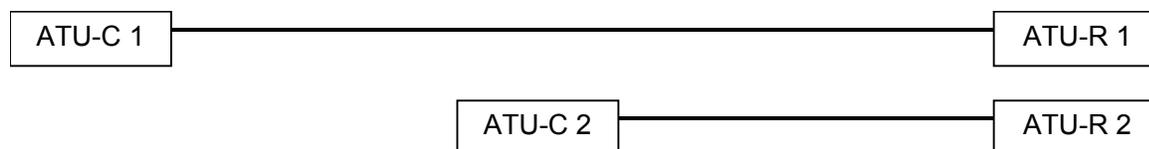


Figure 3: Mixed CO/RT loop topology

Next, the impact on the downstream data rates of CO-deployed ADSL2 is examined through simulation for various levels of power cutback applied at the RT downstream transmitters (presumably as in L2/L3 modes). The RT-deployed services are assumed to be ADSL2 (but results would be practically identical for ADSL2+ also).

The 99 percent worst-case crosstalk models are used in all cases. The rest of the simulation parameters are listed in Table 2.

Table 2: Simulation parameters

| Simulation parameter | Value |
|------------------------------|------------|
| SNR gap | 9.8dB |
| SNR margin | 6dB |
| Coding gain | 4dB |
| Echo attenuation | 80dB |
| Receiver noise floor | -130dBm/Hz |
| AWGN noise | -140dBm/Hz |
| CO-deployed ADSL2 disturbers | 12 |
| RT-deployed ADSL2 disturbers | 12 |

Figure 4, Figure 5 and Figure 6 show the downstream CO data rates versus the loop length. Each figure corresponds to a different loop length for the RT deployed services. In each figure, different curves correspond to different levels of power cutback. The power cutback value of 200dB is not practical for L2, but it is included to compare with the case of the modem turning off completely (as in L3 state).

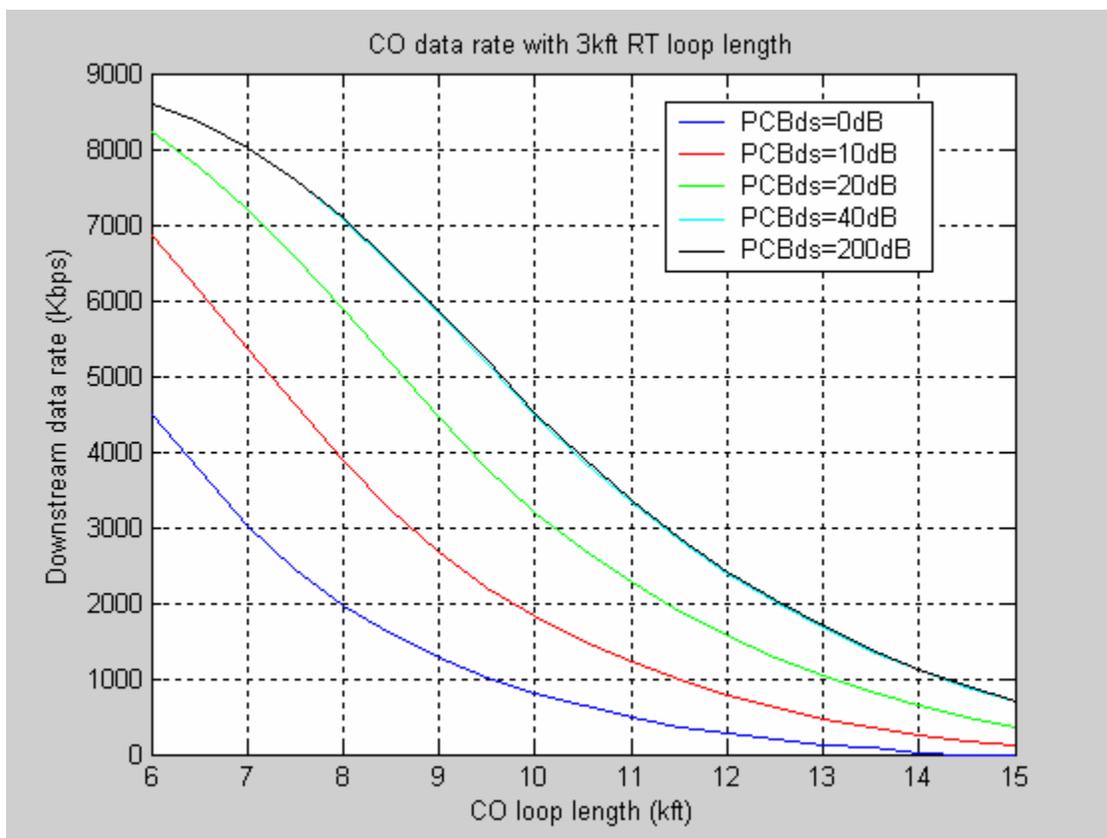


Figure 4: Downstream CO data rate with RT loop equal to 3kft

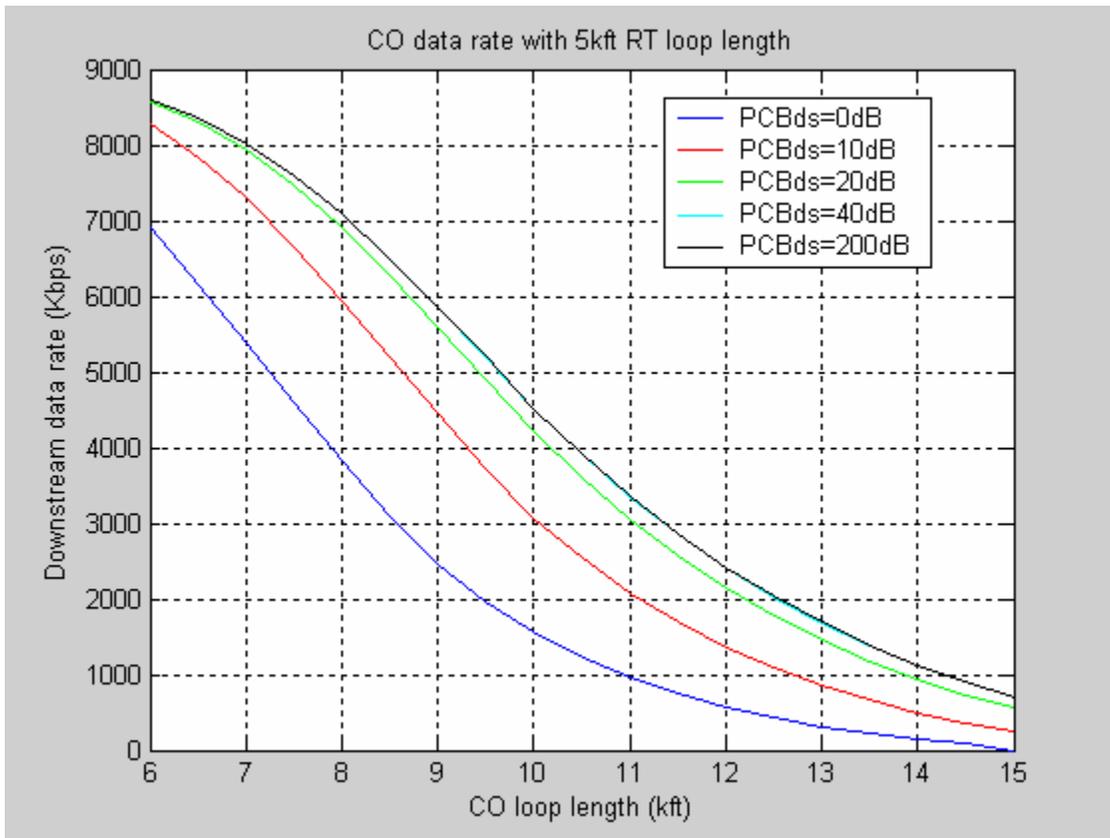


Figure 5: Downstream CO data rate with RT loop equal to 5kft

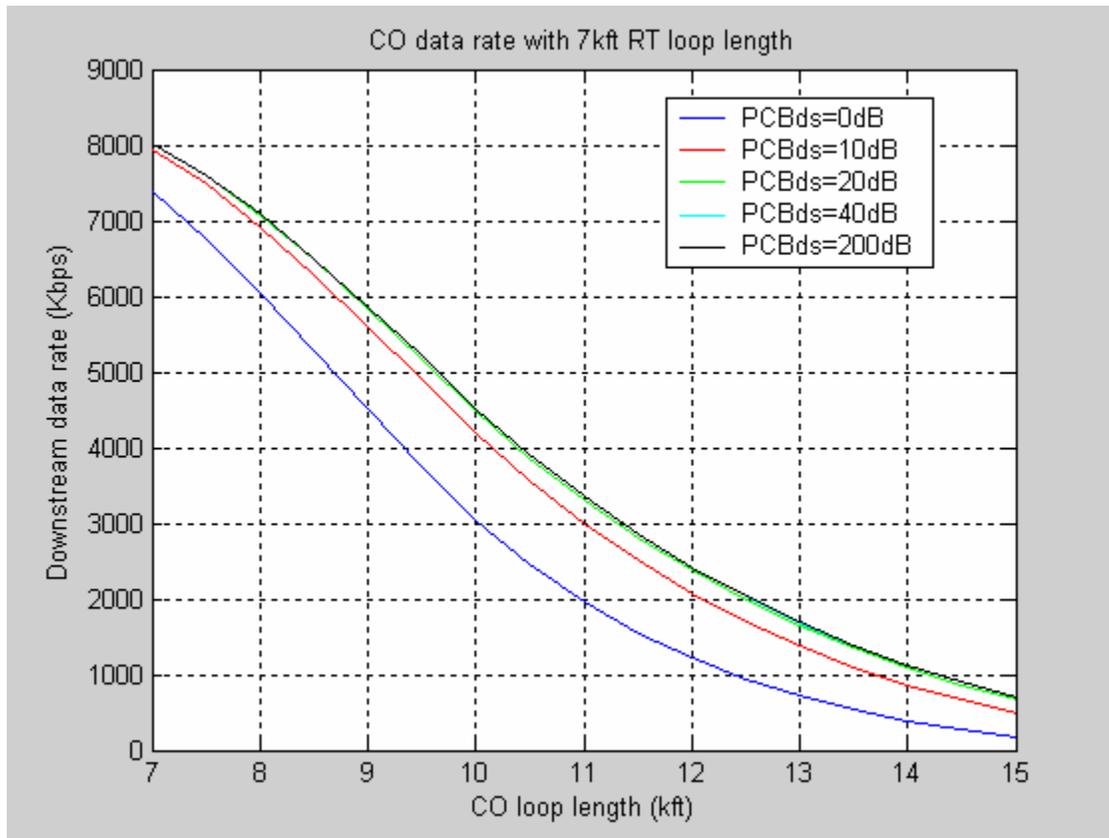


Figure 6: Downstream CO data rate with RT loop equal to 7kft

It can be seen in the above figures that the impact from L2/L3 states used in RT systems with short loops is potentially significant. With a 3kft RT loop, Figure 4 shows that the maximum downstream data rate of a modem may have a large variation depending on the amount of power cutback that is applied at the RT. If the link is initialized with neighboring lines in L2 or L3 states, then the data rate can be twice as large as the data rate with neighboring lines in the L0 state. This proves that in the former case, the link will likely need to reinitialize once one or more of the neighboring lines return to the L0 state.

5. Alternative solutions for power savings: power cutback

The time-varying crosstalk induced by transitions between L2/L3 and L0 states may be undesirable for service providers. An alternative is to achieve power consumption and thermal dissipation savings through the use of power cutback as specified in the ADSL2 and ADSL2+ standards.

Looking at Figure 2, it is seen that power cutback equal to PCB(L0) is applied after the “channel discovery” phase of initialization. Assuming that transitions to L2 state (or L3 state for power savings) are not allowed, the transmitted PSD level during showtime

remains constant and corresponds to a cutback of PCB(L0). This allows for power savings without causing frequent crosstalk variation.

The amount of power cutback can either be forced externally through the management interface or automatically negotiated during initialization between the ATU-C and the ATU-R. This is explained next in more detail.

In some cases, the service provider may prefer to manually set the value of power cutback. For example, the service provider may have performed an analysis through DELT and determined the best transmitted PSD level for the given link. This programming can be made independently for the downstream and the upstream transmission directions.

In most cases, it is desirable to automatically determine the required amount of power cutback. The ATU-C and the ATU-R will then each produce estimates for the downstream and the upstream power cutback amounts. These are:

- a) Downstream power cutback requested by the ATU-C (PCB_DS_ATUC)
- b) Downstream power cutback requested by the ATU-R (PCB_DS_ATUR)
- c) Upstream power cutback requested by the ATU-C (PCB_US_ATUC)
- d) Upstream power cutback requested by the ATU-R (PCB_US_ATUR)

These values are exchanged during the “channel discovery” phase of initialization, and the value of power cutback to be applied is determined as:

- a) Downstream power cutback is largest of PCB_DS_ATUC and PCB_DS_ATUR
- b) Upstream power cutback is largest of PCB_US_ATUC and PCB_US_ATUR

The method employed by each ATU to produce estimates of the power cutback amount is implementation-dependent. However, the intention is to produce values such that the following requirements can be met regarding the SNR margin achieved during showtime:

- a) The margin must be higher than the minimum margin (see MINSNRMds and MINSNRMus in [3]).
- b) The margin should be lower than the maximum margin (see MAXSNRMds and MAXSNRMus in [3]).

The minimum margin requirement is introduced to guarantee a minimum level of service. (When the margin falls below the minimum, the link must reinitialize.)

The maximum margin requirement is introduced to prevent the margin from being unreasonably high. Such large margins may provide little additional robustness in practice compared to moderate levels of margin. Additionally, such large margins imply that the transmitted power is too high for the given connection. Thus, setting the maximum margin indirectly controls the transmitted power and consequently leads to power savings.

Certain constraints of the power cutback mechanism need to be mentioned:

First, the power cutback cannot exceed the value of 40dB. This is not always sufficient to guarantee any configured value for the maximum margin. Even with 40dB of cutback, there may still be some short loops with very low fixed rates that can achieve a margin that is higher than the configured margin. However, it should be noted that the additional power savings from cutting back power beyond 40dB are negligible in the vast majority of cases. Similarly, the crosstalk reduction achieved by power cutback beyond 40dB is insignificant in most cases (see also Figure 4, Figure 5 and Figure 6).

Second, choosing the best power cutback amount may require more than one training attempt. This is because the power cutback value exchange between ATU-C and ATU-R needs to be performed early during initialization, at which time SNR estimates are rather imprecise. Reliable SNR estimates are only available at the end of initialization, when changes in the transmitted PSD are no longer allowed. Thus, the optimization of the value of power cutback may often require a retrain. In the second training attempt, the ATU can utilize the SNR estimates from the previous training attempt and produce much better power cutback estimates.

6. Conclusion

This paper described the low-power modes for ADSL2 and ADSL2+. An overview of the operation and management of these modes was given.

The original intention of defining the L2 state (and the L3 state coupled with fast initialization) was to achieve power consumption and thermal dissipation savings when averaged over a long period of time. However, a number of concerns has been voiced about the use of these states, with the most serious being the impact of time-varying crosstalk.

An alternative approach to achieving transmitted power savings is to utilize the power cutback feature of ADSL2 and ADSL2+. This can be done either in a manual fashion (by directly programming the amount of power cutback), or automatically (by indirectly controlling the transmitted power through the maximum margin parameter).

7. References

- [1] *ITU-T Recommendation G.992.3*, Asymmetric Digital Subscriber Line (ADSL) Transceivers –2 (ADSL2).
- [2] *ITU-T Recommendation G.992.5*, Asymmetric Digital Subscriber Line (ADSL) Transceivers-Extended Bandwidth ADSL2 (ADSL2plus).
- [3] *ITU-T Recommendation G.997.1*, Physical Layer Management for Digital Subscriber Line (DSL) Transceivers.

- [4] Bell Canada, Ikanos, “G.vdsl: VDSL2: Considerations regarding L2 mode,” *ITU/SG15 Temporary Document LC-087*, Lake Tahoe, California, August 2-6, 2004.
- [5] SBC, BT, Qwest, France Telecom, ADTRAN, 2Wire, Ciena, Ikanos, “G.vdsl:VDSL2: Moving the standard forward – Bundled compromise agreements regarding interleaver and L2 mode,” *ITU/SG15 Temporary Document LC-099*, Lake Tahoe, California, August 2-6, 2004.
- [6] Ikanos, “Reprint of G.vdsl:VDSL2: Considerations regarding L2 mode,” *T1E1.4 contribution 2003/550*, Washington, DC, August 9-12, 2004.
- [7] SBC, BellSouth, Qwest, Sprint, “G.vdsl: VDSL2 application,” *ITU/SG15 Temporary Document SI-025*, Stresa, Italy, October 18-22, 2004.
- [8] 2Wire, Bell Canada, BellSouth, BT, France Telecom, Ikanos Communications, Qwest, SBC, Telecom Italia, TeliaSonera, “VDSL2: Proposal to Modify L2 Mode Agreement in VDSL2,” *ITU/SG15 Temporary Document SI-043*, Stresa, Italy, October 18-22, 2004.
- [9] Alcatel, “G.vdsl: VDSL2: Alcatel’s Requirements for VDSL2,” *ITU/SG15 Temporary Document SI-047*, Stresa, Italy, October 18-22, 2004.
- [10] Associate Rapporteur, “G.vdsl: Updated Issues List,” *ITU/SG15 Temporary Document SI-U11R2*, Stresa, Italy, October 18-22, 2004.

© 2005 Texas Instruments Incorporated

Important Notice: The products and services of Texas Instruments Incorporated and its subsidiaries described herein are sold subject to TI’s standard terms and conditions of sale. Customers are advised to obtain the most current and complete information about TI products and services before placing orders. TI assumes no liability for applications assistance, customer’s applications or product designs, software performance, or infringement of patents. The publication of information regarding any other company’s products or services does not constitute TI’s approval, warranty or endorsement thereof.