♦ Aspects of Dynamic Spectrum Management Level 3

Charles Storry, Miroslav Živković, Jan Verlinden, and Adriaan J. de Lind van Wijngaarden

Digital subscriber line (DSL) networks are currently enhanced to provide higher transmission rates and reliability in order to support emerging triple-play services, i.e., voice, video, and data. This is typically achieved by installing digital subscriber line access multiplexers (DSLAMs) in street cabinets to significantly reduce the loop length to the subscribers. In addition, such DSLAMs typically employ DSL technologies that can use a wider spectrum for data transmission and thus effectively provide much higher transmission rates for short distances. The DSL performance is typically limited by crosstalk from the other lines in the bundle. Crosstalk can be mitigated by using spectrum coordination or signal coordination techniques, commonly referred to as dynamic spectrum management (DSM). This paper focuses on the concept and potential gains of signal coordination, known as DSM level 3, as well as the standardization efforts and considerations for future development and deployment of DSM level 3. © 2008 Alcatel-Lucent.

Introduction

Access networks are currently transitioning from best-effort high-speed Internet to high quality Internet Protocol (IP) networks capable of delivering a wide range of services, in particular triple-play services, i.e., voice, video, and data. Existing broadband copper networks typically support high-speed Internet access (HSIA) services using asymmetric digital subscriber line (ADSL) technology with a typical bit rate of 1.5 Mbps. Basic triple-play solutions are likely to require up to 18 Mbps, and enhanced triple-play, supporting multiple high definition television (HDTV) streams, could benefit from 30 Mbps or more of available bandwidth per household. It is obvious that the current ADSL-based networks do not deliver enough bandwidth for the anticipated future needs, in particular for video and IP television (IPTV).

To address this, many operators are migrating toward very high speed digital subscriber line (VDSL) technology, which is able to make use of a wider bandwidth and thus able to provide significantly higher rates over short distances. To reduce the distance to the subscribers, VDSL digital subscriber line access multiplexers (DSLAMs) are used, which are typically strategically co-located with existing neighborhood cross-connect boxes that are within a 500 meter to 1500 meter range of the subscriber's home. Optical fiber is typically used to connect the DSLAMs to the central office (CO), thus effectively creating a fiber to the node (FTTN) network.

While shorter loops offer a significant increase in capacity across the VDSL spectrum, practical deployments will be constrained by crosstalk when

Bell Labs Technical Journal 13(1), 117–128 (2008) © 2008 Alcatel-Lucent. Published by Wiley Periodicals, Inc.
Published online in Wiley InterScience (www.interscience.wiley.com) • DOI: 10.1002/bltj.20286

multiple VDSL users are located within the same cable binder. This could limit the effectiveness of VDSL to address future, more bandwidth-intensive services. Crosstalk occurs when multiple lines co-located in the same group of cables or a binder use signals with overlapping spectrum. A number of technologies known under a common name, dynamic spectrum management (DSM), try to address the issue of reducing interference within digital subscriber line (DSL) binders. While static spectrum management techniques [2] specify static limits per DSL technology used, DSM tries to optimize rate/reach for current conditions in the binder. DSM optimizes performance using increasing levels of coordination (known as level 0, 1, 2, and 3, as outlined in [3]).

This paper introduces the concept of crosstalk mitigation and discusses innovative approaches to increase the usable capacity of telephone companies' copper broadband networks. In particular, signal coordination techniques known as DSM level 3 will be introduced and several possible implementation alternatives will be discussed. The research and standardization efforts that are needed to develop and implement an interoperable, commercially viable solution are identified and addressed.

Panel 1. Abbreviations, Acronyms, and Terms

ADF—Automated distribution frame
ADSL—Asymmetric digital subscriber line
ANSI—American National Standards Institute
ATIS—Alliance for Telecommunications Industry Solutions
CO—Central office
CPE—Customer premises equipment
DA—Distribution area
DSL—Digital subscriber line
DSLAM—Digital subscriber line access multiplexer
DSM—Dynamic spectrum management
FTTC—Fiber to the curb
FTTN—Fiber to the node
HDTV—High-definition television
HSIA—High speed Internet access
IP—Internet Protocol
IPTV—Internet Protocol television
ITU—International Telecommunication Union
ITU-T—ITU Telecommunication Standardization Sector
MIMO—Multiple input, multiple output
NAI—Network Access Interfaces
NIPP—Network Interface Power and Protection
OPEX—Operational expenditure
OSP—Outside plant
VDSL—Very high speed digital subscriber line

Multi-Service Networks

The concept of a multi-service network over a converged infrastructure is not new, but only nowadays are such deployments becoming practical. This has been facilitated by a general reduction in the cost per bit throughout the network, largely due to the adoption of Ethernet technology, as well as a reduction in the bandwidth required to deliver video streams, driven by advanced video compression techniques. Nevertheless, the aggregate bit rate per household continues to rise, and therefore new techniques need to be used to deliver more bandwidth over existing copper networks.

There are several options available to DSL access network operators to increase the data rates:

- Decrease the signal attenuation by moving DSLAMs to the outside plant (OSP) or neighborhood (i.e., FTTN), thus reducing the average loop length and increasing the average rate per customer.
- Increase the number of lines to facilitate multiple pairs per subscriber (e.g., pair bonding [4]).
- Improve technology by adopting newer generation DSLs, e.g., VDSL2 [5], to enable performance enhancing features.
- Increase the signal power by selecting higher power variations of VDSL2 (e.g., VDSL2 profile 8a or 8b), which may be useful for loops longer than approximately 1500 meters where the main limitations are due to the system's background noise, instead of crosstalk noise.
- Increase the spectrum used by DSL technology. For those subscribers whose loops are sufficiently short, rates can be improved through the use of additional spectrum (e.g., VDSL2 profile 17a).
- Decrease noise and crosstalk. DSL performance for shorter loops will be bounded by crosstalk

from other users within the same cable pair groupings or binders. By deploying DSM level 3, signal coordination is used to mitigate crosstalk as discussed in this paper, leading to increased rates for all DSL users in a binder.

These techniques may be employed individually, or in some cases together, to further increase the commercial lifetime of existing copper networks.

Crosstalk Mitigation Using DSM Level 3

DSM level 3 aims to mitigate crosstalk by jointly processing the actual signals of multiple lines in a binder. This requires all transmitters and/or receivers to be co-located. In DSM level 3, the binder effectively becomes a multi-user communication system. DSM level 3 is also often referred to as "vectoring" as the signals of all lines are combined in a vectored signal and jointly processed. When multiple lines are used for an individual subscriber, as in a line-bonding arrangement, DSM level 3 may also include the application of multiple input, multiple output (MIMO) principles [7]. In order to reduce the effects of crosstalk, the system should determine the characteristics of the interaction between wires within the binder, i.e., the "crosstalk channel." The process of determining the characteristics of the crosstalk channel is known as channel estimation. Channel estimation is a complex process which requires feedback from the receivers in order to perform the process effectively.

When all the receivers in a binder are co-located, the channel estimates can be used together with the received information from each line to cancel the effects of crosstalk. If instead, all the transmitters are co-located, then the channel estimates can be used together with the information that is being transmitted on all the active lines to precompensate (or precode) the transmit signal so that the signal appears without crosstalk at each respective receiver. In a typical DSL deployment, all downstream transmitters and upstream receivers are co-located in the DSLAM, and in this situation the majority of the changes to implement DSM level 3 are expected to be in or co-located with the DSLAM. An example of such a DSM level 3 topology is shown in **Figure 1**. A potentially huge legacy base of VDSL2 customer premises equipment

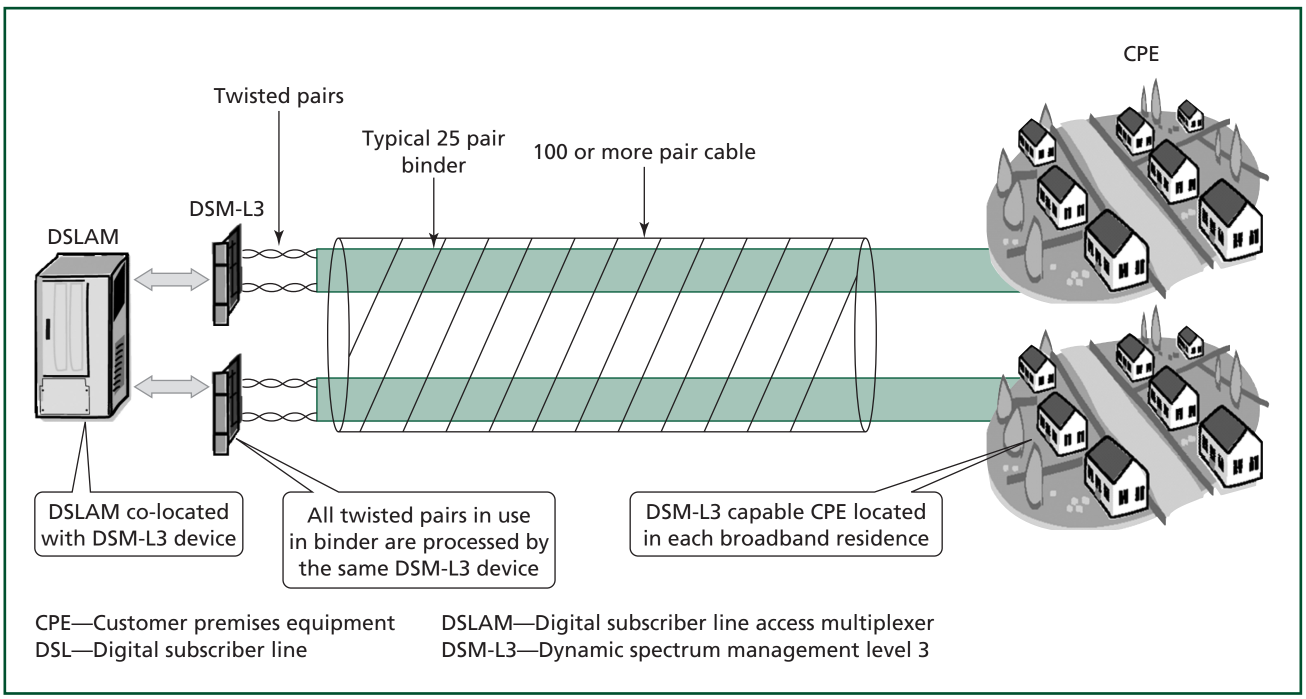

Figure 1.
DSM level 3 topology.

(CPE) is likely to exist at the time of DSM level 3 introduction. Support for channel estimation algorithms via an upgrade to existing CPE is therefore highly desirable. In order to migrate networks toward DSM level 3 support, the channel estimation algorithms should, in any event, be able to estimate the crosstalk from adjacent lines whose CPE is not DSM level 3 enabled. In this way, at least the upgraded lines can obtain the full benefit of crosstalk cancellation. Our research efforts concentrate on the development and refinement of such algorithms. Standardization discussions with respect to the channel estimation algorithms are ongoing.

DSM Level 3 Standardization Progress

The Network Interface Power and Protection–Network Access Interfaces (NIPP-NAI) committee of the Alliance for Telecommunications Industry Solutions (ATIS) has pre-published a technical report on DSM [3]. According to ATIS terminology, such a technical report is strictly informative and does not articulate mandatory capabilities and procedures. However, nearly all elements relevant for the implementation of DSM level 3 are left for further study.

In related work, the International Telecommunication Union Telecommunications Standardization Sector (ITU-T) Study Group 15 Question 4 has chartered a project called G.vector to identify the necessary procedures and modifications to existing DSL recommendations to implement an interoperable DSM level 3 solution.

To assist DSM level 3 standardization, these two organizations are currently studying:

1. Crosstalk channel modeling and performance characterization, and
2. Channel estimation and tracking.

A better understanding of the crosstalk channel is required to estimate the expected gains from crosstalk mitigation, to evaluate implementation options and trade-offs, and to more reliably assess channel estimation algorithms. A commonly used model [1] was created for simulating DSL performance in the presence of crosstalk. Its primary goal is to ensure that under most cases the majority of the effects of crosstalk are considered. This model is often referred to as the 99 percent crosstalk channel model. While this model overestimates the crosstalk, it is also likely to overestimate the potential gains of precompensation. This is not to say that the performance in the absence of crosstalk or following crosstalk mitigation will be lower than for the original (non-precoded) system, but rather that the typical performance may be better than the 99 percent worst case crosstalk model predicts. A more statistically valid crosstalk channel model is required to evaluate performance for different scenarios such as the cancellation of a reduced number of lines in a binder or a reduced number of tones, which may be considered to decrease implementation complexity and cost.

Channel estimation needs to be facilitated through standardization, as the line signals are being modified according to these estimates. The application of a correction factor that results from incorrect channel estimation will appear as a channel distortion and potentially result in worse performance than without a correction factor. While the importance of initial channel estimation cannot be overemphasized, it will be equally critical to track any changes to the channel, which may be caused by thermal or other effects. This is required to have an accurate correction factor for the current conditions at all times. For practical mass-market deployment it will also be important to be able to address situations where the usage of the lines within a binder changes over time such as inactive lines being put into service to support new users or existing users switching their modems off and on. All of this must be accommodated without disrupting existing in-service customers. This may be addressed by estimating the crosstalk from this new line before the line is activated, and, if necessary, the CPE may be put "on hold" while the CO side is estimating crosstalk. As such, a user can be added without bit rate impact on other lines, although this may require a longer start-up time.

Performance Improvements

From our earlier definition of DSM level 3, we know that signals from all the lines in a binder are jointly processed and that the majority of the implementation changes need to be at the DSLAM.

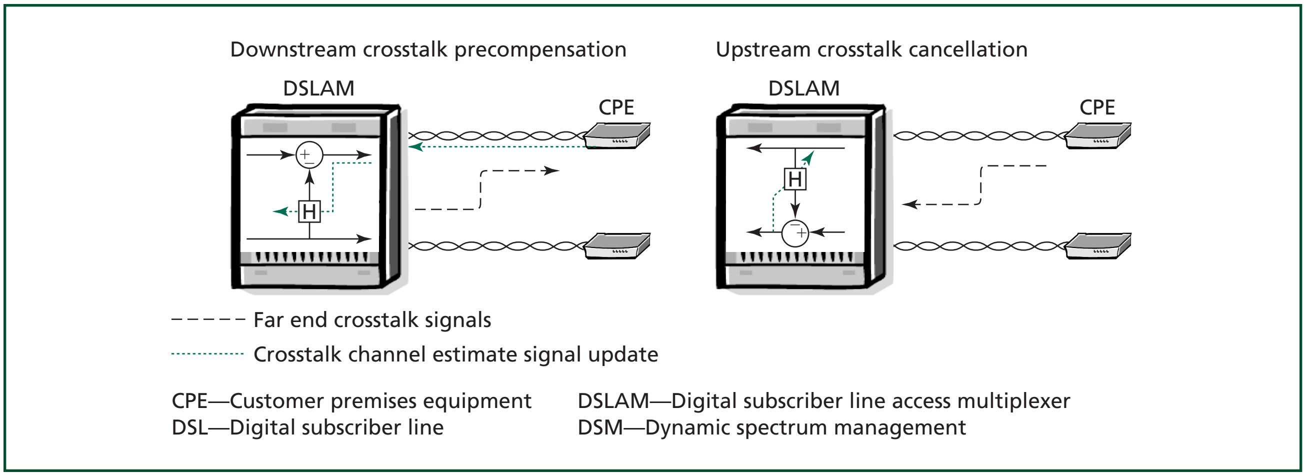

Figure 2.
DSM level 3 architecture.

Such an implementation is shown schematically in **Figure 2**.

A DSM level 3 processing architecture consists of the following steps:

1. Estimate the crosstalk channels.
2. Calculate a compensation factor for each line which may be used.
3. Apply a per-line compensation factor on a line-by-line basis to cancel the crosstalk.

In the downstream direction this involves:

1. The DSLAM sends reference signals to the CPE in a coordinated way in order to estimate the crosstalk channel. Each piece of customer premises equipment analyzes the received signals and sends a report back to the DSLAM, which relays the information to the DSM level 3 device.
2. The DSM level 3 device uses the information retrieved from the CPE to estimate the compensation factors for each downstream line.
3. The precompensation signals are applied to the data and transmitted on each line.

A similar process is followed in the upstream direction. However, in this case, the correction factor is calculated in the DSLAM from the signals it receives from the CPE. The correction factor is applied before decoding the data-carrying received signal.

Since there are no commercial implementations of precoders available today to evaluate the performance for in-service lines in the field, we use simulations and experiments to estimate the performance gains to be expected with DSM level 3. Without the more accurate crosstalk channel model discussed above, we can only reasonably perform simulations with the 99 percent worst case crosstalk model to estimate the gains. The results of such simulations [6] and the resultant performance comparison are shown in **Figure 3**. Note that the conservative and optimistic rate and reach improvements are graphically illustrated for VDSL 8 MHz profile usage. Similar principles can be used to observe improvements for 17 MHz profile VDSL. The optimistic performance gain case assumes complete crosstalk cancellation while the conservative estimate considers the precompensation of all but one (worst case) crosstalker. Observing that this crosstalk model may optimistically estimate performance gains, the simulations show potential rate improvements of about 25 percent on 1000 meter loops, decreasing to no gain on loops longer than about 1500 meters, where other noise sources become dominant. Note that downstream performance following crosstalk cancellation for loops shorter than about 750 meters will benefit from the availability of increased spectrum, e.g., the use of VDSL2 profile 17a. The improvements for the shortest loops may not be as significant if the majority of the frequencies achieve bit-loading maxima, or values close to it, even without DSM solutions. What can be seen from these simulations is that crosstalk has a dominant

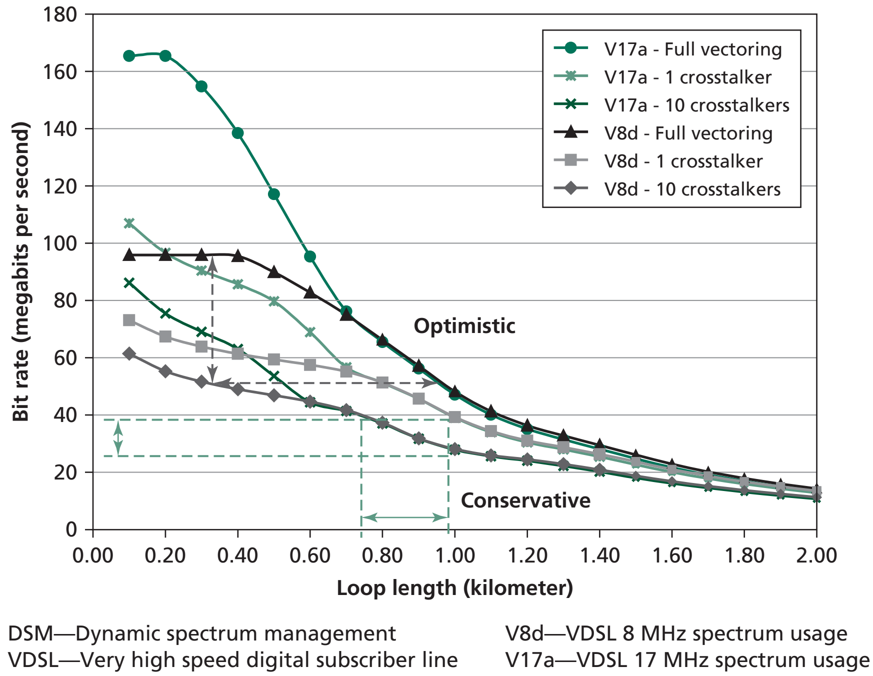

Figure 3.
Typical VDSL performance with and without DSM level 3.

performance-limiting effect on shorter loops, in particular loops less than 1500 meters. A more detailed analysis can be found in [6]. This suggests that the eventual deployment of DSM level 3 solutions should be targeted at DSLAMs located within 1500 meters of the customer premises, which is typical in FTTN deployments.

DSM Level 3 Approach

Given that crosstalk is the dominant performance-limiting factor for loops shorter than 1500 meters and that VDSL delivers increased bit rate in this environment, it is our objective to focus our research efforts on DSM level 3 in conjunction with VDSL technology. Our simulation results suggest that the deployment of crosstalk mitigation solutions is the most relevant when the DSLAM is relatively close to the household (e.g., FTTN). In such deployment scenarios, the characteristics of the distribution cables are of importance—while some older feeder cables may contain as many as 100 wire pairs per binder, most newer distribution cables consist of 25-pair binders. The smaller number of pairs in the binder means a potentially smaller number of crosstalk sources, which relaxes the computational and processing requirements for the DSM level 3 solution.

The effects of crosstalk are frequency dependent and become more dominant for higher frequencies. In the case of VDSL deployments where ADSL deployments remain in place at least for some time, this may be important. ADSL uses just a fraction of the VDSL spectrum at lower frequencies, which means that crosstalk from ADSL is relatively small compared to crosstalk from VDSL. Therefore, DSM level 3 solutions need to address only crosstalk at higher frequencies while ADSL crosstalk impairments could be either ignored or represented by a minimal fixed degradation. This approach facilitates mixed deployments, which is particularly important in markets where competitive access providers have co-located their own DSLAMs in the CO. Therefore, it is safe to allow these DSLAMs to remain in place while deploying VDSL lines which collectively support DSM level 3.

Crosstalk is particularly strong on lines that share the same binder and therefore DSM level 3 will need to scale to the number of lines that are typically present in a binder. In current deployments, inter-binder crosstalk—which is more than 10 dB lower than intra-binder crosstalk—can usually be ignored. Further investigations are required to determine if inter-binder crosstalk may still become a limiting factor on performance when all intra-binder crosstalk is cancelled.

Deployment-Related Aspects

The goal of our DSM level 3 research efforts is to develop a technology that will allow improvements of the rate and reach performance of DSL lines. While more research is required to determine the most efficient implementations and improved performance estimates, there are several deployment-related issues that can be discussed prior to implementation.

It is certain that as more loops are used in a binder, more crosstalk is experienced by each individual loop in that binder. On that basis, any field deployment strategy that increases binder utilization will see a bigger benefit from DSM level 3. Increased binder utilization will occur as customer take rates increase or when multiple pairs are used per customer.

DSM Level 3 Binder Termination

Consider a typical scenario as shown in **Figure 4** where multiple competitive access providers have deployed DSLAMs in the CO to provide HSIA.

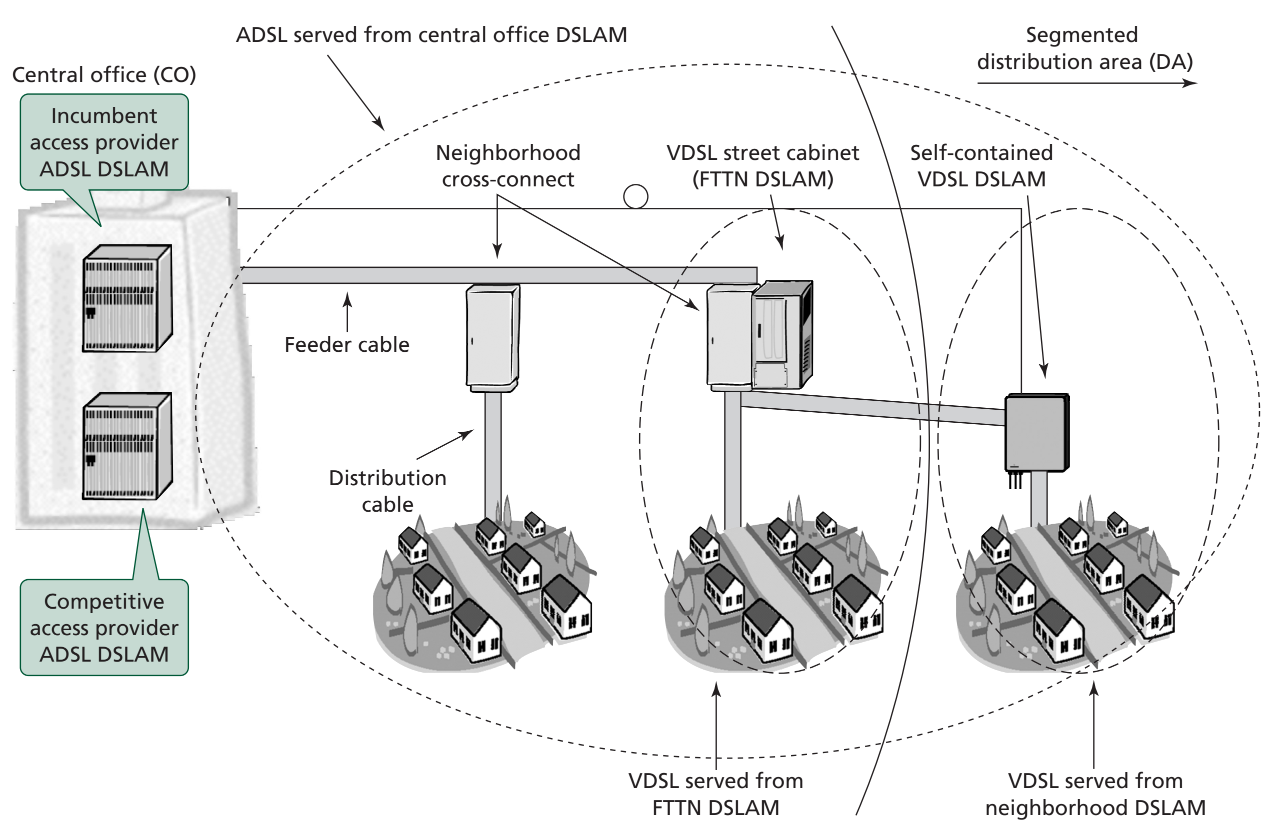

ADSL—Asymmetric digital subscriber line
DSL—Digital subscriber line
DSLAM—Digital subscriber line access multiplexer
FTTN—Fiber to the node
VDSL—Very high speed digital subscriber line

Figure 4.
Typical DSLAM deployment locations.

The incumbent operator may deploy VDSL in street cabinets to address an enhanced broadband service offering and self-contained VDSL DSLAMs in either a fiber to the curb (FTTC) or segmented distribution area (DA) arrangement. Let us assume that all active lines in a binder are co-processed. Even in instances where real-time cancellation need not be applied to all loops, the channel estimation process will likely still require access to all VDSL2 loops. For any VDSL2 line in a binder to benefit from DSM level 3, all VDSL2 lines within that binder must originate from the same location and terminate at the same DSM level 3 device. Any issues with this approach may need to be addressed by deployment rules.

It must be assumed that all VDSL deployments for services on loops within a binder are co-located and under the administration of a single operator. This allows for the DSM level 3 requirement of co-location of all transmitters/receivers at the DSLAM and the ability to apply precoding/cancellation to all relevant lines in the binder. In some situations, not all customers are able to achieve the desired service rate from the FTTN location so either DA segmentation or FTTC deployments may be adopted. Homes which are served by DSM level 3 enabled VDSL must in all cases be served from the nearest upstream DSLAM, and no homes can be served from a downstream DSLAM. These deployment rules seem reasonable assuming that FTTN is deployed to achieve a certain minimum service rate. Whenever this rate is not achievable for all customers, those that are farther away are served from a segmented DA or FTTC.

The above scenario considers the situation where all DSL lines in use are connected to a binder and arrive at the same DSM level 3 capable DSLAM location. However, there is the additional condition that all lines in use in a binder must arrive at the same DSM level 3 device. Depending on the location of this device and its modularity, the lines that require co-processing may not naturally arrive at the same device. In addition, existing provisioning rules which affect initial installations and customer churn may assign cable pairs to ports on a first-come first-served basis rather than grooming the pairs to the appropriate DSM level 3 device. Since all active cable pairs that are located in a binder must be co-processed, the correlation between DSL lines and DSM level 3 co-processing devices must be maintained. If DSM level 3 is implemented completely within the DSL chipset, this implies that chipsets must be at least at the modularity of the number of lines in a binder. If DSM level 3 is implemented outside the chipset but as a co-processor, then support at the line card level may be sufficient, provided, as indicated above, that the line card port density is at least as big as the number of active lines in a binder. However, individual pair routing may not connect all the relevant pairs that are used in a binder to the same line card. If this same strategy is extended across multiple line cards or the DSLAM as a whole, then pair routing may be eased but a significant amount of data will need to be exchanged between line cards and the centralized DSM level 3 server in real time. This may still be insufficient where multiple DSLAMs are deployed at a single location. Finally, options which support more flexible cable routing may be considered such as automated distribution frames (ADFs). Care must be taken here as this technology remains largely unproven in VDSL applications and may introduce inter-binder crosstalk, depending upon wire connection strategy and the ADFs' internal design. Additionally, any solution that allows wiring freedom by supporting some sort of centralized DSM level 3 functionality will require that all modems on multiple line cards, and potentially multiple DSLAMs, be synchronized to the same baud clock. Furthermore, when a DSLAM connection cable (from line card to the main distribution frame) contains multiple wires from one binder together with multiple wires from another binder, there may be additional crosstalk for which the crosstalk cancellation precoder has not been designed. This problem gets worse when the connection cables are long.

One option that avoids the problems mentioned above is the 100 percent provisioning model. As the capital cost of equipment is reduced and customer churn remains, it may be more economical (particularly for OSP deployments) to provision for 100 percent of customers, i.e., to deploy resources to service every customer. If a customer is lost, the service can be administratively disabled and no truck roll will be required. When this household is recaptured, service can simply be re-enabled via administrative commands. This strategy comes with a somewhat higher initial cost as capital is expended for service terminating equipment which may not be required until some later time. These costs can often be easily justified, especially in higher operational expenditure (OPEX) cost countries as they are offset within two or three visits to the remote DSLAM. These and other issues are discussed in the next section, where potential DSM level 3 implementation options are compared.

DSM Level 3 Implementation Scenarios

In this section we discuss several alternatives for implementation of the precoding/cancellation functions. The advantages and disadvantages of each approach and the need to upgrade existing deployments are considered.

We observe the following three implementation scenarios for DSM level 3:

1. Integration of a DSM level 3 solution at the DSL chipset.
2. The usage of a DSM level 3 server.
3. A stand-alone DSM level 3 processor implementation.

DSM level 3 can potentially be integrated in a DSL chipset, i.e., precoding and cancellation are performed entirely within the DSL chipset. A major advantage is the accessibility of the data that is transmitted across the lines.

A DSM level 3 server could alternatively be located outside the chipset to perform crosstalk processing functions. By separating crosstalk processing from the DSL transceiver elements and providing access to the internal signal paths at the appropriate points, a scalable solution may be implemented at various levels and modularity within the DSLAM. Implementing such a server on each line card would simplify the design by keeping all the high-speed chipset interfaces to the DSM server on the line card. Depending upon card density, this may still present challenges such as wiring issues, since once the first line from a binder is terminated on a line card, any subsequent line used in that binder must also terminate on that card. This may be most easily addressed by the 100 percent service model, as already discussed. Raising this crosstalk co-processor concept to the DSLAM level may relieve most wire routing issues, assuming the deployment of a single DSLAM per location. However, this now requires high speed internal interfaces between cards to exchange real-time co-processing data. This may only be practical for DSLAMs with a low port count (e.g., those with up to 24 ports). Taking this concept one step further, an external DSM level 3 server could support a group of co-located DSLAMs, which would eliminate any wire routing issues. As for the previous option, the high-speed real-time data interfaces required may make this solution impractical.

A stand-alone solution does not rely on new chipsets in the DSLAM or chipset interfaces but instead comprises a stand-alone device placed between a CO DSLAM and the lines it serves to perform inline precoding and cancellation. Such an implementation might prove most valuable if it allows existing DSLAMs to be reused without requiring any hardware upgrade. While such a stand-alone DSM level 3 processor would cause each line to be interrupted when the device is inserted in-line, this opportunity could be taken to meet the requirement that the loops are connected such that the used lines within a binder arrive at the same DSM level 3 processing device.

It should be noted that any option that spans multiple line cards and includes the possibility of modems from multiple line cards arriving in the same binder requires modem baud clock synchronization.

Conclusion

This paper discussed several aspects of DSM level 3, a promising new DSL-related technology that may offer a path to provide further rate improvements in order to meet the ever-increasing demands for higher bandwidth. It may therefore further increase the commercial lifespan of deployed copper access networks. In order to address the demands, an evolution from traditional CO-based ADSL deployments may be necessary toward shorter loop deployments using VDSL technology. Typically, these existing ADSL deployments can be left in place and complemented by DSM level 3 enabled VDSL deployments. The support for DSM level 3 offers a potential performance improvement provided the following deployment conditions are met:

- All modems are co-located at one end of the binder and available for coordination,
- Multiple modems are in use in the same binder (no benefit for a single line),
- Loops are shorter than 1000 meters to 1500 meters in length, and
- All lines in a binder appear at the same DSM level 3 device for co-processing.

Before such a technology can be developed and deployed, further research and standardization activities are necessary, including:

- Crosstalk channel modeling that includes both intra- and inter-binder effects, and
- Channel estimation and tracking.

Furthermore, a commercialized solution will require an understanding of not only the underlying DSL technology but also line card and DSLAM system design and field deployment practices. In addition, scalable solutions will require adoption of higher level line management capabilities to complement element management systems.

References

[1] Alliance for Telecommunications Industry Solutions (ATIS), "American National Standard for Telecommunications—Network and Customer Installation Interfaces, Asymmetric Digital Subscriber (ADSL) Metallic Interface," American National Standards Institute (ANSI) T1.413-1998, Nov. 1998.

[2] Alliance for Telecommunications Industry Solutions (ATIS), "American National Standard for Telecommunications—Spectrum Management for Loop Transmission Systems," American National Standards Institute (ANSI) T1.417-2003, Sept. 2003.

[3] Alliance for Telecommunications Industry Solutions (ATIS), "Pre-Published American National Standard for Telecommunications—Dynamic Spectrum Management," prepared by NIPP-NAI, ATIS-PP-0600007, May 2007.

[4] International Telecommunication Union, Telecommunication Standardization Sector, "Ethernet-Based Multi-Pair Bonding," ITU-T Rec. G.998.2, Jan. 2005.

[5] International Telecommunication Union, Telecommunication Standardization Sector, "Very High Speed Digital Subscriber Line Transceivers 2 (VDSL2)," ITU-T Rec. G.993.2, Feb. 2006.

[6] P. Whiting, A. Ashikhmin, S. Borst, J. Jennen, G. Kramer, A. J. de Lind van Wijngaarden, and M. Živković, "Performance Results for Digital Subscriber Line Precoders," Bell Labs Tech. J., 13:1 (2008), 147–161.

[7] W. Yu, Competition and Cooperation in Multi-User Communication Environments, Ph.D. Thesis, Stanford University, June 2002.

(Manuscript approved November 2007)

CHARLES STORRY is a member of the Chief Technology Officer (CTO) team in Alcatel-Lucent's Access Division and is located in Ottawa, Canada. He graduated magna cum laude, receiving an honors degree of bachelor of science in computer science, from the University of Ottawa, Ottawa, Canada. He is a member of the Alcatel Lucent Technical Academy. His core area of responsibility is next-generation DSL/copper access.

MIROSLAV ŽIVKOVIĆ is a member of technical staff in Alcatel-Lucent Bell Labs Europe in Hilversum, The Netherlands. He holds a Dipl. Ing. degree in electronics and telecommunications from the Faculty of Electrical Engineering at the University of Belgrade in Belgrade, Serbia. His current research interests are in the area of dynamic spectrum management for digital subscriber line (DSL) systems.

JAN VERLINDEN is a member of the DSL Experts Team at Alcatel-Lucent Bell in Antwerp, Belgium. He received his master in electrotechnical engineering degree from Katholieke Universiteit Leuven, in Belgium. During his tenure at Alcatel's Research and Innovation division, he has focused on echo canceller techniques as well as dynamic spectrum management (DSM). As such, he participated in the VDSL Olympics by introducing DSM into the VDSL prototype. He also contributes to the American National Standards Institute (ANSI) Network Interface Power and Protection/Network Access Interfaces (NIPP-NAI) standardization effort, which approved the DSM Technical Report in May 2007.

ADRIAAN J. DE LIND VAN WIJNGAARDEN is a member

of technical staff in Bell Labs' Communications and Statistical Sciences Department in Murray Hill, New Jersey. He received his M.S. in electrical engineering from Eindhoven University of Technology, The Netherlands, and his doctorate in engineering from the University of Essen, Germany. He has been deeply engaged in both theoretical and application-driven research in communications, information theory, coding, combinatorial and algorithmic optimization. He provided key contributions to Alcatel-Lucent's 10G and 40G optical systems, in-building wireless and broadband access. He has authored more than 60 technical papers, received five patents, and has several patents pending. Dr. de Lind van Wijngaarden is a Senior Member of the IEEE and serves as a Publications Editor of the IEEE Transactions on Information Theory. He co-organized Shannon Day at Bell Labs in 1998, Shannon's Statue Unveiling at Bell Labs in 2001, DIMACS conferences, and several international workshops. ◆