TR-320
Techniques to Mitigate Uncancelled Crosstalk on Vectored VDSL2 Lines

Notice

The Broadband Forum is a non-profit corporation organized to create guidelines for broadband network system development and deployment. This Broadband Forum Technical Report has been approved by members of the Forum. This Broadband Forum Technical Report is not binding on the Broadband Forum, any of its members, or any developer or service provider. This Broadband Forum Technical Report is subject to change, but only with approval of members of the Forum. This Technical Report is copyrighted by the Broadband Forum, and all rights are reserved. Portions of this Technical Report may be copyrighted by Broadband Forum members.

The Forum draws attention to the fact that it is claimed that compliance with this Specification may involve the use of a patent ("IPR") concerning TR-320. The Forum takes no position concerning the evidence, validity or scope of this IPR.

The holder of this IPR has assured the Forum that it is willing to License all IPR it owns and any third party IPR it has the right to sublicense which might be infringed by any implementation of this Specification to the Forum and those Licensees (Members and non-Members alike) desiring to implement this Specification. Information may be obtained from:

Actelis Networks Inc.
47800 Westinghouse Drive, Fremont, CA 94539 USA

Adaptive Spectrum and Signal Alignment, Inc.
333 Twin Dolphin Drive, Redwood City, CA 94065 USA

Attention is also drawn to the possibility that some of the elements of this Specification may be the subject of IPR other than those identified above. The Forum shall not be responsible for identifying any or all such IPR.

THIS SPECIFICATION IS BEING OFFERED WITHOUT ANY WARRANTY WHATSOEVER, AND IN PARTICULAR, ANY WARRANTY OF NONINFRINGEMENT IS EXPRESSLY DISCLAIMED. ANY USE OF THIS SPECIFICATION SHALL BE MADE ENTIRELY AT THE IMPLEMENTER'S OWN RISK, AND NEITHER the Forum, NOR ANY OF ITS MEMBERS OR SUBMITTERS, SHALL HAVE ANY LIABILITY WHATSOEVER TO ANY IMPLEMENTER OR THIRD PARTY FOR ANY DAMAGES OF ANY NATURE WHATSOEVER, DIRECTLY OR INDIRECTLY, ARISING FROM THE USE OF THIS SPECIFICATION.

Broadband Forum Technical Reports may be copied, downloaded, stored on a server or otherwise re-distributed in their entirety only, and may not be modified without the advance written permission of the Broadband Forum.

The text of this notice must be included in all copies of this Broadband Forum Technical Report.

Issue History

Issue Number	Issue Date	Publication Date	Issue Editor	Changes
1	March 2014	16 May 2014	Peter Silverman	Original

Comments or questions about this Broadband Forum Technical Report should be directed to help@broadband-forum.org.

Editor	Peter Silverman	ASSIA, Inc.	psilverman@assia-inc.com
Metallic Transmission WG Chair	Les Brown	Huawei Technologies	lesbrown@sympatico.ca
Metallic Transmission WG Vice Chair	Lincoln Lavoie	UNH InterOperability Lab	lylavoie@iol.unh.edu
Metallic Transmission WG Vice Chair	Massimo Sorbara	Ikanos Communications	msorbara@ikanos.com

Table of Contents

March 2014   6 of 101

List of Figures

List of Tables

Executive Summary

ITU-T Recommendation G.993.5 [2] defines a transmission method called vectoring, which in combination with G.993.2 can be used to reduce the far-end (self-FEXT) crosstalk levels and improve performance of a group of VDSL2 transceivers. In the downstream direction, vectoring is performed at the CO-side transmitter via the FEXT cancellation pre-coding. In the upstream, it is performed by post-compensating for it at the CO-side receiver via crosstalk cancellation.

Removing self-FEXT greatly improves the performance of VDSL2. Despite the removal of the self-FEXT, vectoring performance may be degraded by the uncancelled crosstalk that is related to some deployment scenarios driven by particular regulatory regime, service deployment strategies or vectoring implementation. Some of these scenarios may be related to the particular regulatory regime, commercial or competitive environments, particular deployment strategies, or to the particular service or CPE choices of customers, and also to vectoring implementation

This Technical Report explores available options for avoiding or mitigating the impact of uncancelled crosstalk on the performance of vectored lines. After clarifying the various types of uncancelled crosstalk that may impact vectoring, it describes such techniques in terms of complexity, effectiveness, operational conditions, and regulatory consequences for effectively deploying them in the network.

Although there are many practical situations where uncancelled crosstalk may degrade vectoring, operators have several tools at their disposal for mitigating its impact on vectored lines. None of these tools by itself can completely remove all types of uncancelled crosstalk, but using a combination of System Level Vectoring (SLV), Cross-DSLAM Vectoring (xDLV), vectoring friendly CPEs, and Dynamic Spectrum Management (DSM) can maintain vectoring gains in most deployments with or without physical Sub-Loop Unbundling (SLU).

Chapter 6 of this Technical Report discusses in detail DSM, Cross-DSLAM Vectoring (xDLV) and Cable Level Vectoring (CLV) as techniques for mitigating or eliminating uncancelled crosstalk. Some of the technologies described in this chapter may be emerging or non-standardized at this time. Chapter 6 highlights the possible trade-offs and limitations, availability in a mono or multi-vendor environment and the status of standardization as applicable to each technique.

1   Purpose and Scope

1.1   Purpose

 The purpose of this Technical Report is to address best practices and analysis regarding techniques for mitigating or avoiding the impact of uncancelled crosstalk on vectored lines, in cases of coexistence of vectored and non-vectored lines in the same cable/binder and coexistence of multiple vectored groups in the same binder. The contents are meant to provide information to the industry and there are no normative requirements stated in this Technical Report.

1.2   Scope

Although vectoring can cancel the crosstalk between vectored lines, the presence of neighboring lines that are non-vectored or belong to a different vectored group can lead to the presence of uncancelled crosstalk, i.e. crosstalk that vectoring does not remove. This Technical Report addresses techniques for mitigating or avoiding the impact of this uncancelled crosstalk on vectored VDSL2 lines. The scope includes techniques to address the coexistence of vectored and non-vectored VDSL2 lines in the same cable/binder as well as crosstalk between multiple vectored groups in the same binder, possibly managed by multiple network operators The primary scope will be techniques to address the coexistence of vectored and non-vectored DSLs in the same cable/binder as well as crosstalk between multiple vectored groups in the same binder, possibly managed by multiple network operators.

The Technical Report addresses best practices and analysis in these areas based upon quantitative analysis, simulation, and possibly field/lab measurements and findings. Techniques based upon management techniques and algorithms will be analyzed. These include among others Dynamic Spectrum Management techniques (DSM Level 1, 2 and 3 tools), Cross-DSLAM Vectoring (xDLV), and Cable Level Vectoring (CLV), and binder management.

2   References and Terminology

2.1   Conventions

In this Technical Report, several words are used to signify the requirements of the specification. These words are always capitalized. More information can be found be in RFC 2119 [15].

SHALL	This word, or the term "REQUIRED", means that the definition is an absolute requirement of the specification.
SHALL NOT	This phrase means that the definition is an absolute prohibition of the specification.
SHOULD	This word, or the term "RECOMMENDED", means that there could exist valid reasons in particular circumstances to ignore this item, but the full implications need to be understood and carefully weighed before choosing a different course.
SHOULD NOT	This phrase, or the phrase "NOT RECOMMENDED" means that there could exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications need to be understood and the case carefully weighed before implementing any behavior described with this label.
MAY	This word, or the term "OPTIONAL", means that this item is one of an allowed set of alternatives. An implementation that does not include this option MUST be prepared to inter-operate with another implementation that does include the option.

2.2   References

The following references are of relevance to this Technical Report. At the time of publication, the editions indicated were valid. All references are subject to revision; users of this Technical Report are therefore encouraged to investigate the possibility of applying the most recent edition of the references listed below.

A list of currently valid Broadband Forum Technical Reports is published at www.broadband-forum.org.

Document		Title	Source	Year
[1]	G993.2	Very high speed digital subscriber line Transceivers 2 (VDSL2)	ITU-T	2011
[2]	G.993.5	Self-FEXT Cancellation (Vectoring) for use with VDSL2 transceivers	ITU-T	2010
[3]	TR-188	DSL Quality Suite	BBF	2010
[4]	TR-198	DQS: DQM systems functional architecture	BBF	2011

		and requirements		
[5]	TR-197	DQS: DSL Quality Management Techniques and Nomenclature	BBF	2012
[6]	ATIS 0600007.2	Dynamic Spectrum Management - Technical Report (Issue 2)	ATIS	2012
[7]	ND1513	Report on Dynamic Spectrum Management (DSM) Methods in the UK Access Network	NICC	2010
[8]	ANIS T1.417	Spectrum Management for Loop Transmission Systems, ANSI standard	ATIS	2003
[9]		S. Huberman, C. Leung, T. Le-Ngoc, "Dynamic Spectrum Management (DSM) Algorithms for Multi-User xDSL	IEEE Communications Surveys & Tutorials, vol. 14, no. 1, First Quarter 2012	2012
[10]		R. Cedrillon, et., al., "Optimal Multiuser Spectrum Balancing for Digital Subscriber Lines"	IEEE Transactions on Communications, Vol. 54, pp. 922-933, May 2006	2006
[11]	ATIS 0600024	Multiple-Input Multiple-Output Crosstalk Channel Model	ATIS-0600024, April 2009	2009
[12]		K. Kerpez, J. Cioffi, S. Galli, G. Ginis, M. Goldburg, M. Mohseni, A. Chowdhery, "Compatibility of Vectored and Non-Vectored VDSL2	IEEE Conference on Information Sciences and Systems (CISS), Princeton, NJ, Mar. 21–23, 2012	
[13]		W. Yu, G. Ginis, J. Cioffi, "Distributed multiuser power control for digital subscriber lines,"	IEEE J. Sel. Areas Commun., vol. 20, pp. 1105–1115, Jun. 2002	2002
[14]		A Colmegna, S Galli, M Goldburg, "Methods for Supporting Vectoring when Multiple Service Providers Share the Cabinet Area,"	FASTWEB/ASSIA White Paper,	April 2012
[15]	RFC 2119	*Key words for use in RFCs to Indicate Requirement Levels*	IETF	1997

2.3 Definitions

The following terminology is used throughout this Technical Report.

Alien Crosstalk   Crosstalk created by alien lines.

Alien Lines	A set of lines is "alien" to a second set of lines within the same cable if its lines carry a DSL signal type that is different from the one carried by the lines in the second set.
	In the context of vectoring, alien lines are lines that carry any non-VDSL2 DSL signal and share the same cable with lines in a Pre-Coded Group.
Alien Noise	Any non-crosstalk noise impairing a DSL line, e.g. impulse noise, RFI, power line communication interference, etc.
Binder	A group of twisted pairs in a telephone cable that are in close proximity to other pairs in the same binder throughout the length of cable due to the process by which the cable was manufactured. Twenty-five pairs is a common size of a binder. Each twisted pair within a binder is identified by a color code unique within a particular binder placed on the wire's insulation.
Binder Group	An individual binder within a telephone cable that contains multiple binders. Each binder group is identified by a ribbon or sheath that surrounds that particular binder and identifies that group with a color code that is unique within that cable to that binder group.
Binder Management	A method to avoid out of domain self-FEXT by assigning DSLAM ports to cable pairs so that only one vectored group resides within a cable binder group.
Board Level Vectoring (BLV)	A vectoring architecture where a vectored group can span at most over the lines terminating on a single line-card. In BLV, there is one vectored group per line-card, and the lines terminating on different line-cards belong to different vectored groups.
Bonding	Use of multiple DSL lines inverse multiplexed at the DSL level to carry a single application payload to a customer over multiple copper loops. DSL Bonding is defined in ITU-T Recommendations G.998.1, G.998.2, and G.998.3
Cable Level Vectoring	In Cable Level Vectoring (CLV), the operation of vectoring is performed across all the pairs in a cable, regardless of whether they terminate on multiple DSLAMs or not.
Cross-DSLAM Vectoring	A vectoring architecture where the operation of vectoring is performed across multiple DSLAMs by coordinating them so that the vector group spans lines that terminate on different vectored DSLAMs.
Crosstalk	Interfering signal received in one copper pair of a cable from services in other copper pairs of the same cable.
DSM	Dynamic Spectrum Management. DSM is an optimization framework incorporating parameters of the subscriber line environment and transmission systems that are time or situation dependent.
Error Sample	The measurement made by a VDSL2 receiver supporting Vectoring that indicates the effect of crosstalk received into loop serving the VDSL2 Line.
Far-End Crosstalk	Crosstalk between DSL services at the far end of the copper loop away from the DSL transmitter.

In-domain Self-FEXT	This type of self-FEXT is generated by lines that belong to the same Vectored Group. There are three notable cases in vectoring:
	1) The in-domain Self-FEXT generated by the lines in the Pre-Coded Group is cancelled by vectoring in both downstream and upstream.
	2) The in-domain Self-FEXT generated within the vectored group but outside of the Pre-Coded Group is cancelled by vectoring in both downstream and upstream if and only if those lines terminate on full vectoring-friendly CPEs.
	3) The in-domain Self-FEXT generated within the vectored group but outside of the Pre-Coded Group is cancelled by vectoring in downstream only if and only if those lines terminate on downstream vectoring-friendly CPEs.
Legacy CPE	A VDSL2 CPE that is neither downstream vectoring-friendly (G.993.2 Annex X), nor full vectoring-friendly (G.993.2 Annex Y), nor vectoring (G.993.5) capable.
Near-End Crosstalk	Crosstalk between DSL services at the near end of the copper loop near the DSL transmitter.
Out-of-domain Self-FEXT	This type of self-FEXT is generated by lines that do not belong to the Vectored Group. The out-of-domain Self-FEXT cannot be cancelled by vectoring.
Pre-Coded Group	The subset of lines in a vectored group on which vectoring is actually performed, i.e. lines that are terminated on both a vectoring-capable DSLAM and on vectoring capable CPEs. In the downstream (upstream), the vectoring is performed at the transmitter (receiver) side via pre-coding (post-compensation).
Pre-coder	The function for the downstream direction that performs the mathematical operation of self-crosstalk cancelation in a vectored group.
Self-Crosstalk	Crosstalk generated by neighboring VDSL2 lines.
Self-FEXT	FEXT created by lines carrying DSL signals of the same type.
	In vectoring context, FEXT generated by neighboring VDSL2 lines, either vectored or not. There are two types of self-FEXT: in-domain and out-of-domain.
Showtime	The state of a DSL connection when application payload data can be transmitted over the connection.
System Level Vectoring (SLV)	A vectoring architecture where a vectored group can span over all the lines terminating on the vectoring capable DSLAM. In SLV, there is only one vectored group per DSLAM.
Vectored Group	The set of lines over which transmission from the AN is eligible to be coordinated by pre-compensation (downstream vectoring), or over which reception at the AN is eligible to be coordinated by post-compensation (upstream vectoring), or both. Depending on the configuration of the vectored group, downstream vectoring, upstream vectoring, both or none may be enabled (see ITU-T Rec. G.993.5 clause 3 - definitions).

Vectoring	The coordinated transmission and/or coordinated reception of signals of multiple DSL transceivers using techniques to mitigate the adverse effects of crosstalk to improve performance (see ITU-T Rec. G.993.5 clause 3 - definitions) [2].
Vectoring Control Entity	The function in a vectored System that manages vectoring for the lines in a DSLAM.
Vectoring Friendly	Vectoring friendly operation is defined in the ITU-T G.993.2 Annex X (downstream friendly operation) and Annex Y (downstream and upstream or "full" friendly operation).

Vectoring friendly operation as defined in Annex X allows cancellation of the downstream in-domain self-FEXT from lines equipped with downstream vectoring-friendly CPEs into the vectored lines of a Pre-Coded Group.

Full vectoring friendly operation as defined in Annex Y allows cancellation of the downstream and upstream in-domain self-FEXT from lines equipped with full vectoring-friendly CPEs into the vectored lines of a pre-Coded Group.

2.4 Abbreviations

This Technical Report uses the following abbreviations:

ADSL	Asymmetric Digital Subscriber Line
AN	Access Node
ANFP	Access Network Frequency Plan
AWG	American Wire Gauge
AWGN	Additive White Gaussian Noise
BLV	Board Level Vectoring
CDF	Cumulative Distribution Function
CLV	Cable Level Vectoring
CLVE	Cable Level Vectoring Equipment
CO	Central Office
CPE	Customer Premises Equipment
DPBO	Downstream Power Back Off
DSEL	D-Side Electrical Length
DSL	Digital Subscriber Line
DSM	Dynamic Spectrum Management
ESEL	E-Side Electrical Length
FEXT	Far-end Crosstalk
FSAN	Full Service Access Network
INP	Impulse Noise Protection

ISB	Iterative Spectrum Balancing
IWF	Iterative Waterfilling
L-CPE	Legacy CPE
MIMO	Multiple In, Multiple Out
MLWF	Multilevel Water Filling
OAM	Operations, Administration and Maintenance
O-CPE	Other than VDSL2 CPE
OSB	Optimum Spectrum Balancing
OSS	Operations Support System
PBO	Power Back-Off
PSD	Power Spectrum Density
RT	Remote Terminal
SHDSL	Single-pair High-speed Digital Subscriber Line
SLU	Sub-Loop Unbundling
SLV	System Level Vectoring
SM	Spectrum Management
SMC	Spectrum Management Center
SNR	Signal to Noise Ratio
SP	Service Provider
UPBO	Upstream Power Back Off
V-CPE	Vectoring CPE
VDSL2	Very High-speed Digital Subscriber Line, Issue 2
VC-CPE	Vectoring Capable CPE
VF-CPE	Vectoring Friendly CPE
xDLV	Cross-DSLAM Level Vectoring

3   Technical Report Impact

3.1   Energy Efficiency

The impact of the techniques described in TR-320 on Energy Efficiency is for further study.

3.2   IPv6

TR-320 has no impact on IPv6.

3.3   Security

TR-320 has no impact on security.

3.4   Privacy

Any issues regarding privacy are not affected by TR-320.

4 Introduction

ITU-T Recommendation G.993.5 [2] defines a transmission method called vectoring, which in combination with G.993.2 can be used to reduce the far-end (self-FEXT) crosstalk levels and improve performance of a group of VDSL2 transceivers. In the downstream direction, vectoring is performed at the CO-side transmitter via the FEXT cancellation pre-coding. In the upstream direction, it is performed by post-compensating for it at the CO-side receiver via crosstalk cancellation.

Removing self-FEXT greatly improves the performance of VDSL2. "Despite the removal of the self-FEXT, vectoring performance may be degraded by the uncancelled crosstalk that is related to some deployment scenarios driven by particular regulatory regime, service deployment strategies or vectoring implementation. This uncancelled self-FEXT may arise from some scenarios commonly found in the field. Some of these scenarios may be related to the particular regulatory regime, commercial or competitive environments, particular deployment strategies, or to the particular service or CPE choices of customers, and also to vectoring implementation.

This Technical Report addresses the techniques for avoiding or mitigating the impact of uncancelled self-FEXT that is not cancelled by vectoring on the performance of vectored lines. After clarifying the various types of uncancelled crosstalk that may impact vectoring, it describes such techniques in terms of complexity, effectiveness, operational conditions, and regulatory consequences for effectively deploying them in the network.

5   The various types of crosstalk that are relevant to vectoring

In vectoring, the transmitters and receivers of all the lines in a Vectored Group can cooperate to remove all the self-FEXT they create if all the lines in the Vectored Group are terminated on vectoring capable CPEs. If all lines in a Vectored Group terminate on vectoring-capable CPEs, then the Vectored Group and the pre-coded group coincide. In the downstream, transmitters collocated at the DSLAM cooperate to eliminate self-FEXT by performing pre-subtraction (pre-coding) of the self-FEXT that will be found at the receiver. As self-FEXT is pre-subtracted at the DSLAM, the modem at the customer premises experiences a signal that is self-FEXT free. In the upstream, receivers collocated at the DSLAM cooperate to cancel the received self-FEXT (post-compensation).

There are scenarios of practical interest where not all crosstalk impairing the vectored group can be cancelled by vectoring. In this Section, the various types of uncancelled crosstalk that may impact the performance of vectored lines are categorized and clarified as to what crosstalk can or cannot be cancelled by vectoring. In the next Section, we will review what techniques are available for avoiding or mitigating the uncancelled crosstalk impairing vectored lines.
Figure 1 illustrates the following situation:

- DSLAM A: this is a vectoring capable VDSL2 DSLAM with SLV architecture; all the lines originating from this DSLAM form a single vectored group(VG-A).

- DSLAM B: this is a VDSL2 DSLAM with BLV architecture, but not all its line cards are vectoring capable; the upper line card #1 is vectoring capable, and the lower line card #2 is a legacy VDSL2 line card.

- DSLAM C: this is a non VDSL2 DSLAM, for example supporting ADSL2 or SHDSL.

- VC-CPE: denotes a vectoring-capable CPE.

- VF-CPE: denotes a vectoring-friendly CPE.

- L-CPE: denotes a legacy VDSL2 CPE.

- O-CPE: denotes an other-than-VDSL2 CPE.

In Figure 1, six groups of lines are highlighted:
- Vectored groups (VG-A and VG-B): set of lines that terminate on a vectoring capable DSLAM/line card.

- Pre-coded groups (PG-A and PG-B): the subset of the vectored group whose lines terminate on a vectoring capable DSLAM/line card and a vectoring capable CPE.

- Legacy VDSL2 groups (LG-A and LG-B): lines in the vectored group originating from a vectoring capable DSLAM/line card that terminate on legacy CPEs..

- Vectoring friendly groups (VFG-A and VFG-B): lines in the vectored group originating from a vectoring capable DSLAM/line card that terminate on vectoring friendly CPEs.

Alien Line group (AG): other-than-VDSL2 lines that originate from DSLAM C.
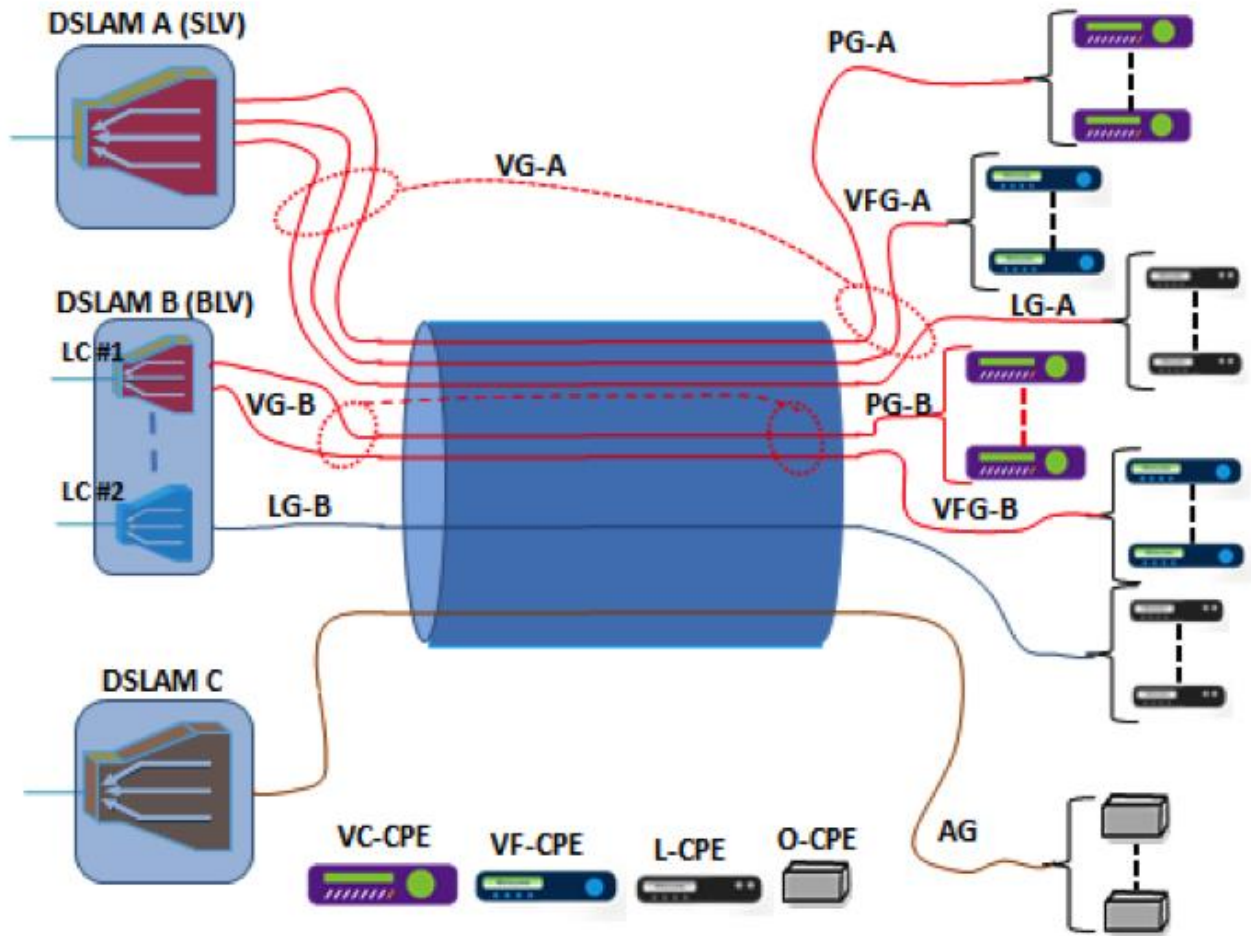
Figure 1 – Illustration of the various groups of lines that create crosstalk.

Figure 1 shows two examples of vectored groups: the vectored group originating from DSLAM A (VG-A), which is formed by the lines in LG-A, VFG-A, and PG-A; the vectored group originating from DSLAM B/Line Card #1 (VG-B). The figure also shows an example of two Pre-Coded Groups: the Pre-Coded Group whose lines originate from DSLAM A (PG-A) and DSLAM B/line card #1 (PG-B), which terminate on vectoring capable CPEs.

The following definitions are made:
- The crosstalk created by any line in VG-A is called "in-domain" self-FEXT, with respect to VG-A.

- The crosstalk created by any line originating from DSLAM B (VG-B and LG-B) is called "out-of-domain" self-FEXT, with respect to VG-A.

- The crosstalk created by any line originating from DSLAM A (VG-A and LG-A) is called "out-of-domain" self-FEXT, with respect to VG-B.

- The crosstalk created by any line in AG is called "alien" crosstalk, with respect to both VG-A and VG-B.

It has to be pointed out that the term "alien" has been used with different meanings in the recent literature; for example, with respect to VG-A, the term "alien lines" has been sometimes used to refer to the crosstalk originating by lines in LG-A, or by the lines in both VG-B and LG-B. This has caused some confusion which is here clarified. In this report, the definitions given are consistent with ITU-T documents (G.993.2 [1], G.993.5 [2]), BBF TR-197 [5], and ANSI Std. T1.417-2001 [8].

5.1   Cancelled and uncancelled crosstalk in vectoring

If lines in a Vectored group terminate on legacy CPEs (the L-group in Figure 1), then the in-domain Self-FEXT originating from these lines degrades the performance of all the lines in the vectored group because it cannot be cancelled by vectoring. In fact, only the in-domain Self-FEXT created by the lines in the Pre-Coded Group is cancelled by vectoring. The lines in the Vectored Group that terminate on legacy CPEs are outside of the Pre-Coded Group and operate as non-vectored VDSL2 lines.

The in-domain Self-FEXT originating within the Vectored Group but outside of the Pre-Coded Group can be eliminated if the legacy CPEs are upgraded to be at least vectoring-friendly, as specified in Annexes X (downstream only) and Y (full vectoring friendly, both downstream and upstream) of ITU-T G.993.2 [1]. For example, the in-domain Self-FEXT created by the lines in VFG-A of Figure 1 can be cancelled by vectoring and does not impact the performance of the lines in the PG-A.

The in-domain Self-FEXT originating within the Vectored Group but from the lines terminating on the legacy CPEs (LG-A in Figure 1) is not cancelled by vectoring and can degrade the performance of the lines in the pre-coded group (PG-A) on those tones where there is a high coupling between lines. In this case, DSM techniques can mitigate this in-domain self-FEXT as described in Sect. 6.1, for example by reducing the transmit power on the tones of lines characterized by high coupling. This in-domain Self-FEXT may also be cancelled (in the downstream only) by vendor specific/proprietary techniques.

If multiple Vectored Groups are formed when additional VDSL2 DSLAMs are present in the cabinet, then the out-of-domain Self-FEXT is present and this can further degrade the performance of the vectored lines. Since this Self-FEXT is generated outside of the vectored group, it cannot be canceled by vectoring regardless of the type of CPEs that terminate the out-of-domain lines. For example, the crosstalk generated by lines in VFG-B is canceled on lines in PG-B, but not by lines in PG-A. Techniques that can avoid the out-of-domain self-FEXT are cross-DSLAM vectoring (xDLV) and cable level vectoring (CLV), which are addressed in sections 6.2 and 6.3, respectively. DSM techniques that can be used to mitigate the effects of out-of-domain self-FEXT are described in Sect. 6.1.

Out-of-domain Self-FEXT also exits when multiple Board Level Vectoring (BLV) line cards i.e. multiple vectored groups originate from the same DSLAM. System Level Vectoring (SLV) can eliminate this out-of-domain crosstalk by avoiding creation of multiple vectored groups. Neither SLV nor xDLV can help cancel or mitigate alien crosstalk or in-domain crosstalk created within the vectored group but outside the pre-coded group, e.g. the crosstalk created by lines in LG-A cannot be cancelled by lines in PG-A.

Binder management is a method for mitigating out of domain noise by assigning customer lines within one or more binder groups to a vectoring group. In doing so, multiple binders can be assigned to multiple vectoring groups such that no two different vector groups transmit in the same binder, and crosstalk coupling for out of domain disturbers have the benefit of lower crosstalk due to physical separation within the cable. In board level vectoring deployments, this may require pre-deployment of all of the line cards within the DSLAM to be wired to disjoint sections of the cross connect which connect to separate binders, or rewiring of all of the customers as line cards are incrementally added. Binder management may be impractical since existing network operations practices and systems often do not enforce this type of mapping of DSLAM ports to cable pairs. If the network management system does support the mapping, customer churn could result in very inefficient utilization of cable pairs for DSLAM ports. Furthermore, out-of-domain self-FEXT would still exist due to crosstalk between binder groups.

Alien crosstalk includes crosstalk from any non-VDSL2 sources, and it is not cancelled by vectoring. However, DSM techniques can mitigate its effects.

5.2   Impact of uncancelled crosstalk on vectoring

There are several cases where the above types of uncancelled crosstalk are present.
- Cases where in-domain self-FEXT arises:

    o Gradual deployment, for example when the service on all the lines in a Vectored Group is not simultaneously upgraded and some lines may still be terminated on legacy CPEs, which are allowed to train without being placed in a vectoring friendly mode.

    o Customer's service choices, for example some customers may not want to change service or their CPE, or technological choices driven by the operator, or remote firmware update of CPEs to vectoring-friendly is not possible.

    o Regulatory Regime or Commercial Practices, when the legacy CPE cannot be upgraded because it is not owned by the same service provider deploying the vectored DSLAM.

- Cases where out-of-domain self-FEXT arises:

    o Vectoring implementation, for example when BLV is used and multiple Vectored Groups (one per line card) are created.

    o Deployment, for example when lines in a cable are terminated on multiple DSLAMs. Note that whether the additional DSLAMs are vectored or not, they still create out-of-domain self-FEXT to the Vectored Group terminated on the first DSLAM unless xDLV [14] or CLV are used.

    o Regulatory regime, for example when SLU is allowed and multiple operators own different DSLAMs (vectored or not) connected to the same cable.

- Cases where alien crosstalk arises:

     o   Presence of different services in the same cable, for example when the same cable carries ADSL/2/2plus or SHDSL in addition to VDSL2.

If the above cases occur, then vectoring cannot cancel all the crosstalk impairing the Pre-Coded Group. If this uncancelled crosstalk is not mitigated, it can severely impact the performance of the lines in the Pre-Coded Group as shown in Figure 2.

The black curve in Figure 2 shows the ideal (single line) performance that vectoring could attain. The red curve shows the 1% worst-case downstream data rates of twelve vectored lines (the Pre-Coded Group) in the presence of 12 VDSL2 legacy lines when no mitigation technique is used. The impact of the legacy lines on the Pre-Coded Group is the same regardless of whether they create in-domain or out-of-domain self-FEXT. Thus, the red curve would also be the performance that an xDLV or a CLV system would achieve in the presence of in-domain Self-FEXT generated when the 12 VDSL2 lines of the simulation scenario are terminated on legacy CPEs.
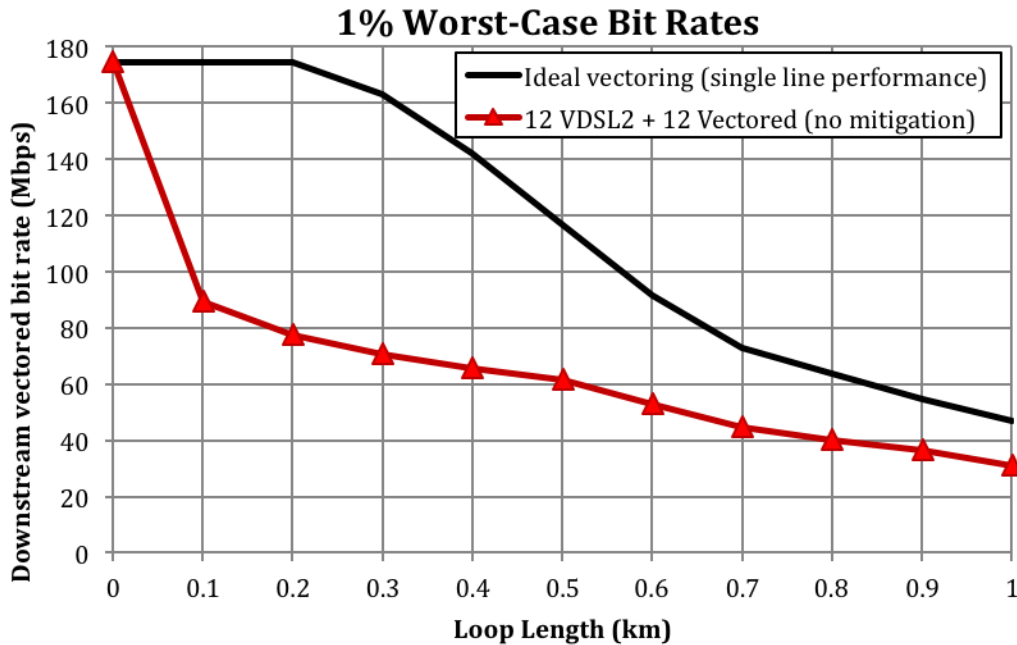
Figure 2 - Downstream vectored data rates, 1% worst case for: (a) Ideal Vectoring; (b) Mixed binder case of 12 legacy VDSL2 and 12 Vectored lines when no action to mitigate the uncancelled self-FEXT is undertaken. See Appendix VI for simulation assumptions.

When all self-FEXT is cancelled by vectoring, the remaining sources of interference (e.g. Radio Frequency Interference (RFI), impulse noise from electrical services in the home, interference from power line communications, etc.), referred to here as "alien noise," and alien crosstalk will become the dominant noise source(s) because they are not cancelled by vectoring. Mitigation of alien noise and alien crosstalk are not addressed in this report.

March 2014   25 of 101

6 Techniques for mitigating or avoiding the impact of uncancelled crosstalk on vectored lines

As clarified in Sect. 5, there are three kinds of uncancelled crosstalk:
1. in-domain Self-FEXT generated within the Vectored Group but outside of the Pre-Coded Group, i.e. lines terminated on downstream vectoring-friendly CPEs (upstream self-FEXT is not cancelled) or legacy CPEs (both downstream and upstream self-FEXT is not cancelled);
2. out-of-domain Self-FEXT;
3. alien crosstalk

The main techniques for *mitigating* the impact of some or all the types of uncancelled crosstalk on vectored lines are:

- Avoidance of multiple Vectored groups in the same cable (e.g. avoidance of sub-loop unbundling with vectoring)
- Dynamic Spectrum Management, Level 1 and Level 2, [4], [5], [6], [7]

The main techniques for avoiding completely the impact of out-of-domain self-FEXT are:

- Binder Management
- Cross-DSLAM Level Vectoring (xDLV) [14]
- Cable Level Vectoring (CLV)

In the next subsections, three techniques will be analyzed in detail: 1), the use of DSM techniques for mitigating all types of crosstalk is described in Section 6.1; 2) the use of Cross-DSLAM Level Vectoring (xDLV) for avoiding out-of-domain Self-FEXT; 3) the concept of Cable Level Vectoring (CLV) for avoiding out-of-domain Self-FEXT. Furthermore, the following key topics are developed per each technique:

- Description and terminology:

Describes the technique in its nature (e.g. technological, technical, network based practice, etc.), the relevant terminology associated to the technique and the way it allows benefits in mitigating and/or avoiding the impact of uncancelled crosstalk on the Vectored ones

- Actual benefit on Vectoring performance with respect to its ideal performances:

Qualitatively and quantitatively describes the benefits in mitigating and/or avoiding the impact of uncancelled crosstalk on the vectored lines; this topic is mainly focus on the theoretical benefit obtainable from the technique regardless of the network, operation and regulatory conditions described under the subsequent topics; results in support of the described benefits or reference to relevant literature may be reported.

- Network conditions and constraints to deploy this technique:

Describes the possible network conditions and constraints (e.g. architectural, system, hw/sw interoperability, scalability aspects) to be put in place (fully or partly) by the Operator(s) to effectively obtain (fully or partly) the described benefits; the sensitivity of the benefits to the degree of application of the described conditions and constraints is described as well.

- Operational conditions and constraints to operate this technique:

Describes the possible operational conditions and constraints (e.g. static/dynamic behaviour, configuration, profiling, scalability aspects) to be put in place (fully or partly) by the Operator(s) to

effectively obtain (fully or partly) the described benefits; the sensitivity of the benefits to the degree of application of the described conditions and constraints is described as well.

- Regulatory provisions needed to mandate a certain degree of coordination among operators: Describes the possible regulatory provisions (e.g. coordination and information exchange among different operators, spectral rules, rate limits, etc.) to be put in place by the Operator(s) to effectively obtain the described benefits; the sensitivity of the benefits to the degree of application of the described provisions is described as well.

- Foreseen availability of this technique or degree of maturity, if already available: Provides an indicative assessment of the state of the art of the technologies and /or solutions necessary to enable a certain technique within the framework of BBF anti-trust provisions.

- Technical considerations about the above constraints and trade-offs with respect to the usefulness of the technique: Provides an overall assessment of the technique in terms of benefits versus complexity for effectively deploying it.

6.1   Dynamic Spectrum Management (DSM)

The DSL environment is a multiuser environment where users interfere with each other by creating crosstalk (such an environment is known as an 'interference channel'). Dynamic Spectrum Management (DSM) is a framework under which techniques that reduce the effects of crosstalk are defined as solutions to optimization problems. Basically, DSM techniques allow each user to attempt to achieve their desired data rate while minimizing the disturbance to the other users sharing the same cable. DSM has been defined in both academic publication and the work of Standards Development Organizations, see references [3] through [7] and references therein. DSM techniques have been deployed in the field.

6.1.1   Description

Three levels of DSM are defined in the ATIS *DSM Technical Report*, [6] depending on what information is exploited and what level of control is exercised on the lines.

Below is an overview of the three DSM Levels as given in TR-197 [5].

DSM Level 1
"DSM Level 1 occurs where each line is monitored and configured independently to assess the line and noise conditions of that single line. DSM Level 1 is also known as Dynamic Line Management (DLM). DSM Level 1/DLM techniques typically improve line performance, that is rate/reach, stability or power consumption by manipulation of scalar parameters such as the Line Rate, Margin controls, FEC and INP parameters in DSL configuration profiles. DSM Level 1/DLM techniques currently are widely deployed in a number of Network Providers DSL networks worldwide" [5].

A good review of the benefits of DSM Level 1 can be found in [4], [6], [7].

DSM Level 2

"DSM Level 2 techniques optimize the spectrum and performance of multiple lines in a cable/binder as the lines are monitored to assess the noise and line conditions in the cable. In DSM Level 2 the interaction from one line (disturber) into another line (victim) is taken into account when configuring individual lines in the binder. The line that is configured can be either victim or disturber. The goal of DSM Level 2 optimization is coordinated optimization of all lines in the cable/binder. DSM Level 2 techniques can utilize all DSM Level 1/DLM techniques while also using coordinated modification of the spectrum parameters on DSL lines in order to achieve the benefits across multiple lines in a cable/binder. DSM Level 2 is currently emerging into deployment" [5].

A good review of DSM Level 2 and its benefits can be found in [6]. A more formal overview of DSM Level 2 optimization algorithms can be found in the recent tutorial [9], where most of known DSM Level 2 algorithms are compared in terms of performance (achievable data rate) and computational complexity.

DSM Level 3
"DSM Level 3 is the management of Vectoring techniques, such as those described in G.993.5 [2], to cancel crosstalk between DSL lines" [5].

6.1.2 Benefits on Vectoring performance

Over the last decade several studies reported in the literature have confirmed that DSM can mitigate the effects of uncancelled crosstalk and allow lines in a vector group to retain some of the vectoring gains that would be otherwise lost if no mitigation technique were used (see [9] and references therein). Specifically, these studies report that mitigation techniques based on DSM Level 1 (rate limiting) and DSM Level 2 (spectrum balancing) can be used for mixed deployments (mixture of vectored and non-vectored lines and of multiple vector groups) in both upstream and downstream directions.

DSM Levels 1 and 2 can mitigate the impact of uncancelled crosstalk by implementing trade-offs between the speeds supported by the vectored and non-vectored lines or between multiple vector groups. These DSM techniques can mitigate the reduction of vectoring gains due to uncancelled crosstalk and these trade-offs generally require that the data rates achievable by one set of lines be reduced to increase the data rates achievable by the other set. This kind of trade-off in not unique to DSM, Power Back Off (UPBO/DPBO) is another example where a trade-off between the data rates of lines is introduced.

DSM techniques use various optimization procedures to establish trade-offs between lines. The entity that actuates the DSM techniques is commonly referred to as a Spectrum Management Center (SMC). SMCs that manage multiple DSL networks in a single physical plant can operate autonomously, where *a priori* inter-network politeness rules are pre-defined. They can also be tightly coupled, where the various SMCs exchange information about their respective network's performance. Centralized DSM techniques are also possible, wherein a single SMC manages multiple networks – see Annex A of [4]. An SMC, as defined in the ATIS DSM TR [6], is a centralized system to which line information and spectrum control information is reported and which implements DSM techniques by finding the best configuration of the line. In fully autonomous DSM, the DSL equipment could be allowed to autonomously adapt to its specific

crosstalk environment, however it would not be aware of any impact that their operation might have on other lines.

Centralized systems using a SMC generally exhibit better performance at the cost of increased complexity and line coordination. On the other hand, distributed systems can operate fully autonomously without the need of explicit message passing between SMCs, but often sacrifice achieving configurations that best optimize performance.

A full description of these DSM techniques is out of the scope of this Report. The reader can find such details, their achievable performance, and complexity requirements in various standards and the open literature. [5], [6], [7], [9], [10], [12].

6.1.3 Network conditions and possible constraints

The effect of crosstalk is not uniform across DSLs, and depends on a large number of factors such as loop length, frequencies used, cable geometry and the density of DSLs in a cable binder. As a result, the data rate performance of DSL can have a very wide variation in the field, especially for the shorter loops, where crosstalk dominates. As lines exhibit time-variation (channel faults, noises coming and going, etc.) the best practice would be to decide standard type and profile through long-term observations and thus adapt choices to achieve better long-term target stability and data rate.

Vectored DSL also expands the capabilities of loop re-profiling. Improved practices for re-profiling vectored lines lead to improved overall network operation and a reduction of line instability issues, which could otherwise generate customer calls and consequent technician dispatches. Such practices include prioritizing lines for allocating vectoring resources, configuring transmit spectra and power to enable coexistence among vectored and non-vectored lines, and management of non-crosstalk noise sources. These practices always require a dynamic re-profiling capability using historical data, and sometimes require access to multiple lines since optimization cannot be performed on a line-by-line basis. The proper implementation of DSM algorithms (either used as a technique for mitigating uncancelled crosstalk, or for other goals) requires leveraging the knowledge of this historical data that can be collected by the SMC.

In the case of distributed DSM, the various SMCs collect historical data and analyze it in the context of network-wide operational objectives. There can be differing degrees of coupling among the various SMCs. In the most loosely coupled case, the SMCs do not communicate with one another, but share a common set of rules with respect to PSDs, rate limits, etc. With distributed DSM, transmit parameters, such as the transmit power or transmit PSD, are determined for each line within constraints such as the limit PSD mask.

In a much more tightly coupled case, the SMCs also exchange quasi real-time information regarding Hlog, Xlin and other parameters. This is really a physically distributed implementation of a logically centralized SMC, which is considered to be a "de-centralized" implementation. This inter-SMC information exchange requires a higher level of Service Provider ("SP") coordination than distributed DSM. Currently there is no standard for the protocol, rules and data exchange among SMCs so that de-centralized DSM implementations are practically implemented via an agreed set of pre-defined rules in a multi-operator environment at the cost of reduced performance with respect to the case where SMCs communicate with each other. Furthermore, SP coordination is needed when

updating the rule-set over time. At the time of publication, the UK NICC DSL Technical Group had commenced a study (ND1518) on requirements related to data sharing between SPs when decentralized SMCs are utilized.

An example is to start up within transmit PSD or power limits that are set to bound the estimated crosstalk impact on other lines, and then self-optimize transmit parameters to maximize performance and minimize crosstalk; this should not require more than a couple of retrains per line. This operation may be limited by the number of available DSLAM profiles in some cases, but newer DSLAMs support many profiles. No new chipset capability or particular network condition is necessary for most DSM algorithms of practical interest, but in some cases there are bit-loading algorithms (such as Multi-Level Water Filling) that would require applying a firmware upgrade to existing transceivers. The OSS provisioning and assurance processes need to integrate the autonomous line reconfiguration to avoid conflicts with the default set of DSL profiles.

6.1.3.1 Single SP scenario

Centralized DSM, with a single SMC, can be easily implemented when there is a single Service Provider (SP) managing the copper network, and the network planning and deployment of the SMC is less onerous with a single Service Provider (SP). As shown in Figure 3, a single SP is exclusively responsible for the use of copper twisted pairs within an area, which is the case when there is no physical unbundling. In this case, all lines (both vectored and non-vectored) are controlled by a single entity. The management system is also under the full control of the single SP and DSM use is straightforward as the SMC has a full view of the network.
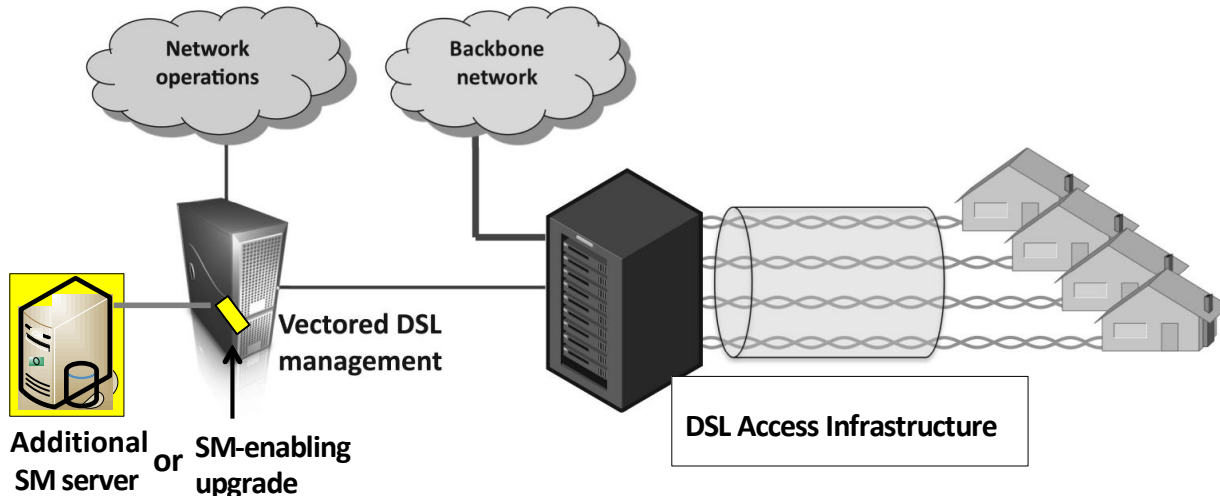
**Figure 3 - Centralized DSM deployment in a Vectored DSL scenario.
Case of single SP managing a cable.**

Centralized DSM requires the introduction of an SMC server in the network. This can be achieved via a Spectrum Management enabling upgrade of the processing/management features of existing network management systems or via an additional dedicated SMC. Depending on SMC implementation, there may be cases where scalability issues arise (i.e. depending on the maximum number of supported DSLAMs and/or lines per SMC with respect to the total number of DSLAMs) so that it may be necessary to deploy multiple SMCs. Using dedicated SMC servers is just one

possibility; an alternative is to utilize shared computing resources in virtual environments or in a cloud computing system.

In this scenario of a single operator it is still possible that coexistence issues and the need for mitigation techniques arise. Classical examples are when the SP carries out a gradual deployment of vectoring, when there are technical limitations that limit the size of the vector group, or when there are lines terminating on legacy CPEs.

For example, when multiple DSLAMs are present in a cabinet, such as in serving areas where the cabinet serves a very large number of customers, and the SP decides to upgrade to vectoring only a subset of the DSLAMs, then the lines terminating on the legacy VDSL2 DSLAMs that share the same cable where vectored lines are present will create out-of-domain crosstalk that will degrade the performance of the vectored lines. Upgrading all DSLAM in the cabinet to xDLV will avoid out of domain crosstalk, however when this is not accomplished implementation of DSM may be beneficial to the vectored lines as the impact of this out-of-domain crosstalk would be mitigated.

In another example, there may be practical limitations on complexity that may allow only partial vectoring (not all lines in the vector group are cancelled, not all crosstalk of a cancelled line is removed completely by the vectoring engine) thus leaving some lines in the access suffering from crosstalk. DSM could be used to mitigate the effects of this uncancelled crosstalk.

6.1.3.2 Multiple SPs scenario

In the case where requirements for physical unbundling are present and two SPs share the copper twisted pairs, and choose to use separate DSLAMs and separate spectrum management systems, these systems do not have a full view of the entirety of twisted pairs in the network. For this scenario, upgrading all vectored DSLAMs in the cabinet to support xDLV will avoid out-of-domain crosstalk and allow full vectoring gains (if no legacy CPE creating uncancelled in-domain crosstalk is present). However, when this is not accomplished, implementation of DSM may be beneficial to the vectored lines as the impact of this out-of-domain crosstalk would be mitigated.

Having separate spectrum management systems limits diagnostics and reprofiling capabilities, when compared to the case of a single system managing all lines in the network (see Figure 4). If an SP already has an SMC, the complexity of adding a DSM functionality for managing coexistence and mitigation of uncancelled crosstalk is negligible. In case an SP does not already have an SMC, then an SMC functionality must be added to existing network management systems, or cloud-based resources can also be used.

With separate management systems, some performance loss is expected for vectored DSL, because of a reduced ability to coordinate vectored and non-vectored VDSL2 systems of provider A and provider B. Provider A's lines create out-of-domain self-FEXT to the lines of provider B and vice-versa. Still, the respective management systems can be very effective with 'policing' actions, such as detecting whether the systems of provider A are inadvertently causing disruption to the systems of provider B, and also for ensuring that each network is polite overall to the other which can substantially reduce inter-system crosstalk. Such policing can allow provider B to request provider A to correct the situation. In this case, distributed DSM techniques

can be very valuable as no communication is required between the two SPs except for sharing a common set of rules with respect to PSDs, rate limits, etc. If the various operators decide to share certain network information with each other, e.g. via explicit inter-SMC message passing or a centralized database, network performance can be improved even when multiple SMCs are present as information sharing allows each SMC to have a more complete view of the physical network.
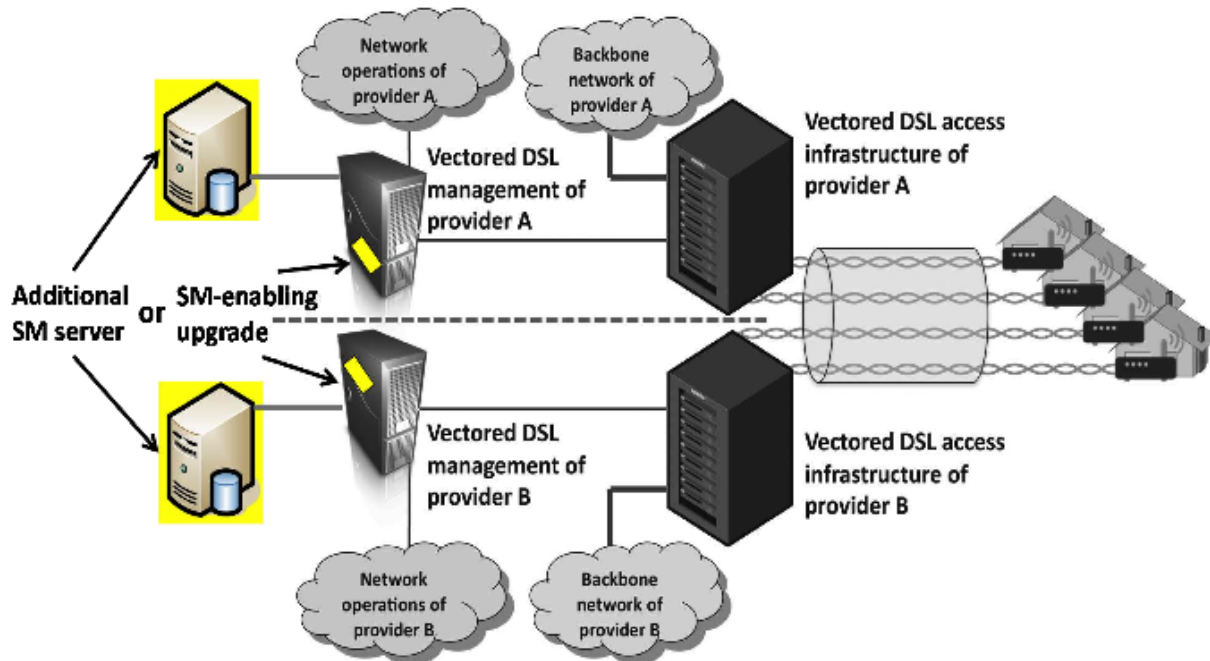
Figure 4 - Distributed DSM deployment in a Vectored DSL scenario (independent SM servers). Case of multiple SPs managing a cable.

A second unbundled architecture that results in better efficiencies than the previous example occurs when operators agree to share a single SMC with full view of the network rather than having multiple SMCs exchanging information – see Figure 5. In this case, the vectored DSL access infrastructure is shared among the providers and is managed by a single management system enabling monitoring and control of both vectored and non-vectored DSL lines. This single management system could also be administered by a trusted third party to ensure fairness among operators.

The scenario depicted in Figure 5 also applies to the case when virtual unbundling (sometimes also known as bit-stream access) is used, i.e. where a single SP (sometimes named the 'network access provider') has physical control of the network and competitive SPs deploy their services over the network access provider's physical network.

As shown in Figure 5, each of the SPs sharing the cabinet area still has access to the management functions but with appropriate restrictions to prevent disclosure of sensitive information of other providers or to affect the services offered by other providers. The benefits from diagnostics and reprofiling can be maximized to the same extent as with only a single provider – and this is because the network access provider has a complete view and control of the physical network and the competitive SP's agree to share certain line data. At the same time, each provider can define and manage its services independently, subject to overall rules on fairness.

This deployment scenario requires an SMC shared among different operators or the deployment of a shared network management server with SMC support. Coordination among the involved SPs is needed in the different deployment phases depending on the degree of ownership and/or control of the shared SMC and also based on the existing SP agreements about network operations and maintenance procedures.
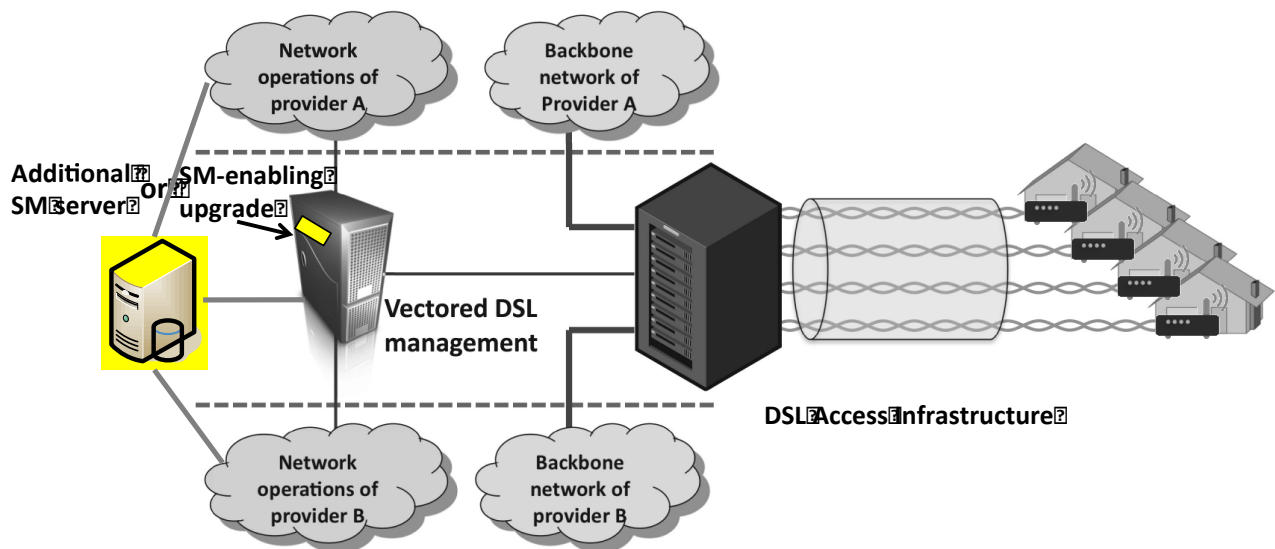
Figure 5 - Centralized DSM deployment in a Vectored DSL scenario. Case of multiple SPs with shared cable management.

Note: the DSLAM icon may represent a single DSLAM or also multiple ones.

6.1.4 Operational conditions and possible constraints

Along with the network and management architectural upgrades described in the previous section, the effective operation of DSM techniques depends also on the associated operational procedures. Essentially these can be divided into two categories:
- Conventional DSL management functionalities
  - monitoring of DSL parameters of the lines in the cable
  - reprofiling of the DSL lines

- DSM-related functionalities
  - computation of optimal spectrum profiles via the DSM algorithms
  - for Centralized DSM in virtual unbundling, coordination of SPs sharing data
  - for decentralized DSM, coordination of inter-SMC communication

Monitoring of DSL parameters and DSL line data availability:
The DSL data necessary as input to DSM algorithms needs to be provided via the DSL monitoring functionalities and related activities normally performed for network operation. The overall load and scalability of these data collection activities need to be assessed when planning and executing the deployment of the DSM techniques for the following reasons:

- the optimality of the spectrum profiles determined via the DSM algorithms depends on the availability of DSL data information, i.e. what parameters are available and how frequently they are collected;

- the efficacy of the DSM technique is influenced by the availability of DSL data for all the lines in the binder/cable.

Reprofiling of the DSL lines
The DSM related line reconfiguration is not expected to be very frequent.

From an Operation, Administration, and Maintenance (OAM) point of view, the DSM related line reconfiguration process needs to be integrated with the existing OSS processes that control the line profiling associated with existing network operation procedures such as service provisioning, service assurance, and line troubleshooting.

Further to the discussion above, the references [3], [4], [6] through [9]describe some of these operational aspects.

Computation of optimal spectrum profiles

Along with the architectural and feature allocation considerations reported in the previous section, the frequency of the computation of the optimal spectrum profiles is a relevant element of operations. The spectrum profiles determined via DSM are optimal from a maximization of the weighted sum-rate point of view and depend on the electromagnetic characteristics of a given cable and the filling/distribution of DSL systems in the cable itself. Hence it can be said that the set of chosen profiles to be applied to a set of lines to mitigate their impact on those of a Vectored group typically varies in a quasi-static way as much as the SMC monitoring and computation algorithms need to routinely control the current cable conditions and assess the degree of optimality of the applied profiles.

As described above DSM is a constraint-based optimization technique which computes a set of profiles to be applied to each line under DSM management, i.e. an n-tuple of DSL spectrum parameters values per line. Operationally this means that the set of "optimal" profiles is likely to be different on a site-by-site basis. At each site the profile instances of the "optimal" set will likely differ on a line-by-line basis.

Operators typically define a set of network-wide default profiles at the time of activation of a given DSL line whose rate, spectrum and DSL link parameters take default values depending on the service characteristics (e.g. bitrate, latency, INP), on the site characteristics (e.g. DPBO, UPBO) and other aspects. Similarly other network-wide profiles are defined for service assurance and other network operation purposes.

The network-wide set of profiles and those determined using DSM algorithms need to be managed and stored in the network databases, in the network management systems and on the DSLAMs themselves.

Depending on the SP strategies for profile handling and network management, it may be necessary to find a good trade-off between optimal configuration and a manageable number of profiles. However, it is desirable to avoid limiting too much the number of network-wide profiles as this would require introducing performance trade-offs anyway and independently of using DSM as a mitigation technique.

Coordination of SPs for Centralized DSM in an unbundled environment

In this scenario, a shared management platform with an associated SMC is put in place by the involved SPs who agree to share network data to achieve a complete view of the network. This requires SP coordination of operational activities that are related to the sharing of the above resources.

Coordination of operators in de-centralized implementations of DSM
Decentralized DSM in a multi-operator environment requires operational coordination among the management domains of different SPs for the exchange of the rules and line profiling constraints, up to the extent of exchanging quasi real-time information as discussed previously. However, the lack of standards for such an architecture limits the practicality of de-centralized implementations of Centralized DSM.

6.1.5 Possible regulatory requirements

Distributed DSM techniques do not necessarily raise regulatory issues as SMCs can operate autonomously without the need of message passing between users or between SMCs.

In the case centralized DSM where SMCs exchange information, regulatory issues may arise in some cases. For example, for the case of the scenario depicted by Figure 5, regulations as to what information the various SPs share about their network may be required. Regulation, if needed, might specify what information is shared not how it is shared or what SP should do. Also, regulatory requirements are only necessary if the parties cannot come to an agreement among themselves on how to coordinate their systems.

Furthermore, performance trade-offs and allowable spectral impact levels may need to be selected by the industry and regulators. Mitigation of uncancelled crosstalk on vectored lines implies controlling the rate and power of the non-vectored lines (or the other vectored groups) to limit the degradation suffered by the vectored lines. These trade-offs generally require that the data rates achievable by one vector group are reduced to increase the data rates achievable by the other group (or group of non-vectored lines). In some cases, this trade-off can entail a significant reduction of the attainable peak speeds of one vector group in order to maintain good vectoring gains in the other vector group.

If lines are controlled by multiple operators this may require regulatory provisions on how operators exchange information or on how trade-offs among the peak speeds supported by each operator are selected; if not regulatory provisions, at least agreements between operators are needed. However, if the non-vectored lines (or the additional vector groups) are properly managed, the impact on the vectored lines can be limited and made predictable so that various levels of coexistence can be achieved on the basis of the trade-offs agreed upon by operators or set by regulators.

The effectiveness and acceptance of these trade-offs will also depend on the definition of fairness for the competitive environment that physical loop unbundling seeks to enable. In certain jurisdictions and service environments, the requirement of data rate trade-offs may not be compatible with the existing competitive dynamics.

6.1.6 Foreseen availability

There are many well-known DSM techniques, which are implemented as software and DSM systems have been deployed by network operators. The capability of handling in certain situations the coexistence of multiple vector groups as well as vectored and non-vectored lines already appears to be a feature available in some deployed DSM solutions.

6.1.7 Technical considerations: trade-offs vs. usefulness

6.1.7.1 Case of Multiple Vector Groups

Coexistence is challenging for the two-vector group case, and two vector groups with strong crosstalk between them cannot both be run at the very high speeds of vectored lines. If the data rate

of one of the vectored groups is maximized, then the data rate of the other vectored group will be at VDSL2 speeds. In this case, a better approach is to use xDLV across all vectored DSLAMs so that a single xDLV vectored group is present and out-of-domain crosstalk is eliminated, thus allowing full vectoring gains to all lines in the xDLV group (if no legacy CPE creating uncancelled in-domain crosstalk is present).

Some illustrative performance results of deployments with multiple vector groups are provided in Appendix II and VI. For example, if the data rates of both vector groups are maximized, then in scenarios with two vectored groups:

- Vectored lines suffer ~35% degradation from realistically achievable rates, while still maintaining ~30% or more improvement over dense VDSL2 deployments.

- If all vectored lines are set to operate at a 100 Mbps target, then vectored lines suffer minimal degradation up to 200m and around 30% degradation on longer lines, while still maintaining ~30% improvement over dense VDSL2.

6.1.7.2 Case of Mixed vectored VDSL2 and Non-Vectored VDSL2

In mixed vectored/non-vectored scenarios, DSM can mitigate the impact of uncancelled crosstalk by implementing trade-offs between the speeds supported by the vectored and the non-vectored lines. DSM provides a framework for preserving most of vectoring gains while limiting non-vectored lines to data rates that are often typical of legacy non-vectored VDSL2 service levels. Some illustrative performance results of deployments with mixed vectored and non-vectored VDSL2 lines are provided in Appendix I-VI.

- Vectored lines suffer < 15% data rate degradation from realistically achievable rates, while non-vectored lines are capped at speeds typical of legacy non-vectored VDSL2 services (40-50 Mbps).
- If all vectored lines are set to operate at a 100 Mbps target, then vectored lines suffer no degradation at all up to 500 m while non-vectored lines can be capped at even higher speeds (40-95 Mbps).

For more data, see Appendices I-VI.

6.1.8 Summary of informative appendices analyzing DSM as a technique

Appendixes I though VI provide simulations of use of DSM techniques to mitigate uncancelled crosstalk in a vectored environment. An overview of the contents of these six appendixes is provided in an introductory section placed before these appendixes.

6.2 Cross-DSLAM Level Vectoring (xDLV)

6.2.1 Description

Cross-DSLAM level vectoring [14] is a technique for avoiding the impact of out-of-domain self-FEXT and is based on the coordination among different vectored DSLAMs to perform the Vectoring functions in a distributed fashion via a master-slave architecture or distributed processing.

This technology can be seen as an evolution of the so called System Level Vectoring (SLV) which operates over a single vectored DSLAM and it is conceived to extend its operation and, in perspective, its performance advantages to a multi-equipment scenario.

It relies on the sharing of a Vectoring Engine resource among the vectored DSLAMs and on the exchange of data to/from this processing resource and all the vectored DSLAMs that have to be coordinated. This technology can be seen as an evolution of System Level Vectoring (SLV) and DSLAM coordination is achieved by connecting a high-speed cable between the vectored DSLAMs that participate in xDLV [14] for exchanging information necessary for vectoring lines on the interconnected DSLAMs, including clock, symbol, and crosstalk information.

6.2.2 Benefits on Vectoring performance

Cross-DSLAM Level Vectoring allows extending the vectoring capabilities of self-FEXT cancellation beyond the boundaries of a single DSLAM chassis and to a larger vector group and potentially to collocated equipment managed by different SPs. Basically, xDLV brings in-domain all those lines that are out-of-domain, de facto eliminating out-of-domain self-FEXT. However, xDLV does not cancel the in-domain self-FEXT created by lines terminated on legacy CPEs. This residual in-domain self-FEXT can impact the performance of the Pre-Coded Group so that it would have to be mitigated with some other mitigation technique or eliminated upgrading all legacy CPEs to vectoring-friendly or vectoring-capable CPEs.

Since it is an extension of the vectoring group size to potentially all the lines in a cable, xDLV's benefit to vectoring performance is optimal if and only if there are no lines terminated on legacy CPEs. In the case some lines terminate on legacy CPEs, then these lines would have to be managed in order to restore full vectoring gain.

6.2.3 Network conditions and possible constraints

Cross-DSLAM Level Vectoring has no specific network constraints related to the technology in itself, beyond the installation and logistic activities related to deploying or upgrading multiple chassis and interconnecting them. Among the installation constraints there is the type of interconnecting cable (i.e. metallic or optic) and its maximum length, and equipment grounding issues that in turn influences the maximum distance among "collocated" chassis.

Today Cross-DSLAM Level Vectoring requires that equipment of the same vendor are deployed at a given cabinet location where the xDLV coordination is required. In some cases this co-location may not be trivial. Currently there is no standardization development for the cross-chassis data exchange/protocol to allow multi-vendor interoperability. Instead the deployed equipment technology may vary for each cabinet location. This requires coordination among the SPs deploying

xDLV capable equipment as related to the crosstalk cancellation capabilities of the xDLV vector group.

6.2.4 Operational conditions and possible constraints

Cross-DSLAM Level Vectoring requires specific activities for managing and operating multiple chassis with a XDLV type of coordination.

In a multi-operator environment, xDLV requires the coexisting SPs make some common choices on a site basis:
   a) vectoring both in upstream and downstream directions
   b) same bandplan and PBO settings
   c) adoption of compatible equipment releases
   d) same fallback policies (for example with respect to non-vectoring friendly CPEs)
   e) SPs' coordination on maintenance windows and release changes

All the other VDSL2 line settings remain under the full control of each SP.

6.2.5 Possible regulatory requirements

As described above xDLV encompasses a certain degree of coordination among SPs with collocated xDLV chassis. This does not require any regulatory provisions to be put in place in order to facilitate such coordination.

6.2.6 Foreseen availability

Cross-DSLAM Level Vectoring is currently under research and development by the industry and it is planned in the roadmap of a number of equipment manufacturers. There are currently no industry standards to support interoperability of this functionality among products of different manufacturers

6.2.6.1 Case of Multiple Vector Groups

If a subset of the vectored DSLAMs sharing the same cable does not support xDLV or if some operators do not agree to support xDLV, then multiple vectored groups are existing present in the cable. These vectored groups create out-of-domain crosstalk onto each other, reducing the vectoring gain of the xDLV group. This deployment model is not advisable as it leads to sub-optimal performance. For example, for the scenario depicted by Figure 2, the performance achieved by the xDLV vectored group in the presence of unmitigated out-of-domain crosstalk is given by the red curve.

Mitigation techniques (DSM) can be used to reduce the impact of this out-of-domain crosstalk, although it is often not possible to restore vectoring gains to all the vectored groups. For optimal performance, xDLV must be deployed on all the vectored DSLAMs in a cabinet that share the same cable so that a single xDLV vectored group is formed and out-of-domain crosstalk is eliminated.

March 2014   39 of 101

6.2.6.2 Case of Mixed vectored VDSL2 and Non-Vectored VDSL2

Even if all the vectored DSLAMs sharing the same cable support xDLV, in-domain crosstalk can still be present when there are lines terminating on legacy CPEs. This in-domain crosstalk is not cancelled by vectoring and the vectoring gains of the xDLV group may be reduced. For example, for the scenario depicted by Figure 2, the performance achieved by the xDLV vectored group in the presence of uncancelled in-domain crosstalk is given by the red curve.
Mitigation techniques (DSM, upgrade to vectoring friendly CPEs) can be used successfully to reduce the impact of this in-domain crosstalk and preserve in most cases the vectoring gains of the whole xDLV vectored group. The deployment of these techniques introduces some additional operational conditions and possible constraints which are addressed in detail in Section 6.1.

xDLV can be used to remove out-of-domain self-FEXT by extending the Vectored Group to span across multiple DSLAMs, thus yielding to potentially optimal performance. CLV may be used in future for the same purpose, even if its early stage of development makes availability of mature products still unclear. xDLV and CLV do not cancel the in-domain self-FEXT due to eventual presence of legacy CPEs. The deployment of xDLV has the benefit of a wider crosstalk cancellation but introduces some additional operational conditions and possible constraints with regard to those of BLV or SLV and these are addressed in detail in Section 6.2 and 6.3.

6.3 Cable Level Vectoring

6.3.1 Description

Cable level vectoring (CLV), a technique that is currently under research and development, is a technique for avoiding the impact of out-of-domain self-FEXT. In the context of this section, the term vectoring is used to describe a proprietary, and not ITU-T Recommendation G.993.5 compliant, technology for reduction of crosstalk levels performed across all pairs in the cable.
. CLV does not require vectored DSLAMs or DSLAM coordination, but requires the deployment of special CLV equipment (CLVE) at the cabinet where all the VDSL2 lines from a particular cable are terminated. The CLVE performs the operation of vectoring by allocating all the lines in the cable to a single Vectored Group even though lines in that cable may be terminated on two or more DSLAMs as shown in Figure 1. While vectoring functionality is not necessary in the DSLAMs connected to the CLVE, the DSLAMs may also include vectoring functionality. With CLV, out-of-domain self-FEXT crosstalk is avoided for all lines in the cable connected to the CLVE.

March 2014   40 of 101
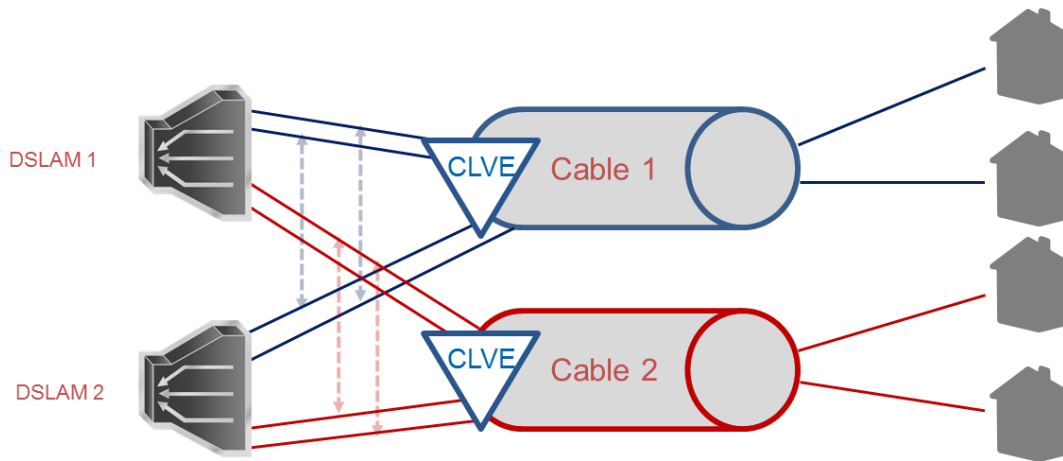
Figure 6 - CLV architecture for multiple DSLAMs example

6.3.2 Benefits on Vectoring Performance

CLV terminates all the lines in a cable hence all the lines benefit vectoring performance gain regardless of the number of DSLAMs deployed at the cabinet as out-of-domain self-FEXT crosstalk is avoided. It is not necessary to deploy vectored DSLAMs to enjoy the vectoring performance gain hence legacy deployed DSLAM can be used with the addition of CLVE. In case some lines are terminated on legacy CPEs (neither vectoring capable nor vectoring friendly) then the in-domain self-FEXT generated outside of the Pre-Coded Group cannot be cancelled by vectoring. In this case, CLV distributed nature will reduce the complexity of other mitigation techniques (e.g. DSM,…) so these lines will not create harmful in-domain self-FEXT crosstalk. In fact, a crosstalk mitigation technique like DSM will need to handle the impact of the legacy CPEs at the cable level rather than handling multiple DSLAMs that may be also owned by different operators and the necessity of identifying neighboring lines in a cable across different DSLAMs can also be averted.

6.3.3 Network conditions and possible constraints

The deployment of CLV as a technique requires the introduction of an additional dedicated equipment (CLVE) in the network per vectored cable and additional cross-connection panels. This requires the availability of adequate space, powering and heat dissipation capability.

The CLVE should be collocated with the DSLAMs. If the distance between the DSLAM and the CLVE is excessive, than crosstalk in the Tie cable might limit the vectoring performance gain. If multiple SPs are involved the plans and execution of the CLVE deployment and maintenance require coordination among the SPs.

The CLVE should be able to support a number of ports up to the cable size.

6.3.4 Operational conditions and possible constraints

The deployment of CLV as a technique requires a change to the typical management architecture and the deployment of additional systems for the management of CLVEs.

The CLVE could be managed by messages from a management system that are interpreted and forwarded by the DSLAMs or alternatively the CLVE could be managed by messages directly addressed to the CLVE from the management system.

As mentioned above, the CLV architecture requires the support of two-segment functionality. In the first alternative in-band management messages are interpreted and forwarded by the DSLAM. In the second alternative out-of-band management messages are directly interpreted by the CLVE. In the latter case, regulatory definitions might be required to determine information sharing to allow mutual operation of different service providers.

Additional OAM activities are required as related to the CLV equipment and a number of these activities need to be associated with the DSLAMs connected to given CLVE, e.g.: line profiling, monitoring and data collection, service assurance and troubleshooting.
The degree of management coordination increases at a multi-SPs site where one or more CLVE are a shared resource.

The addition of CLVE will need to be addressed in the network management systems for inventory management and service provisioning.

6.3.5 Possible regulatory requirements

In a multi-SP environment there may be regulatory requirements related to the deployment of shared CLVEs and the need of SPs coordination.

As said above regulatory definitions might be required about information sharing among different service providers.

6.3.6 Foreseen availability

CLV is currently under research and development.

6.3.7 Case of Multiple Vector Groups

The CLVE performs the operation of vectoring by allocating all the lines in the cable to a single vectored group hence it is indifferent to the number of DSLAMs and originated vectoring groups.

March 2014   42 of 101

7   Summary of mitigation techniques

In Sections 6.1through 6.3, three techniques for handling uncancelled crosstalk have been discussed in detail: DSM-based techniques, xDLV, and CLV.

These three techniques add to the other techniques mentioned in this Technical Report: availability of SLV vectoring implementations, vectoring-friendly CPEs, avoidance of SLU at sites where vectoring is deployed, and binder management.

xDLV can be used to remove out-of-domain self-FEXT by extending the Vectored Group to span across multiple DSLAMs, thus yielding to potentially optimal performance. CLV may be used in the future for the same purpose, although CLV is currently under research and development and availability of mature products is still unclear. xDLV, and CLV do not cancel the in-domain self-FEXT due to eventual presence of legacy CPEs. No proposal for xDLV standardization has been made as of the date of publication. The deployment of xDLV aside with the benefit of a wider crosstalk cancellation introduces some additional operational conditions and possible constraints with regard to those of BLV or SLV, and these are addressed in detail in Section 6.2.

DSM is an effective technique for the mitigation of all types of uncancelled crosstalk, i.e. in-domain/out-of-domain self-FEXT and alien crosstalk. In the case of mixed vectored VDSL2/non-vectored VDSL2 scenarios, applying DSM-based management to both the vectored and non-vectored lines preserves most of vectoring gains while non-vectored lines are capped at speeds typical of legacy non-vectored VDSL2 services (40-50 Mbps). The speed cap imposed on the legacy non-vectored VDSL2 lines can even be higher than 40-50 Mbps if the vectored lines are capped to a maximum rate of 100 Mbps. This applies both to the case that the non-vectored VDSL2 lines create in-domain or out-of-domain self-FEXT, i.e. regardless of whether they belong to the same Vectored Group or are terminated on different DSLAMs.

In the case where out-of-domain self-FEXT is created by multiple Vectored Groups, then the multiple Vectored Groups cannot all operate at vectoring speed. Thus, DSM-based techniques can allow the lines of only one Vectored Group to operate at vectoring speeds whereas the lines in the other Vectored Groups would have to be capped to speeds typical of legacy non-vectored VDSL2 services. In this case, other solutions like xDLV and CLV – when available on the market – could be more effective as they would eliminate out-of-domain self-FEXT.

Along with the described performance trade-offs, the deployment of DSM as a technique in a brownfield scenario introduces network and operational conditions and possible constraints in the processes and procedures required to effectively run such technique and these are addressed in detail in Section. 6.1.

Avoidance of physical SLU could also be an effective solution depending on the specific network environment. However, even without physical SLU there are cases where uncancelled crosstalk would still be present. For example: when legacy CPEs are present and cannot be upgraded, when a single operator may have deployed two or more DSLAMs at the cabinet and upgrades only one of them to vectoring. In these cases, DSM-based techniques are an effective interim solution.

March 2014   43 of 101

Deployment of Vectored DSL raises new operational issues for the network operator. If not mitigated, uncancelled crosstalk could decrease the vectored lines performances. Operators have several tools at their disposal for mitigating the impact of uncancelled crosstalk. Although none of these tools by itself can completely remove all types of uncancelled crosstalk, using a combination of SLV, xDLV, vectoring friendly CPEs, and DSM-based management can maintain vectoring gains in most deployments with or without physical SLU. In the future, solutions currently in research and development such as CLV could further improve the set of tools.

March 2014   44 of 101

Overview of Appendixes

Appendix I shows rate regions of vectored and non-vectored mixtures achieved by the optimal spectrum balancing (OSB) DSM technique. These regions show good trade-offs, and in particular the non-vectored lines can transmit reasonable bit rates for legacy VDSL2 services with very little impact on the vectored group, if the non-vectored lines are properly managed.

Appendix II shows rate regions of deployments of multiple vector groups achieved by OSB. This concludes that the trade-offs for the multiple vectored group case are not as good as they are for the mixed vectored / non-vectored case, but the tradeoff can still be controlled with DSM.

Appendix III analyzes distributed user scenarios, with vectored and non-vectored lines of differing lengths in the same cables. These scenarios generally have similar or better performance at the same loop length than scenarios with collocated subscribers and lines all of equal length, although it is difficult to directly compare distributed scenarios and equal length scenarios. It is concluded that the equal length scenarios are good representatives since they are both conceptually simple and represent a type of worst-case.

Appendix IV compares the iterative waterfilling (IWF) DSM technique to using flat power backoff. While IWF does have advantages, flat power backoff provides nearly as good performance in most cases while being simpler.

Appendix V shows rate regions of vectored and non-vectored mixtures similar to those in Appendix I, but derived using the simpler IWF technique instead of OSB. This simpler algorithm provides performance similar to that of OSB; and allows the analyses to be extended to cover both ideal cancellation as well as more realistic conditions, and with histogram plots.

Appendix VI shows some specific rate-reach plots for vectored and non-vectored mixtures, giving examples of attractive tradeoffs that could be implemented in practice. Some high-level overall conclusions are made about DSM for managing vectored and non-vectored mixtures.

Much of the analysis in the appendices plots achievable tradeoffs between vectored groups and non-vectored lines. A typical example is highlighted here in Figure 7, which was calculated with simulation details as given in Section IV.2 using the IWF DSM technique. The case of a vectored group mixed with non-vectored lines shows that non-vectored lines may be run at 20-40 Mbps with little impact on the vectored lines' speed. The case of two vectored groups shows that the sum bit rate of the two groups is substantially higher than the maximum bit rate of either group, and so the tradeoff is better than a zero-sum game.

Downstream, 0.5 km, 1% Worst-Case
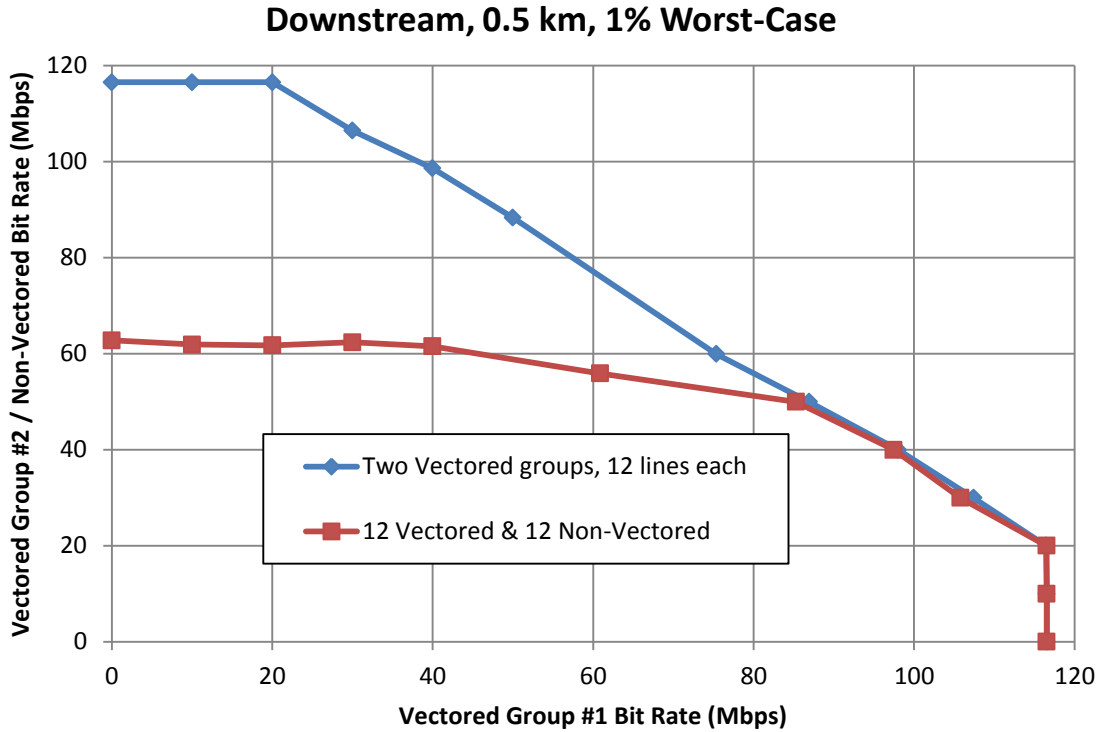
Figure 7 - Downstream 1% worst case bit rates for DSM performed by IWF with the two cases of mixed vectored and non-vectored VDSL2 lines, and two vectored groups, 0.5 km loop, collocated end-points.

While Figure 8 shows the various tradeoffs that may be chosen at a particular loop length, Figure 8 then chooses a particular tradeoff at each loop length and shows a plot of bit rate as a function of loop length using the simulations as described in Appendix VI. This trade-off has minimal impact on the vectored lines.
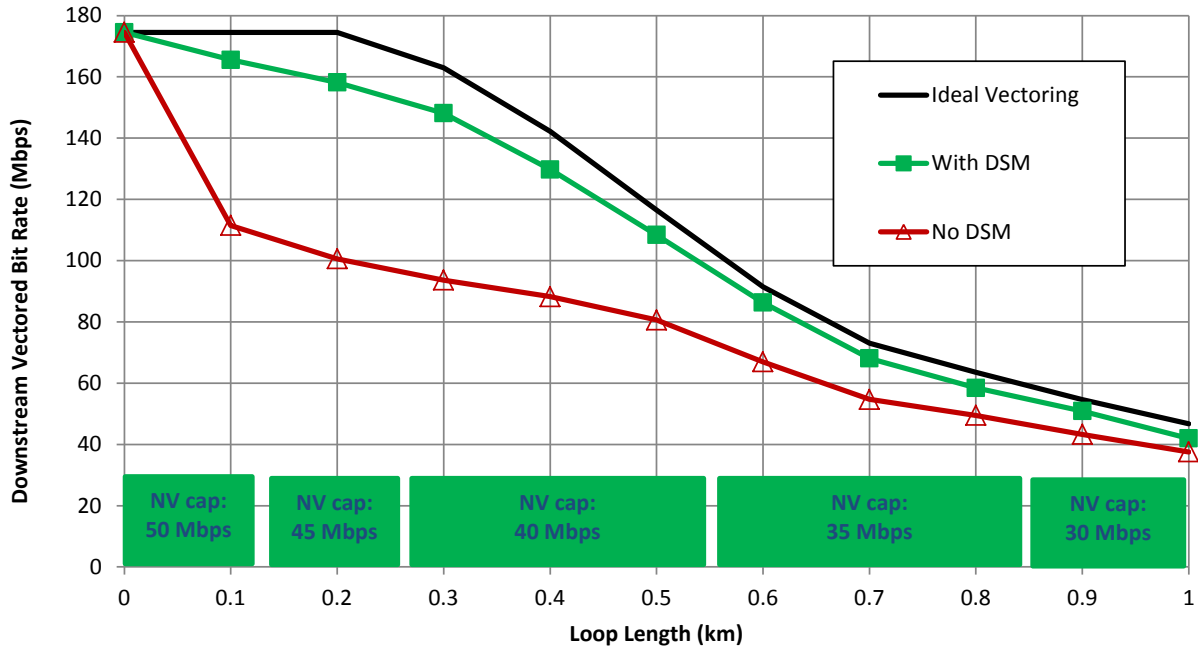
Figure 8 - Average downstream vectored bit rates, with non-vectored (NV) lines rate limited by a cap and DSM performed by iterative waterfilling (IWF).

Appendix I.   DSM Optimal Spectral Balancing (OSB) Simulation Results - Mix of Vectored and Non-vectored Lines

This appendix explores mixed deployments of vectored and non-vectored VDSL2. Simulations are run using the DSM level-2 spectral optimization technique known as Optimal Spectrum Balancing (OSB) [10]. These simulations assume that all subscribers are collocated. The achievable rate region is found, showing the trade-offs between vectored lines and non-vectored lines that can be achieved by joint spectral optimization.

In general, the rate region trade-off for the vectored group, and the group of non-vectored lines, is "squarish," showing a good trade-off between the two groups. It is shown to be generally possible to run non-vectored lines at appreciable bit rates in the same binder as vectored lines, with almost no degradation to the vectored lines, if the non-vectored lines are properly managed.

I.1   VDSL2 Bit Rate Calculations

The simulations here calculate VDSL2 downstream bit rates with both vectored VDSL2 and non-Vectored VDSL2.

The simulations use typical models and parameters, including the BT loop model, that are generally found in the NICC DSM technical report [7], and the ANSI spectrum management standard, T1.417 [8] Downstream VDSL2 Profile 17a is simulated. The transmit PSD is at most 3.5 dB below the VDSL2 profile 998ADE17-M2x-A PSD limit mask defined in Annex B of G.993.2 [1]. Additionally, transmit PSDs are limited to at most -55 dBm/Hz to meet the G.993.2 average power constraint.

The margin is 6 dB and the total coding gain is 3 dB. Bit rates are calculated by summing the capacity calculation of each 4.3125 kHz tone with a 9.75 dB SNR gap, with bits per Hz per sub-carrier lower limited to at least one bit and upper limited to 14 bits per Hz per sub-carrier. As defined in T1.417, for VDSL2 6% guard-band is imposed at the edge of each passband below 12 MHz and these guard-bands carry no bits. Above 12 MHz, a guard band of 175 kHz is assumed at the edge of each passband. Loops are all 0.4 mm / 26 AWG. AWGN at -140 dBm/Hz is added to each receiver, unless noted otherwise.

Simulations with Optimal Spectrum Balancing (OSB) [10] have 1% worst-case same-binder FEXT plus -140 dBm/Hz noise, with two Cabinet-based VDSL2 and two Exchange-based VDSL2 crosstalkers. The transmit PSDs of all VDSL2 lines are shaped by OSB. It is ideally assumed that self-FEXT cancellation is perfect with vectoring, so there is no residual crosstalk between the vectored lines, but there are two non-vectored lines worst-case crosstalkers into the vectored lines; and there are also two vectored worst-case crosstalkers and one non-vectored worst-case crosstalker into the non-vectored lines. The worst-case FEXT from the non-vectored and the vectored lines into the non-vectored lines is FSAN summed.

Note that average crosstalk from 24 disturbers is roughly equal to 99% worst-case crosstalk from two disturbers [See Appendix III] so the results here have slightly worse crosstalk than the typical case with a filled 25-pair binder.

March 2014   48 of 101

For most simulations, it is assumed that the non-vectored lines and the vectored lines all originate at the same cabinet. Some additional cases were run with non-vectored lines deployed from an exchange, and vectored lines deployed from a cabinet. The E-side length is the distance from the exchange to the cabinet. The D-side length is the distance from the cabinet to the TU-R (km).

I.2   Rate Regions

This section presents a general discussion of "rate regions," since these are the fundamental results of this appendix. Two extremes are presented for illustration.
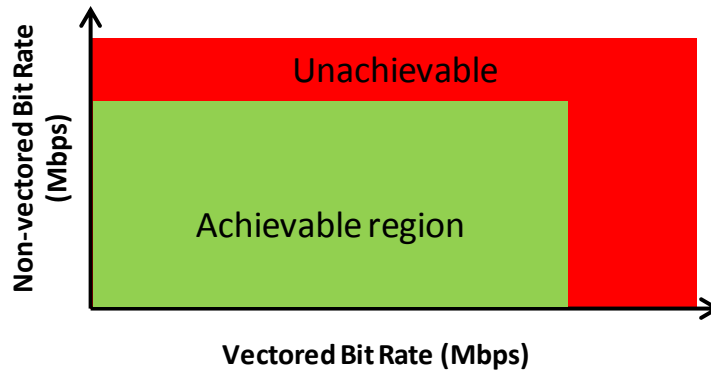
Figure I-1 - Fanciful illustration of a "square" rate-region. The unachievable region extends to infinity.

Figure I-1 shows a fictional illustration of a "square" rate region. Here the two systems (in this case a vectored group and a set of non-vectored lines) have no interaction. Either set of lines can transmit whatever rate they want without disturbing each other.
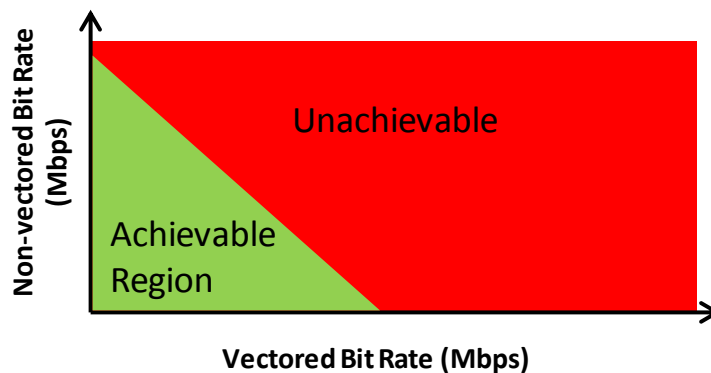
Figure I-2 - Fanciful illustration of a "zero-sum" rate-region. The unachievable region extends to infinity.

Figure I-2 is a fictional illustration of a "zero-sum" rate region for two systems: a vectored group, and a set of non-vectored lines. As the bit rate of one system increases, the bit rate of the other system decreases proportionally. Here one system can only increase speed by directly cutting into the speed of the other system.

Actual achievable rate-regions are shown here for mixed deployments of a vectored group and a set of non-vectored lines, and these are seen to generally offer favorable trade-offs somewhat similar to the "square" rate region.

I.3   Downstream OSB Results

Optimal Spectrum Balancing (OSB) [10] is a DSM algorithm that finds the best transmit PSDs for given bit rate trade-offs. While OSB is computationally intensive and generally suited to centralized DSM implementations, there are a number of alternative algorithms that can run in a non-centralized, distributed, manner that either meet the OSB bit rates or come very close [6].
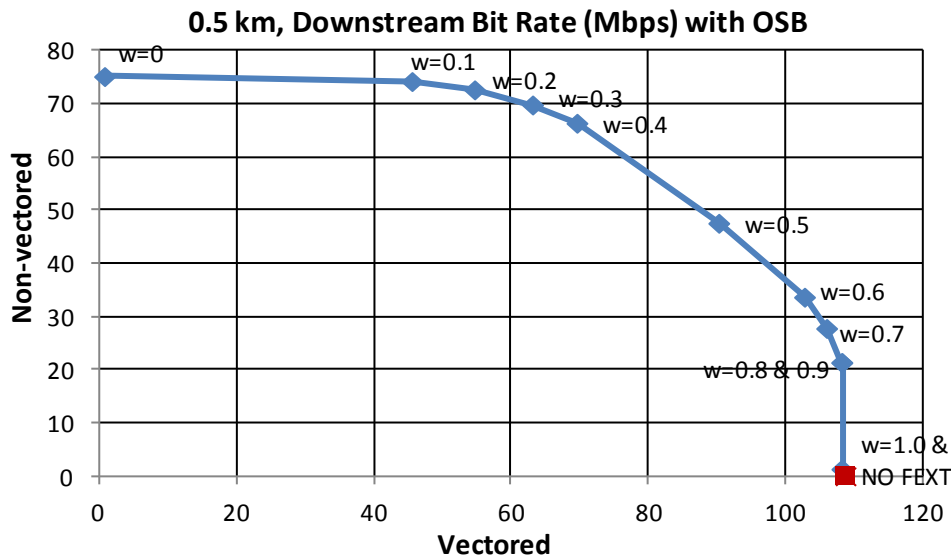
Figure I-3 - Vectored and non-vectored downstream bit rates with OSB, 0.5 km loop, collocated end-points.

Figure I-3 shows the downstream bit rates that are achievable with OSB on a 0.5 km loop. An input to the OSB algorithm is the trade-off parameter "w," $0 < w < 1$; which determines how much to weight each system. Here w is varied from 0 to 1 in increments of 0.1, tracing out the points as shown in Figure I-3.

Figure I-4 shows the downstream bit rates that are achievable with OSB on a 0.3 km loop, and Figure I-5 shows the downstream bit rates that are achievable with OSB on a 1.0 km loop.

Downstream Bit Rate (Mbps) with OSB, 0.3 km
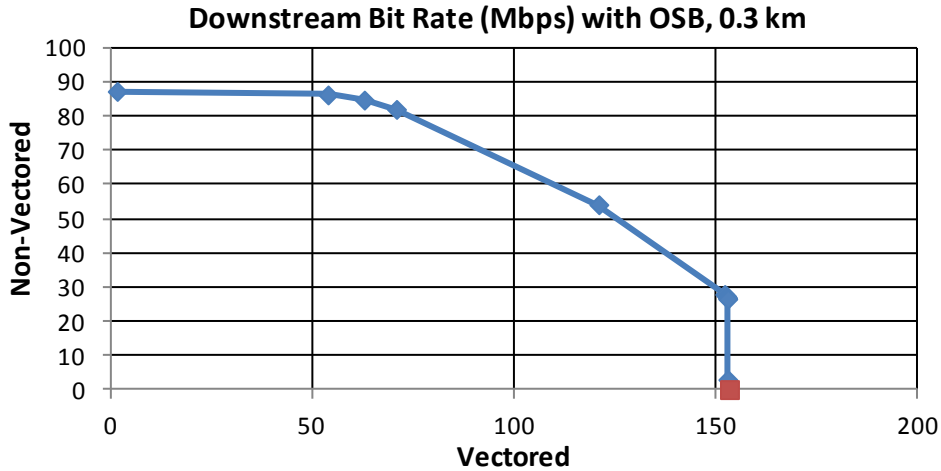
Figure I-4 - Vectored and non-vectored downstream bit rates with OSB, 0.3 km loop, collocated end-points.

Downstream Bit Rate (Mbps) with OSB

1 km

Figure I-5 - Vectored and non-vectored downstream bit rates with OSB, 1.0 km loop, collocated end-points.

A case of mixed cabinet and exchange-based lines is shown. Here the vectored lines are deployed from the cabinet, the non-vectored lines are deployed from the exchange, and both sets of lines transmit in the same distribution cable for the full length from the cabinet. Figure I-6 shows the downstream bit rates that are achievable with OSB with 0.5 km E-side length from exchange to cabinet, and 0.5 km D-side length from cabinet to customers.
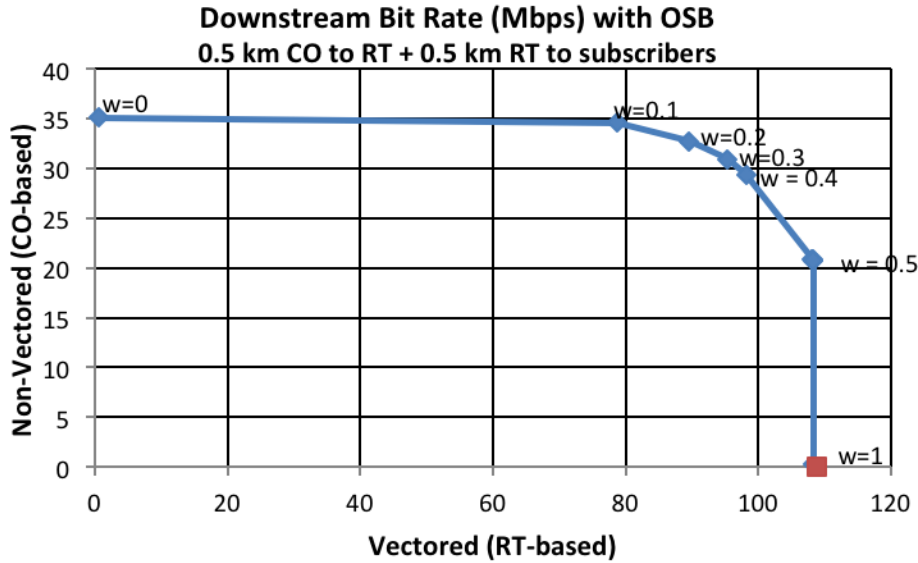
Figure I-6 - Vectored and non-vectored downstream bit rates with OSB, 0.5 km E-side length from exchange to cabinet, 0.5 km D-side length from cabinet to customers, collocated end-points.

I.4   Upstream OSB Results

Figure I-7 shows the upstream bit rates that are achievable with OSB on a 0.5 km loop.
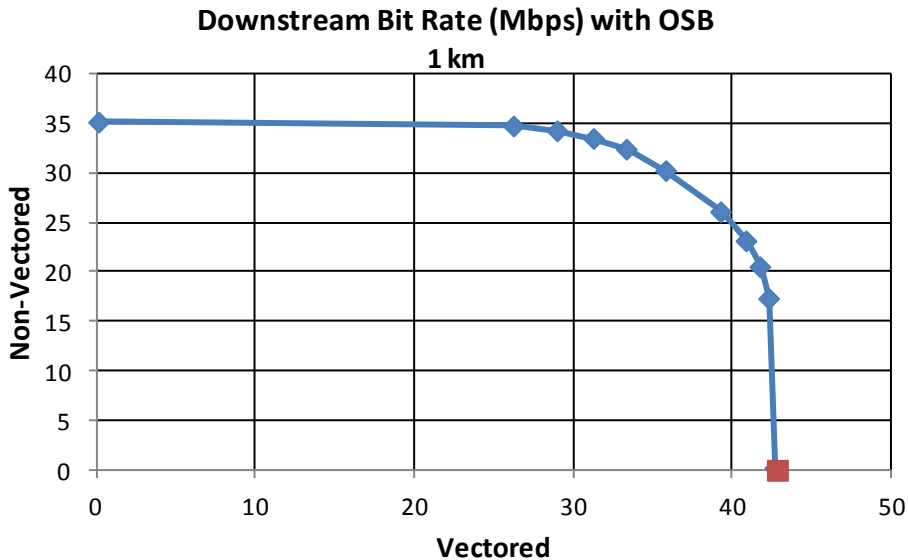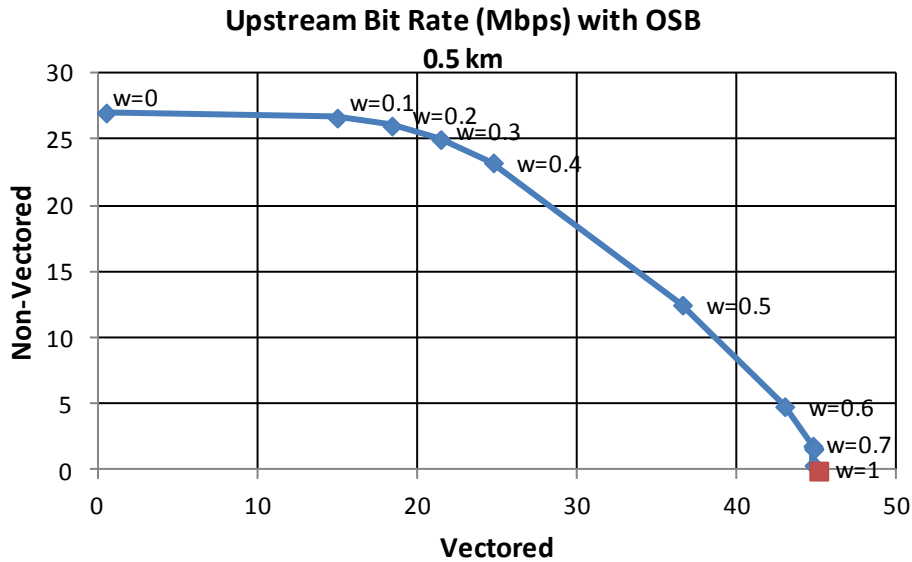
Figure I-7 - Vectored and non-vectored upstream bit rates with OSB, 0.5 km loop, collocated end-points.

Figure I-8 shows the upstream bit rates that are achievable with OSB on a 1.0 km loop.
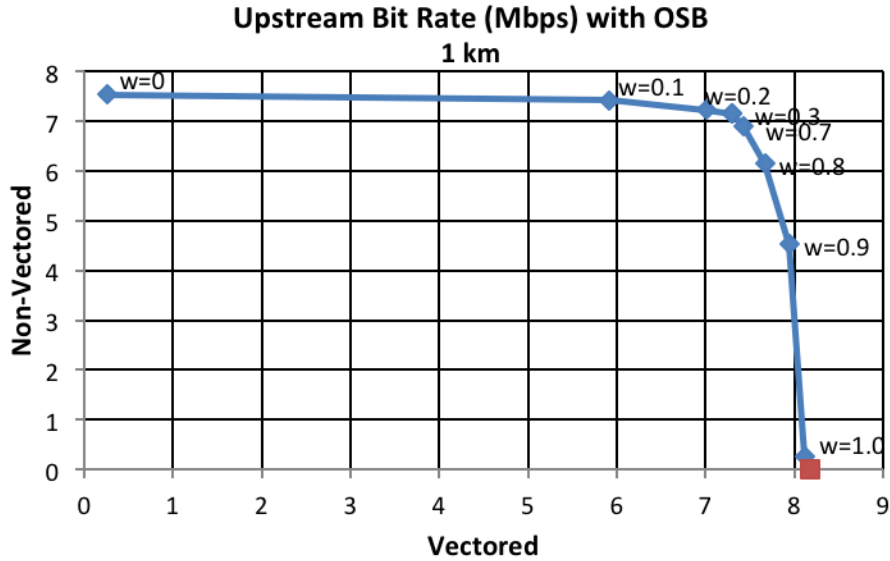
Figure I-8 - Vectored and non-vectored upstream bit rates with OSB, 1.0 km loop, collocated end-points.

A case of mixed cabinet and exchange-based lines is shown. Here the vectored lines are deployed from the cabinet, the non-vectored lines are deployed from the exchange, and both sets of lines transmit in the same distribution cable for the full length from the cabinet. Figure I-9 shows the upstream bit rates that are achievable with OSB with 0.5 km E-side length from exchange to cabinet, and 0.5 km D-side length from cabinet to customers.
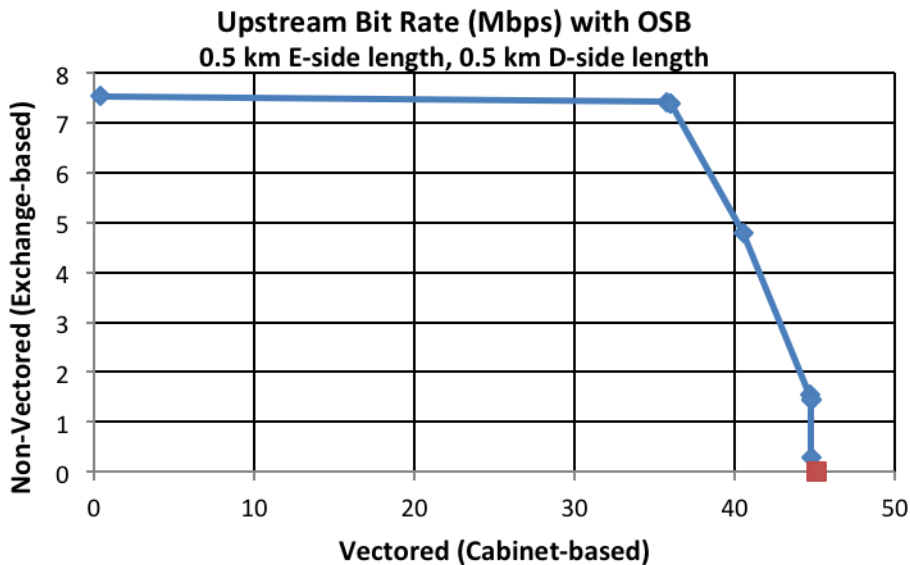
Figure I-9 - Vectored and non-vectored upstream bit rates with OSB, 0.5 km E-side length from exchange to cabinet, 0.5 km D-side length from cabinet to customers.

I.5   Comparative Plots - Downstream

In this Section multiple curves are plotted on the same figure to show trends of downstream data rates as a function of distance, number of vectored lines, and varying background noise levels (e.g., the level of flat AWGN+background noise).
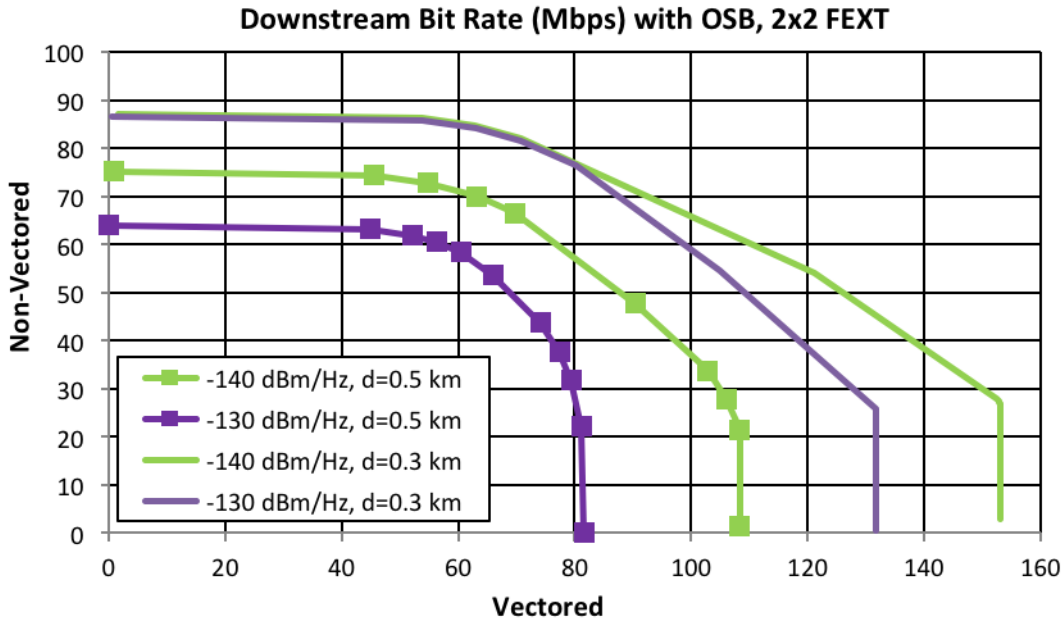
Figure I-10 - Vectored and non-vectored downstream bit rates with OSB, collocated end-points, with varying background noise levels and loop lengths. Network end points are collocated and the loop length is d.

Figure I-10 confirms that the dominant impairment for vectored lines is additive background noise, whereas the dominant impairment for non-vectored lines is crosstalk at shorter distances.

A variation of 10 dB in the level of external background noise can produce wide variations of the maximum data rate achievable by vectored lines, as shown in Figure I-11. For example, at d=300m and with no non-vectored FEXT the rates of vectored lines go from 155 Mbps down to 135 Mbps, and at d=500m the rates of vectored lines go from 110 Mbps down to 80 Mbps, as the background noise increases by 10 dB.

Besides background noise, other non-idealities such as imperfect crosstalk cancellation (not all crosstalk is ideally cancelled by the vectoring engine) and partial cancellation (not all disturbers are cancelled) can further reduce the maximum rate achievable by vectoring.

Figure I-11 - Vectored and non-vectored downstream bit rates with OSB, collocated end-points, with varying loop length.
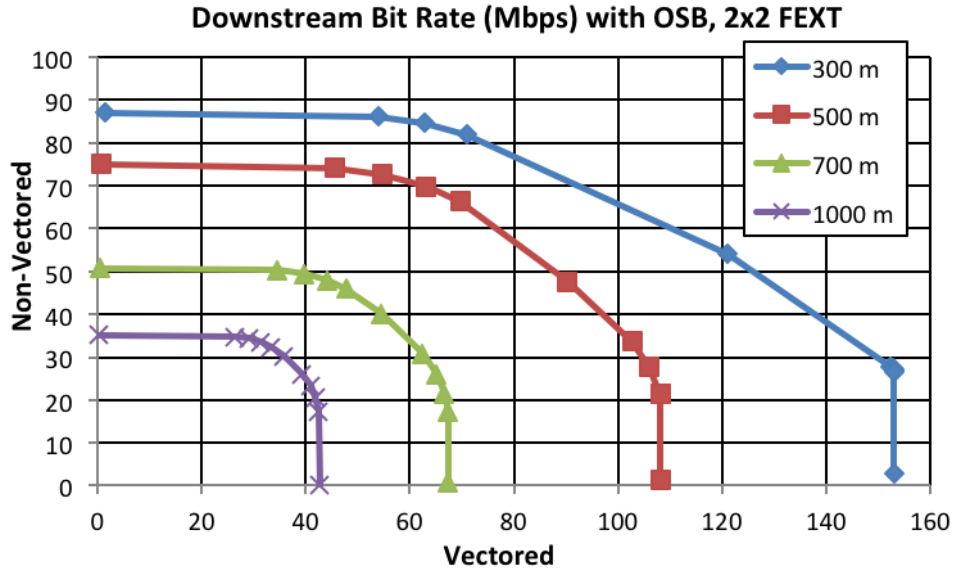
Figure I-11 shows that rather regular trends characterize the trade-offs between non-vectored and vectored lines as loop length varies.

At d=300m, vectored lines can achieve 100 Mbps (an EU mandate) in worst case crosstalk conditions and ideal vectoring conditions (total self-FEXT cancellation) while non-vectored lines run at most at 65 Mbps which is a very high speed compared to what is deployed today with ADSL2plus or even current offerings of VDSL2.

At d=500m, vectored lines can achieve 100 Mbps in worst case crosstalk conditions and ideal vectoring conditions (total self-FEXT cancellation) while non-vectored lines can run at most at 35 Mbps which is a very high speed compared to what is deployed today with ADSL2+, and which is similar to current offerings of VDSL2.

At longer and longer distances, the data rates of non-vectored and vectored lines decrease rapidly but the achievable rate region becomes more "squarish" confirming that the two systems can operate almost independently.

Downstream Bit Rate (Mbps) with OSB, 2x2 FEXT, 500 m DSEL
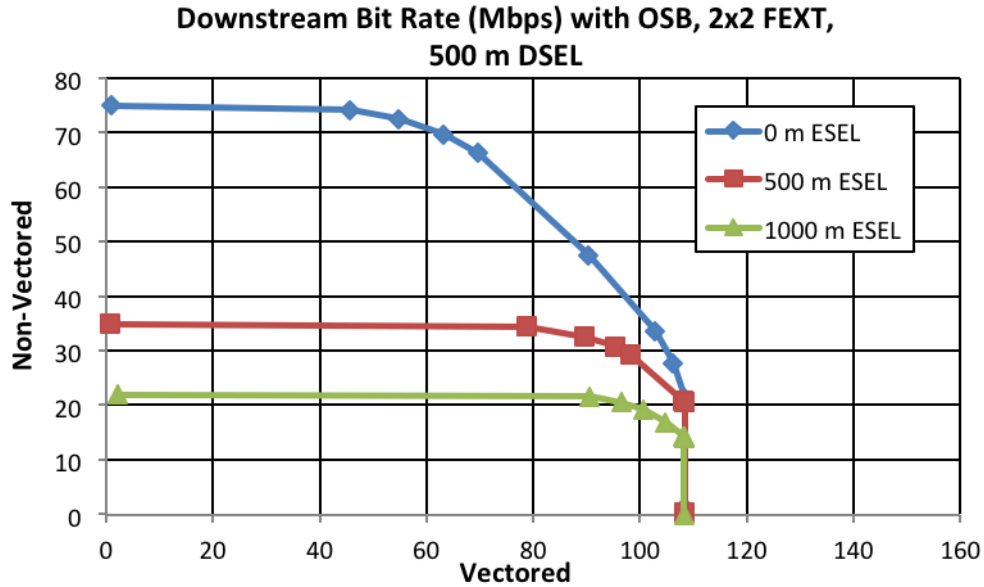
Figure I-12 - Vectored and non-vectored downstream bit rates with OSB, collocated end-points, with the cabinet located DSEL=500m from the customers, and varying the distance from the CO to the cabinet (ESEL).

Figure I-12 shows a case with non-vectored VDSL2 deployed from a CO and vectored VDSL2 deployed from a cabinet, with the systems transmitting in the same distribution binder. Here the D-Side Electrical Length (DSEL, distance from the cabinet to the customers) is 500m, and the E-Side Electrical Length (ESEL, distance from the CO to the cabinet) varies from 0 (CO-deployment) to 1000 m. A regular trend is seen, where the non-vectored lines simply have lower speeds as the distance from the CO increases. In the case where lines originate at different locations, the CO-based non-vectored lines suffer most degradation. Furthermore, as the CO-to-cabinet distance increases, the trade-off between vectored and non-vectored lines improves (the rate region is more squarish).
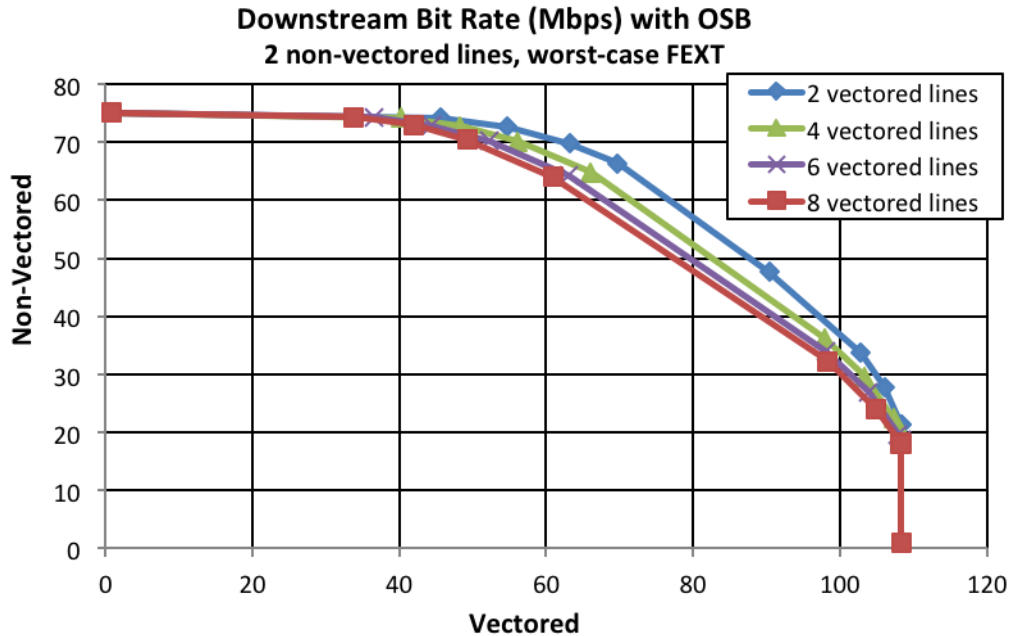
Figure I-13 - Vectored and non-vectored downstream bit rates with OSB, collocated end-points, with 500m loop length, 2 non vectored lines and varying number of vectored lines. 1% worst-case FEXT.

Figure I-13 shows the change in data rate trade-offs when more and more vectored lines are added. The effect is small and it is important to point out that the effect shown would be even smaller in practice as here we have worst case crosstalk which is clearly pessimistic. This plot confirms that as a deployment of vectored lines gradually increases, there is only a small change in compatibility tradeoffs.

I.6   Comparative Plots - Upstream

In this Section we plot multiple curves on the same Figure to show trends of upstream data rates as a function of distance, number of vectored lines, and varying background noise.

Figure I-14 - Vectored and non-vectored upstream bit rates with OSB, collocated end-points, with varying AWGN levels and loop lengths.

Figure I-14 confirms that the dominant impairment for vectored lines is noise, whereas the dominant impairment for non-vectored lines is crosstalk at shorter distances - just as for the downstream case.

A variation of 10 dB in the level of background noise can produce wide variations of the maximum data rate achievable by vectored lines. For example, with no non-vectored crosstalk, at d=300m the upstream rates of vectored lines drop from 62 Mbps to 52 Mbps, and at d=500m the rates of vectored lines can drop from 45 Mbps to 31 Mbps as the noise increases by 10 dB. However, the trade-offs between vectored and non-vectored lines exhibit similar trends just like the downstream case.
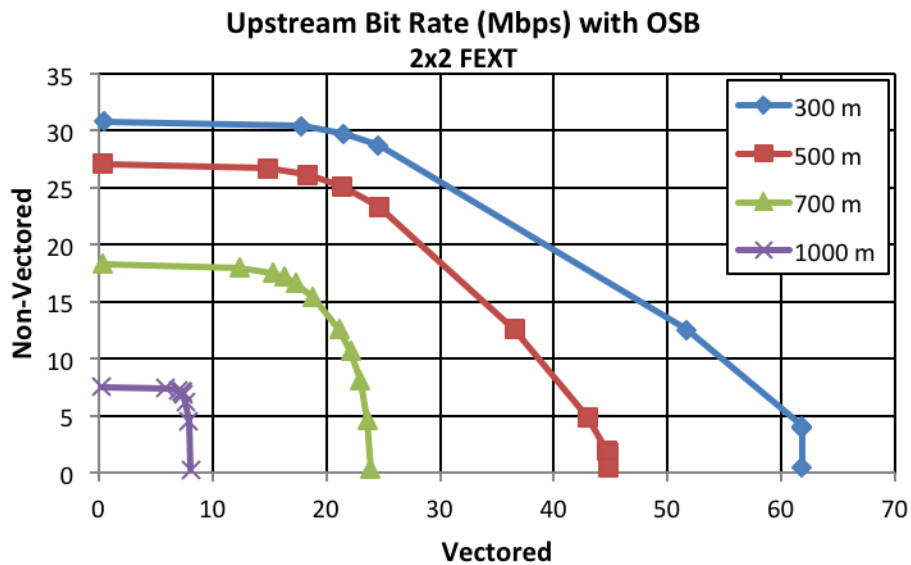
Figure I-15 - Vectored and non-vectored upstream bit rates with OSB, collocated end-points, with varying loop length.

Figure I-15 shows that rather regular trends characterize the trade-offs between non-vectored and vectored lines as loop length varies, similar to downstream.

At d=300m, vectored lines can achieve 50 Mbps with worst case crosstalk and ideal vectoring, while non-vectored lines run at 13 Mbps. Other achievable data rate combinations are 40 Mbps for vectored lines and 20 Mbps for non-vectored lines. These are very high speeds for upstream, particularly compared to the mere 1 Mbps or available today with ADSL2. Also, vectored lines run at 2x or 3x the speed of non-vectored lines, which shows that the investment in vectoring offers a payback.

At longer and longer distances, the data rates of non-vectored and vectored lines decrease rapidly but the achievable rate region becomes more "squarish" confirming that the two systems can operate almost independently just like the downstream case
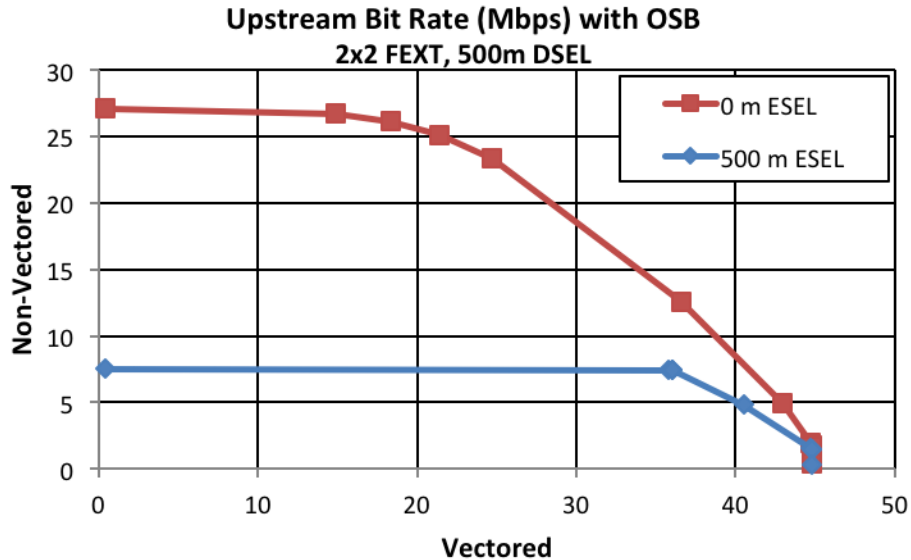
Figure I-16 - Vectored and non-vectored upstream bit rates with OSB, collocated end-points, with the cabinet located DSEL=500m from the customers, and varying the distance from the CO to the cabinet (ESEL).

Figure I-16 shows a case with non-vectored VDSL2 deployed from a CO and vectored VDSL2 deployed from a cabinet, with the systems transmitting in the same distribution binder. Here the D-Side Electrical Length (DSEL, distance from the cabinet to the customers) is 500m, and the E-Side Electrical Length (ESEL, distance from the CO to the cabinet) is either 0 (CO-deployment) or 500 m. A regular trend is seen, where the non-vectored lines simply have lower speeds as the distance from the CO increases. In the case where lines originate at different locations, the CO-based non-vectored lines suffer most degradation. Furthermore, as the CO-to-cabinet distance increases, the trade-off between vectored and non-vectored lines improves (the rate region is more squarish). This trend was the same for downstream.
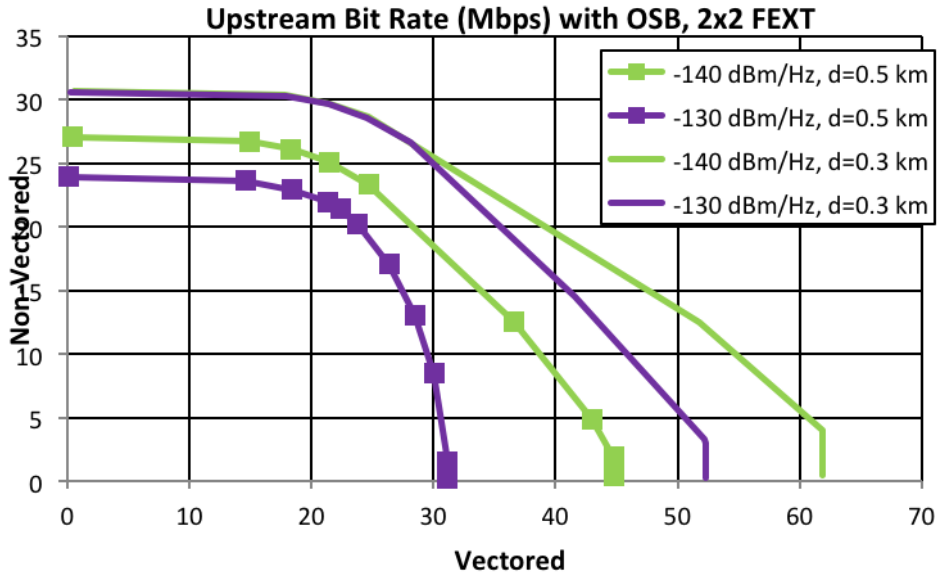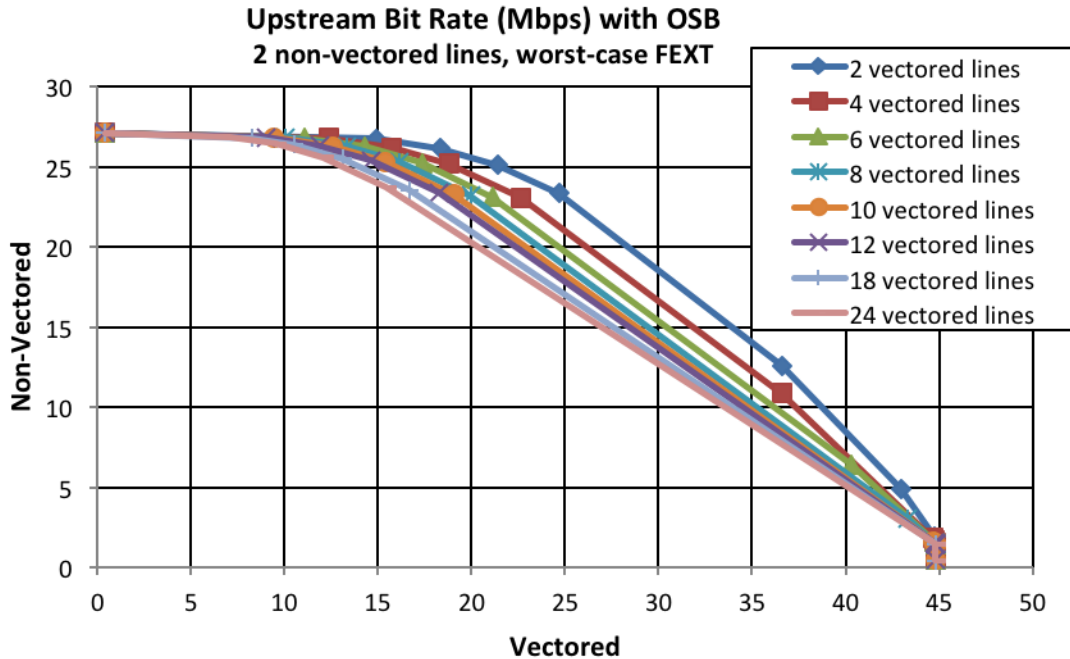
Figure I-17 - and non-vectored upstream bit rates with OSB, collocated end-points, with 500m loop length, 2 non vectored lines and varying number of vectored lines.

Figure I-17 shows the change in data rate trade-offs as more and more vectored lines are added. The effect is moderate, and it is important to point out that this would smaller in practice as here a worst-case crosstalk is considered which is clearly overly pessimistic. Similar to downstream, it is seen that the rate trade-offs are relatively insensitive to the number of vectored lines.

I.7   Results

This section shows the range of bit rates due to typical crosstalk coupling variations. Crosstalk is modeled here with the ATIS MIMO crosstalk model [11], with a filled 25-pair binder. There is no DPBO. The iterative waterfilling (IWF) DSM spectral optimization technique is used [6]; this technique can run in autonomously across multiple distributed lines with different crosstalk characteristics. Otherwise, simulation parameters are the same as elsewhere in this appendix.

A set of same-binder crosstalk couplings is drawn from the ATIS MIMO model, and the bit rates of all the lines in the binder are computed. This is repeated 1000 times and histograms of all bit rates are plotted. The drawn couplings are increased by 1 dB to align the ATIS MIMO model 1% worst case with the ATIS 1% worst-case crosstalk model. The case of all non-vectored lines has 24 crosstalkers. The cases with vectored and non-vectored lines has crosstalk from 12 non-vectored lines and from 12 vectored lines; with compatibility between them optimized by IWF.

500 m, All non-vectored lines, 24 Crosstalkers
Triangle is 1% worst-case point
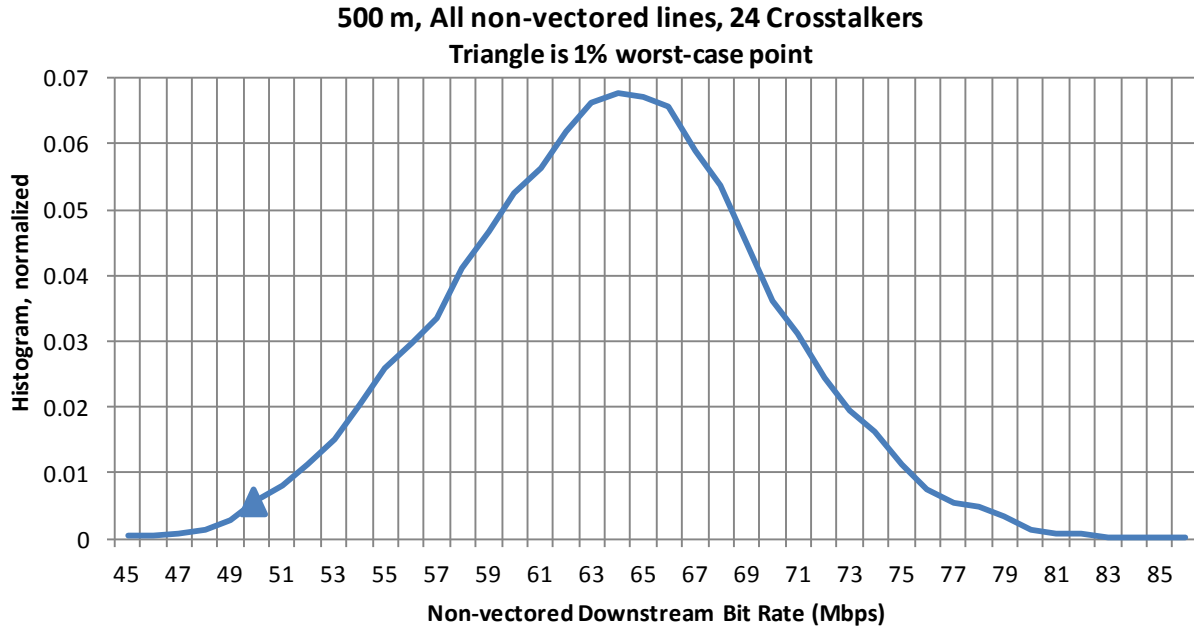
Figure I-18 - Filled binder of non-vectored lines, showing the variation due to different crosstalk couplings.

500 m, 12 vectored and 12 non-vectored lines
Triangles are 1% worst-case points
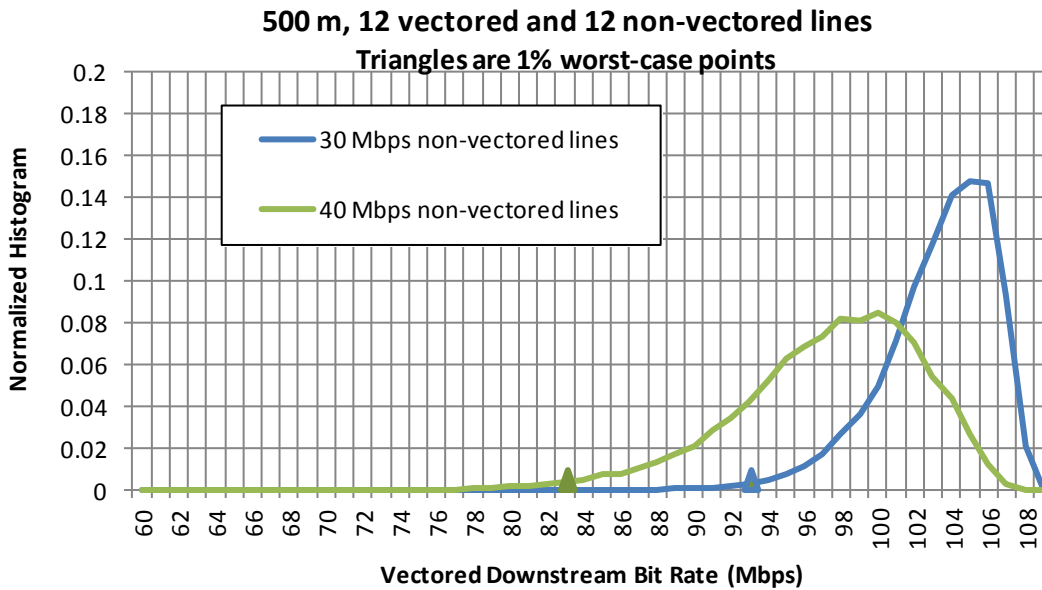
Figure I-19 - Filled binder of equal numbers of non-vectored lines and vectored lines, showing the variation due to different crosstalk couplings.

Histograms of downstream VDSL2 bit rates are shown in Figure I-18 and Figure I-19. It is seen that vectoring substantially improves performance, even with the presence of many non-vectored lines in the same binder. It is also shown that the 1% worst-case bit rate is far below typically achieved bit rates.
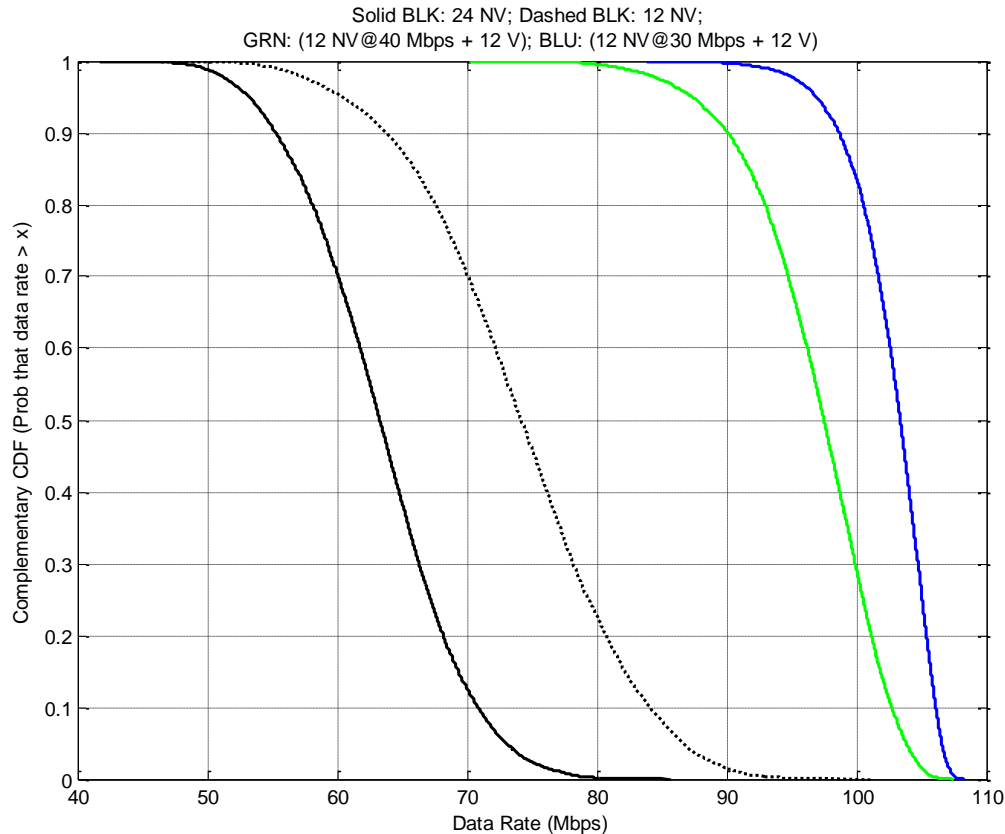
Figure I-20 - Complementary Cumulative Distribution Function (III-CDF) of vectored or non-vectored data rates for the following cases: 24 non-vectored lines (Solid BLK); 12 non-vectored lines (Dashed BLK); 12 vectored lines plus 12 non-vectored lines capped at 40 Mbps (GRN); 12 vectored lines plus 12 non-vectored lines capped at 30 Mbps (BLU).

In addition to histograms, it is also interesting to look at the Complementary Cumulative Distribution Functions (III-CDF) which gives the probability that a certain data rate R or more is achieved. Figure I-20 shows the III-CDF for the scenarios reported at the beginning of this section plus the case when there are only 12 non-vectored disturbers (dashed black curve).
Figure I-20 shows that 99% of lines can achieve a rate of 55 (50) Mbps or more when the binder is filled with 12 (24) non-vectored lines but only 1% of lines would achieve between 91-101 (79-86) Mbps.

If 12 vectored lines were added to the same binder and the 12 non-vectored lines were capped at 40 Mbps, then 99% of lines would achieve a rate of 82 Mbps or more while 30% of lines would achieve more than 100 Mbps. In the case that the 12 non-vectored lines were capped at the lower rate if 30 Mbps, then 99% of lines would achieve a rate of 93 Mbps or more while 84% of lines would achieve more than 100 Mbps.

This confirms that the spreading of the data rates of the mixed binder can be considered narrower than the case when the binder is 50% or 100% full of non-vectored lines only. This spread is also inversely proportional to the cap imposed on the non-vectored lines.

March 2014   62 of 101

I.8   Additional considerations on OSB

OSB is a centralized DSM algorithm. Centralized systems require a central hub with full knowledge of the network, and this allows for better performance at a cost of increasing the complexity and computational time. OSB maximizes the weighted sum rate across all the users and requires solving for the optimal transmit powers for each user separately on each frequency tone by exhaustive search. OSB is not computationally tractable for many users or many groups of users, and instead the main use of OSB is for determining the upper bounds on the performance of other DSM algorithms. OSB also produces transmit PSDs that achieve these upper bounds, and these PSDs are sometimes insightful. Other centralized DSM algorithms with lower complexity than OSB have been presented in the literature, e.g. Grouping Spectrum Management (GSM), Monotonic Optimization Based Power Control (MARL), Iterative Spectrum Balancing (ISB), Generalized ISB (GISB), or Centralized Multi-Level Waterfilling (MLWF). These are often better suited to field implementation than OSB.

There are also distributed DSM algorithms which can be run independently by multiple Spectrum Management Centers (SMCs) and are often more useful in the field as they are computationally much simpler, do not require any inter-SMC communication, and provide results that are often very close to the ones of OSB. Some well-known-distributed DSM algorithms are: Autonomous Spectrum Balancing (ASB), Distributed Spectrum Balancing (DSB), Iterative Waterfilling (IWF), Selective Iterative Waterfilling (SIW), Semi-Blind Spectrum Balancing (2SB), Successive Convex Approximation for Low-complexity (SCALE).

The results found by OSB are dependent on the number of users (see Figure I-21), and the chosen value of the weight w (see Figures I-3, and I-21).
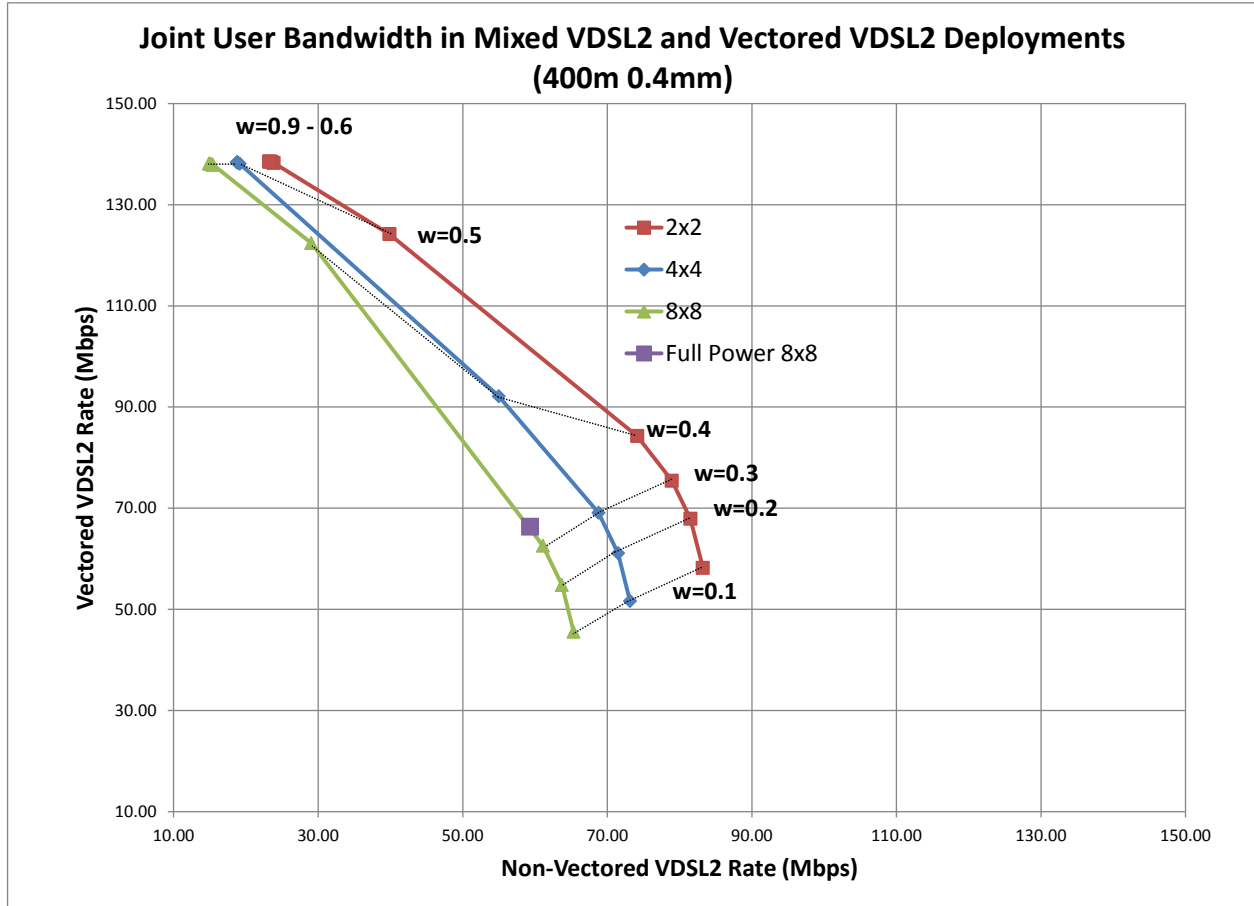
Figure I- 21 - Vectored and non-vectored downstream bit rates with OSB, 0.4 km loop.

The relationship between the users' data rate and the weight w is not straightforward and intuitive choices for the value of the OSB weighting parameter w do not always lead to expected results, e.g. w=0.5 does not always imply that the two users/groups of users achieve similar data rates or half their maximum data rate. Instead, many values of w should be evaluated to determine the achievable rate region, then this rate region can be used to select operational bit rates.

OSB is not meant to be used by selecting a priori a specific value of w upon which regulatory policy can be written, but indeed the opposite: it is regulatory policy that would determine what the proper value of w is. Regulatory policy may change from country to country, and also definitions of fairness between operators can be dependent on the jurisdiction where regulations are enacted. And in the absence of regulation, any voluntary agreement between operators to support a centralized DSM scheme like OSB would have to first define a definition of fairness upon which DSM would find the best line profiles.

From this point of view, OSB (like other DSM techniques) provides a framework where various definitions of fairness or various regulatory policies can be supported. Furthermore, OSB provides the optimal (from a weighed sum rate metric) rate region that is actually achievable thus identifying the region of all the data rates achievable by the vectored and non-vectored lines. The identification of this region allows the selection of feasible and desirable operating points, however, there is no "magic number" that specifies a meaningful coexistence solution for a heterogeneous and dynamic

network. Since things in the field vary over time, a dynamic approach can (and should) always tune the network to the best operating point as defined by coexistence objectives.

I.9   Conclusions

This appendix showed achievable rate regions for some cases of deployments where a vectored group and a set of non-vectored lines share the same cable binder. These regions show good trade-offs, and in particular the non-vectored lines can transmit reasonable bit rates with very little impact on the vectored group, if the non-vectored lines are properly managed.

The presented statistical results also suggest that the assessment of the performance of vectored lines in the presence of non-vectored lines should be evaluated beyond the traditional 1% worst case data rate and that a better and more complete picture can be given when taking into consideration higher percentile points and the distribution of data rates.

Appendix III shows that rate region trade-offs with users at distributed distances are very similar to the trade-offs shown here for collocated users.

The overall conclusion is that it is technically feasible for mixed deployments of vectored and non-vectored lines to successfully coexist, and it was shown here that this can be implemented via Dynamic Spectrum Management (DSM).

March 2014   65 of 101

Appendix II.   DSM Optimal Spectral Balancing (OSB) Simulation Results – Multiple Vectored Groups

II.1   Introduction

This Appendix explores deployments of multiple vector groups. Simulations are run using the DSM level-2 spectral optimization technique known as Optimum Spectral Balancing (OSB). These simulations assume that all customers are collocated. The achievable rate region is found, showing the trade-offs between vectored groups.

II.2   VDSL2 Bit Rate Calculations

The simulations here calculate VDSL2 downstream bit rates with vectored VDSL2 only. The simulations use typical models and parameters that are generally found in the ANSI spectrum management standard, T1.417. Downstream VDSL2 Profile 17a is simulated. The transmit PSD is at most 3.5 dB below the VDSL2 profile 998ADE17-M2x-A PSD limit mask defined in Annex B of G.993.2. Additionally, transmit PSDs are limited to at most -55 dBm/Hz to meet the G.993.2 average power constraint.

The margin is 6 dB and the total coding gain is 3 dB. Bit rates are calculated by summing the capacity calculation of each 4.3125 kHz tone with a 9.75 dB SNR gap, with bits per Hz per sub-carrier lower limited to at least one bit and upper limited to 14 bits per Hz per sub-carrier. As defined in T1.417, for VDSL2 6% guard-band is imposed at the edge of each passband below 12 MHz and these guard-bands carry no bits. Above 12 MHz, a guard band of 175 kHz is assumed at the edge of each passband. Loops are all 0.4 mm / 26 AWG. AWGN at -140 dBm/Hz is added to each receiver, unless noted otherwise.

Simulations with Optimal Spectrum Balancing (OSB) have 1% worst-case same-binder FEXT plus -140 dBm/Hz noise, with two Cabinet-based vectored groups and each group with two lines. The transmit PSDs of all VDSL2 lines are shaped by OSB. It is ideally assumed that self-FEXT cancellation is perfect with vectoring, so there is no residual crosstalk between the vectored lines. The worst-case FEXT from the vectored lines is FSAN summed.

Note that average crosstalk from 24 disturbers is roughly equal to 99% worst-case crosstalk from two disturbers, so the results here have slightly worse crosstalk than the typical case with a filled 25-pair binder.

For most simulations, it is assumed that all vectored lines originate at the same cabinet. Some additional cases were run with a group of vectored lines deployed from an exchange, and another group of vectored lines deployed from a cabinet.

II.3   Downstream Rate Regions

Optimal Spectrum Balancing (OSB) is a DSM algorithm that finds the best transmit PSDs for given bit rate trade-offs. While OSB is computationally intensive and generally suited to centralized DSM implementations, there are a number of alternative algorithms that can run in a non-centralized, distributed, manner that either meet the OSB bit rates or come very close.
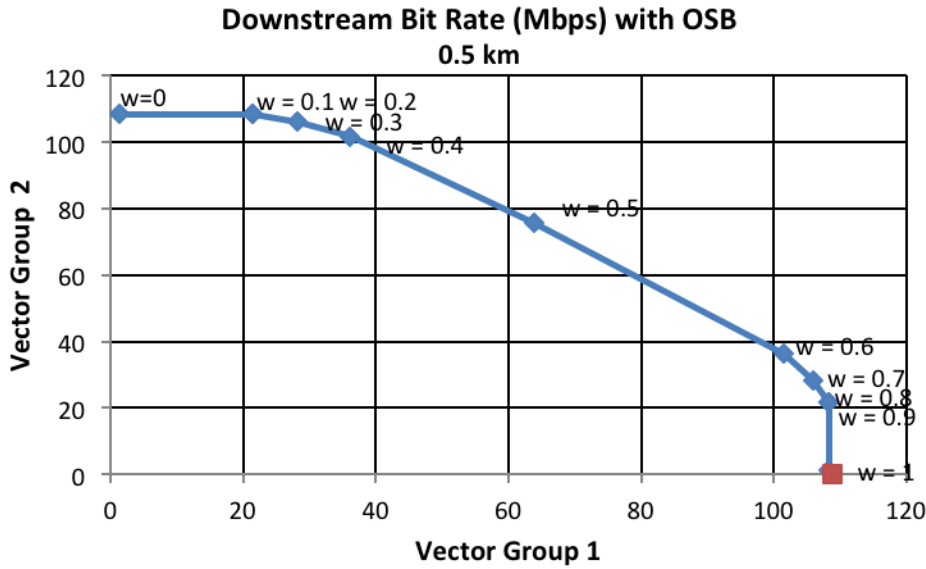
Figure II-1 - Downstream bit rates for two vectored groups each with two lines with OSB, 0.5 km loop, collocated end-points.

Figure II-1 shows the downstream bit rates that are achievable with OSB on a 0.5 km loop. An input to the OSB algorithm is the trade-off parameter "w," 0 < w < 1; which determines how much to weight each system. Here w is varied from 0 to 1 in increments of 0.1, tracing out the points as shown in Figure I- 3

Figure II-2 shows the downstream bit rates that are achievable with OSB on a 1.0 km loop.
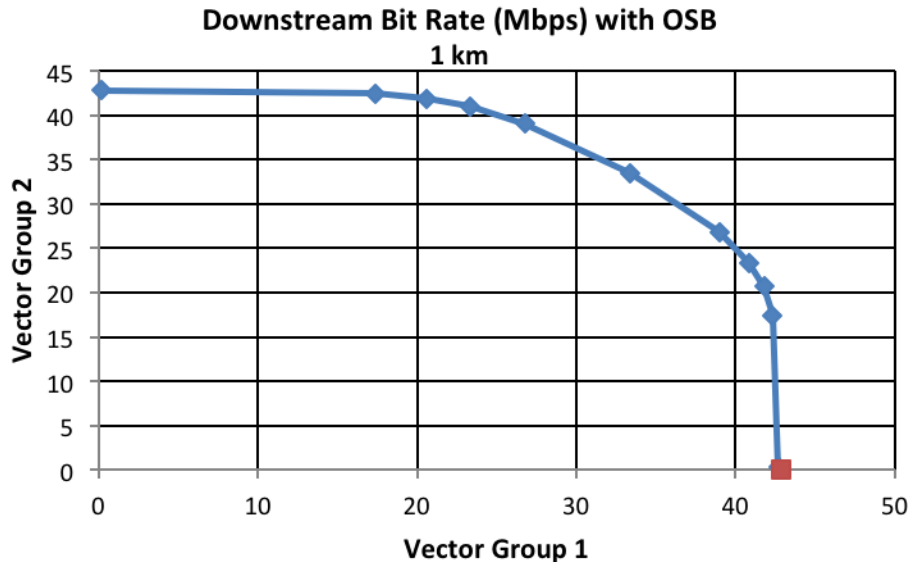
Figure II-2 - Same as Figure II-1, but for a 1 km loop.

A case of mixed cabinet and CO-based lines is shown in the next two Figures. Here a group of vectored lines is deployed from the cabinet, and another group of vectored lines is deployed from

the CO. Both groups of lines transmit in the same distribution cable for the full length from the cabinet.
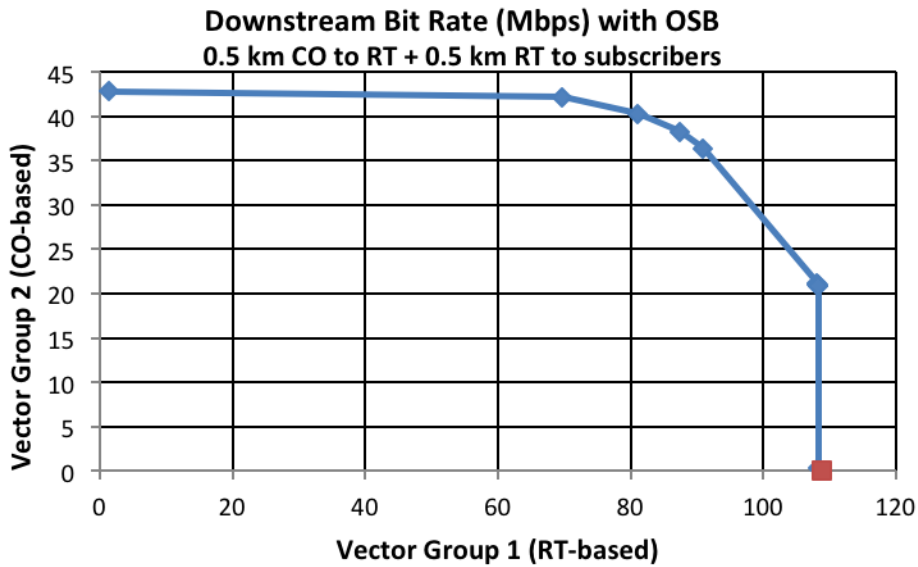
Figure II-3 - Downstream bit rates with OSB, 0.5 km length from CO to cabinet, and 0.5 km from cabinet to customers, collocated end-points.

With this vector-vector case, the rate region becomes more "squarish" as loop length grows, confirming that the two groups become increasingly independent as loop length grows. This behavior also occurs in the case of mixed deployments of vectored and non-vectored lines.

II.4   Upstream Rate Regions

This Section shows upstream rate region results.

Figure II-4 - Upstream bit rates for two vectored groups each with two lines with OSB, 0.5 km loop, collocated end-points.
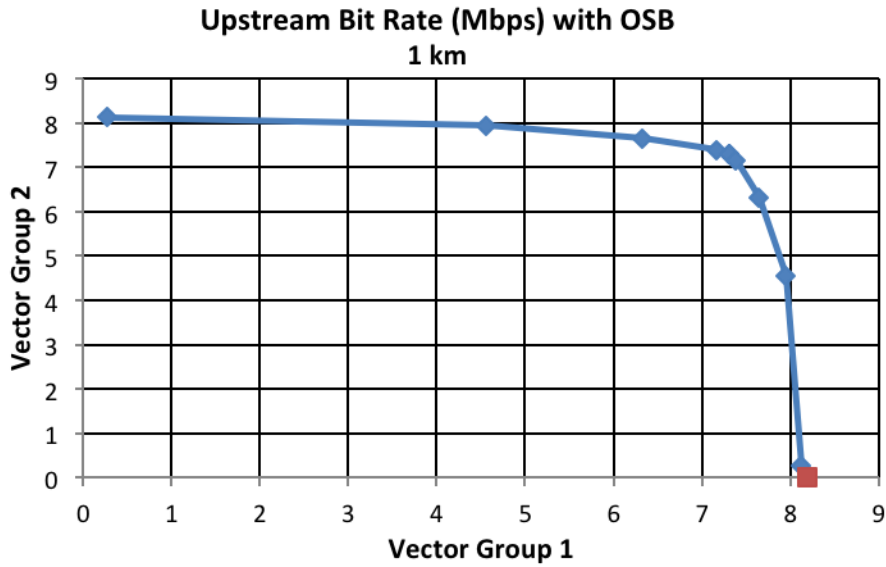
Figure II-5 - Same as Figure I-6, but with 1 km loop.

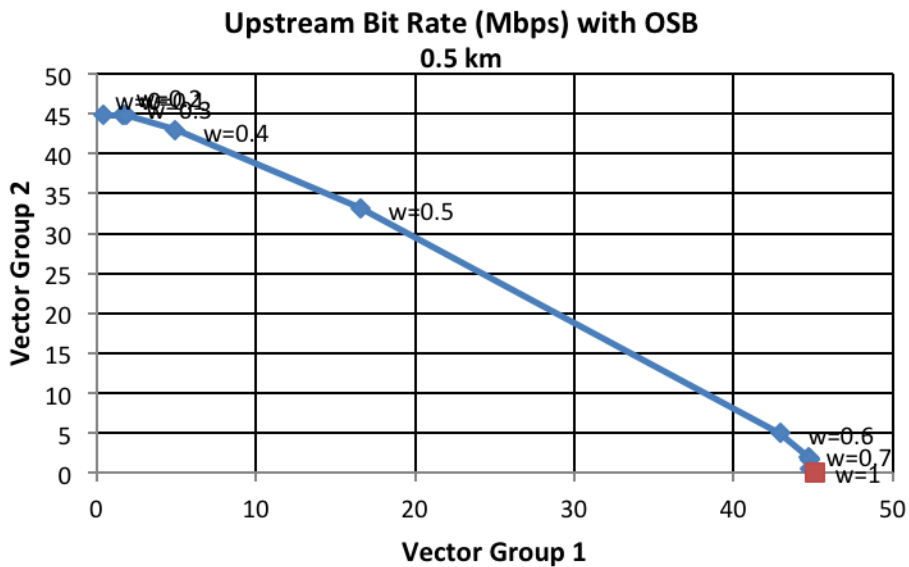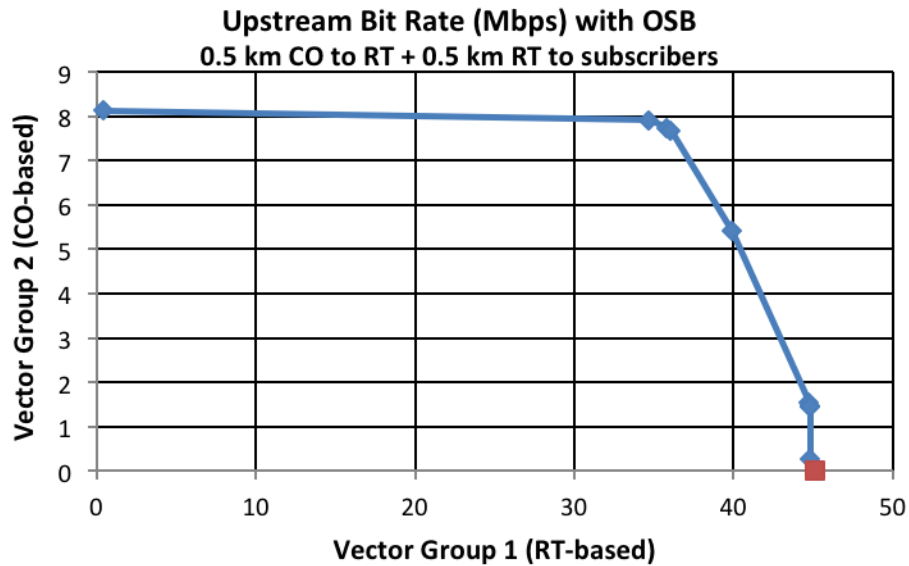**Figure II-6 - Upstream bit rates with OSB, 0.5 km length from CO to cabinet, and 0.5 km from cabinet to customers, collocated end-points.**

II.5   Comparison Between the Multiple Vector Group and the Mixed Deployment Cases

In the following Figures, the rate regions are compared for the case of mixed deployments of vectored and non-vectored lines, and for the case of two vector groups. The Figures confirm similar trends between the mixed deployment case and the multiple vector group case, with the difference that the rate region of the multiple vector case extends beyond the mixed deployment case. This is not surprising as the lines in the second vector group can achieve better rates than the lines in the non-vector group.
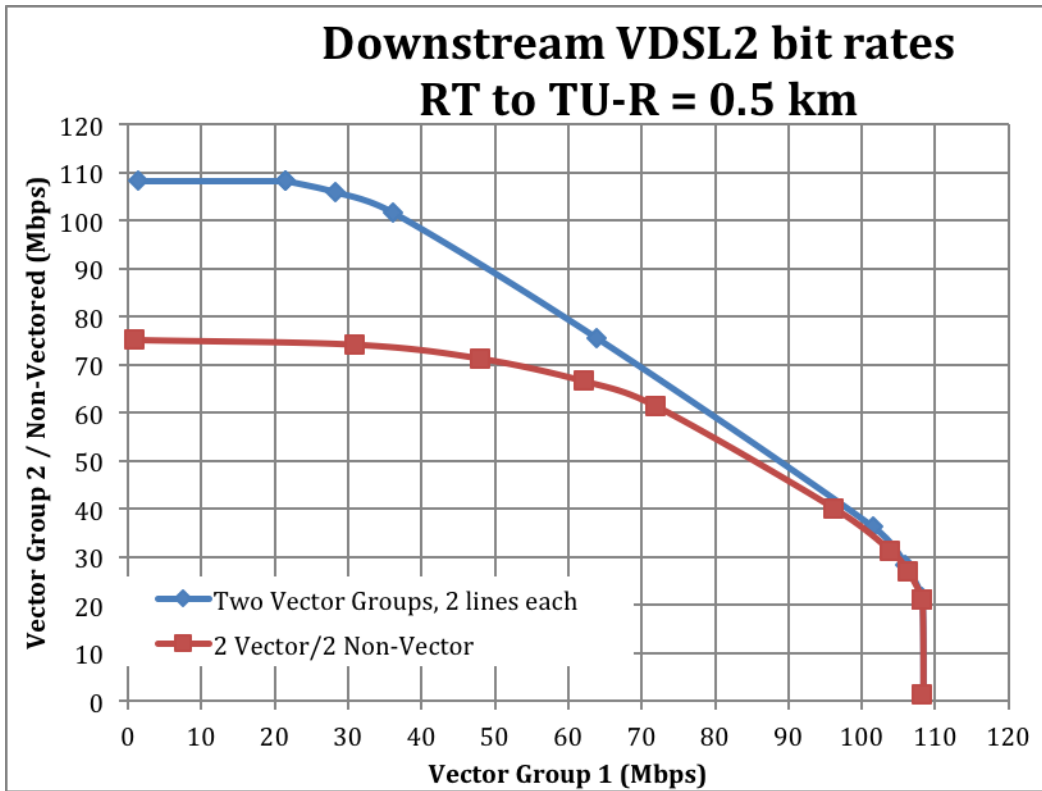
Figure II-7 - Downstream bit rates for the two cases of mixed vectored/non-vectored deployment and two vectored groups with OSB, 0.5 km loop, collocated end-points.
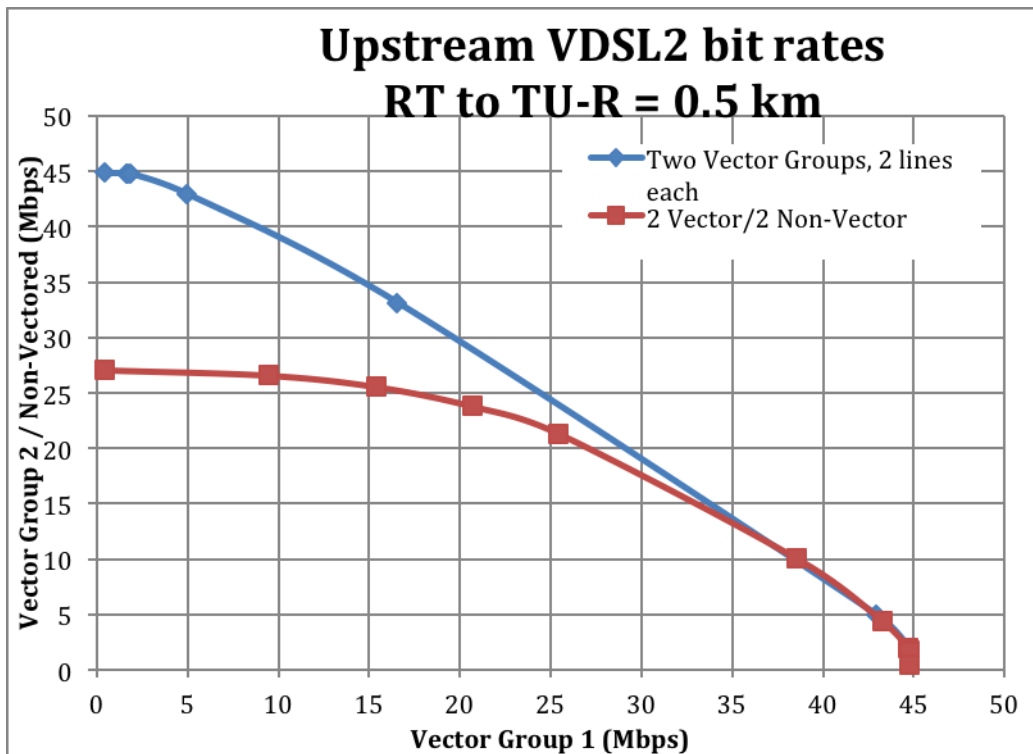
Figure II-8 - Same as Figure I-9, but for upstream.

II.6   Conclusions

This appendix showed achievable rate regions for some cases of deployments where multiple vectored group coexist in the same cable binder. These regions show trade-offs with similar trends to the case of mixed deployment. However, it is not feasible for both vectored groups to achieve 100 Mbps downstream data rate simultaneously and a low data rate cap would need to be put on the second vector group to allow vector-like data rates for the lines in the first vector group. Basically, in order to achieve vector-like data rates on the first vector group it is necessary to operate the lines in the second group at non-vectored data rates. This tradeoff is worst on shorter loops, and improves on longer loops where vectoring is less effective.

It can be concluded that that the trade-offs for the multiple vector case are harsher than for the mixed deployment case where vectored lines can still achieve vector-like performances in the presence of non-vectored lines when those are properly managed.

March 2014   71 of 101

Appendix III. DSM Iterative Waterfilling (IWF) Simulation Results - Mix of Vectored and Non-vectored Lines (Distributed Length Case)

III.1   Introduction

This appendix explores mixed deployments of vectored and non-vectored VDSL2. Simulations use the DSM spectral optimization technique known as Iterative Waterfilling (IWF) [See Appendix I]. These simulations present cases of distributed users. For distributed users, a subset of the rate regions is found which shows trade-offs between vectored and non-vectored lines and also trade-offs between lines of different lengths. Upstream results are presented both with and without UPBO for comparison.

This appendix builds on the results of [See Appendix I] (collocated case) by adding results based on IWF, which is computationally much simpler than OSB, works in a distributed manner, and does not require centralized control. A more complicated loop topology is simulated with distributed users.

Results show that trade-offs in these distributed-user cases are similar to trade-offs previously shown for collocated-user cases [See Appendix I].

III.2   VDSL2 Bit Rate Calculations

The simulations here calculate VDSL2 downstream and upstream bit rates in mixed deployments of vectored VDSL2 and non-Vectored VDSL2.

The simulations use the models and parameters as used previously, including the BT loop model, which are generally found in the NICC DSM technical report [7]. Downstream VDSL2 Profile 17a is simulated. The transmit PSD is at most 3.5 dB below the VDSL2 profile 998ADE17-M2x-A PSD limit mask defined in Annex B of G.993.2 [1]. Additionally, transmit PSDs are limited to at most -55 dBm/Hz to meet the G.993.2 average power constraint.

The margin is 6 dB and the total coding gain is 3 dB. Bit rates are calculated by summing the capacity calculation of each 4.3125 kHz tone with a 9.75 dB SNR gap, with bits per sub-carrier lower limited to at least one bit and upper limited to 14 bits per sub-carrier. As defined in T1.417, for VDSL2 6% guard-band is imposed at the edge of each passband below 12 MHz and these guard-bands carry no bits. Above 12 MHz, a guard band of 175 kHz is assumed at the edge of each passband. Loops are all 0.4 mm / 26 AWG. AWGN at -140 dBm/Hz is added to each receiver.

All lines are cabinet-based. Simulation results presented here have 1% worst-case same-binder FEXT plus -140 dBm/Hz noise, with twelve vectored VDSL2 lines and twelve non-vectored VDSL2 lines. The transmit PSDs of all VDSL2 lines are shaped by IWF. It is ideally assumed that self-FEXT cancellation is perfect, so there is no residual crosstalk between the vectored lines, but the twelve non-vectored lines crosstalk into the vectored lines. For the non-vectored lines, there are eleven non-vectored crosstalkers and twelve vectored crosstalkers. The worst-case FEXT into the non-vectored and the vectored lines is FSAN summed. It is assumed that the non-vectored lines and the vectored lines all originate at the same cabinet.

March 2014   72 of 101

The upstream results with UPBO use parameters as defined in the UK Access Network Frequency Plan (ANFP) V5.1.1, with A1=60, A2=60, B1=21, and B2=8, and kl0 is ideally based on the length of the loop that each system transmits over [7].

III.3   Distributed Users

This section presents upstream and downstream results for a situation with distributed users in a binder. These simulations model cases of distributed users as shown in Figure III- 1, except that the simulations have 24 users, not the 6 used for illustrative purposes in the figure. All lines originate at the same location and the shortest vectored line and the shortest non-vectored line are both of length d0 meters. Longer lines are also present in pairs such that there is one vectored and one non-vectored line each of lengths d0+d, d0+2d, …, d0+11d.
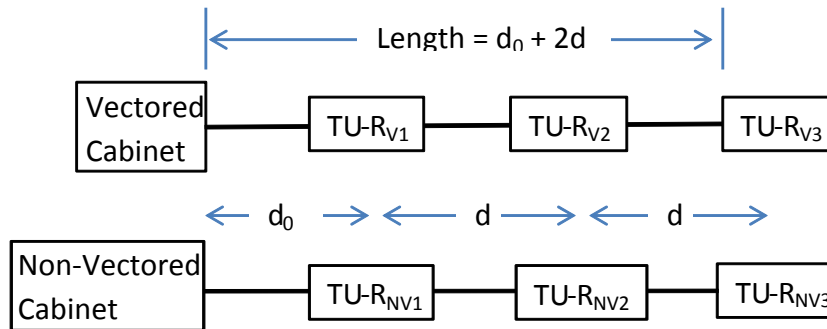
**Figure III-1 - Distributed user scenario: Vectored and non-vectored crosstalkers.**

This appendix presents results for distributed users with Iterative Waterfilling (IWF), which is a DSM algorithm that finds the best transmit PSDs for given bit rate trade-offs. While OSB is computationally intensive and generally suited to centralized DSM implementations, IWF is an alternative algorithm that can run in a non-centralized, distributed, manner that comes very close to meeting the OSB bit rates [6].

III.4   Distributed Rate Region Representation

This section explains how subsets of rate regions for distributed length binder cases have been found and plotted for good visual representation. The rate region for a 24 user distributed case is 24 dimensional and so it can only be plotted for a given set of trade-offs. A subset of the rate region is obtained by constraining a group of lines to certain fixed rates and letting the other lines do rate adaptive waterfilling in presence of crosstalk seen by other lines. The rates of different lines can then be plotted in pairs, with connected curves for simultaneously achievable rate-tuples, giving a 'spider-web' like rate-region plot.

In the simulation results presented below, the rates achievable by IWF are shown. Other DSM algorithms can also be used to plot rate-region subsets attainable by their use.

III.5   Upstream Distributed Results
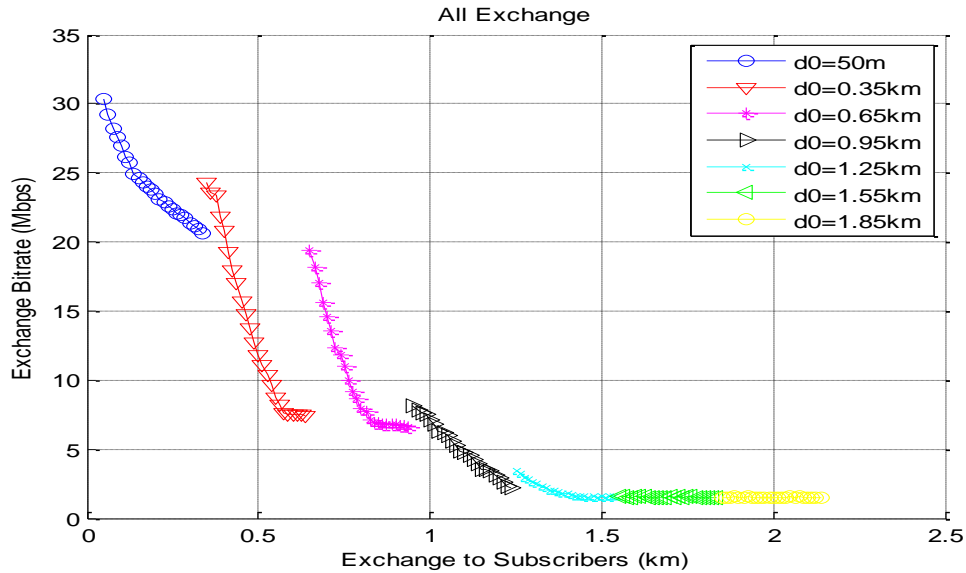
Figure III-2 - Upstream bit rates, distributed subscribers with d = 25m, all lines originating at the same network end-point (exchange or cabinet) with static UPBO.

Figure III-2 shows cases of increasing d0 with static UPBO for the case of all non-vectored VDSL2 lines originating at the same network end-point. This shows that bit rates decrease rapidly with loop length within any one route, with 24 worst-case FEXT.

III.6   Upstream Distributed Results with No UPBO

Figure III-3 shows the IWF rate region for the case of d0 = 300m and d = 25 m. Each rate region curve represents the trade-off between the rates of exactly one vectored and one non-vectored line for the length indicated in the legend; that is, these 2 lines, one vectored and one non-vectored, have the same length. Thus, the outermost curve shows the rate trade-off between the shortest vectored and non-vectored lines of length 300m, while the innermost curve draws the rate curve for the longest lines of length (d0+11d =) 575m. However, these curves are valid only for the particular rates and PSDs for all the other lines in the binder. The dotted blue lines illustrate the simultaneously achievable rate points for all the 24 lines. Each such dotted blue line intersects each rate region at one point, and the collection of such points (rate pairs for non-vectored and vectored lines of same length) is the overall 24-dimensional rate vector for the 24 users. Each dotted blue line thus represents a possible set of DSM IWF rates.
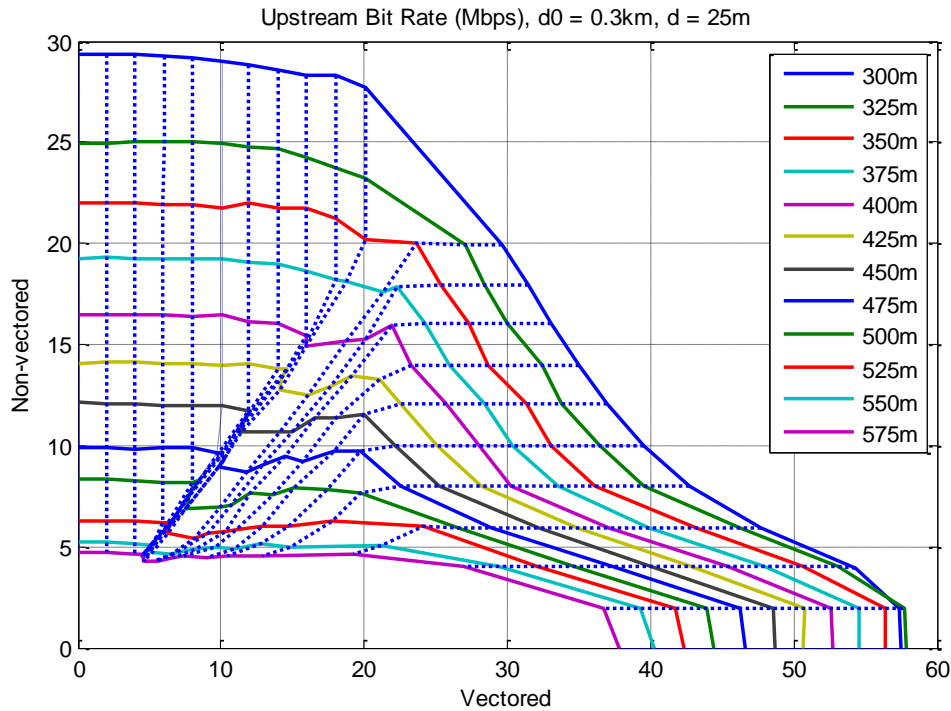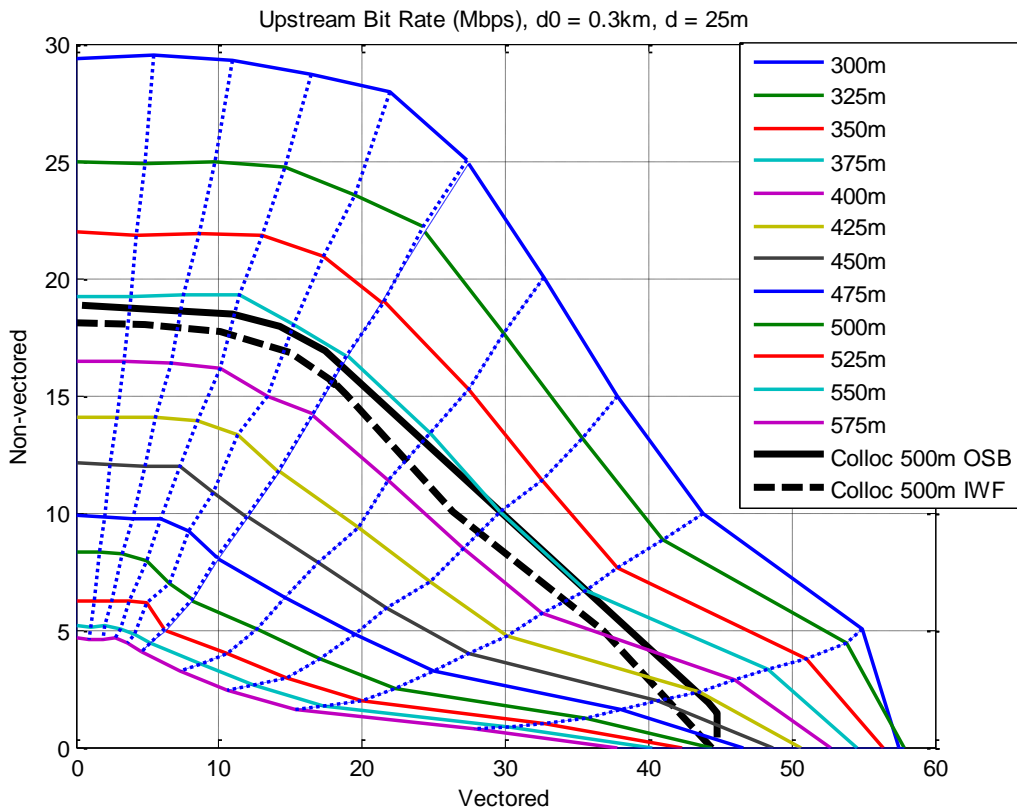
Figure III-3 - Vectored and non-vectored upstream bit rates with IWF, distributed end-points, d0 = 300m, d = 25m, equal rate targets.

The rate-region in Figure III-3 is found by constraining half the lines (either all vectored lines or all non-vectored lines) to equal rate targets. Thus the dotted blue lines are either horizontal or vertical correspondingly. However, if the equal rate cannot be met beyond a certain length, then these dotted blues lines do sometimes deviate below that equal rate. The straight horizontal dotted blue lines represent the cases where all the non-vectored rates are given certain equal rate targets. The straight vertical lines on the other hand represent the cases where all the vectored lines are given equal rate targets that vary between multiple iterations. Again, the fact that the lines in the middle of the plot do not remain straight means that the rate targets assigned are too high to be simultaneously achievable. The lines that cannot achieve their rate targets then converge to whatever best rates they can achieve, given the crosstalk from other lines, using IWF.

Figure III-3 implies for instance that vectored lines could not offer an upstream speed (uniformly) of more than 38 Mbps even if all were vectored and the non-vectored lines would obtain only 2.5 Mbps upstream. If the non-vectored lines are however polite with 5 Mbps, then vectored lines could still all achieve more than 30 Mbps. If the non-vectored lines were polite with 7.5 Mbps, then all of the vectored lines except 3 would still also get 30 Mbps upstream, and those last 3 would still get at least 20 Mbps. This illustrates that vectored lines can achieve most of the high speed they promise in an unbundled environment where non-vectored lines simply operate politely at current operating speeds.

Figure III-4 shows another rate region curve for the same scenario, where the rate targets are differently assigned. Instead of assigning equal rates to all the fixed-rate lines, linearly increasing rate targets are used, i.e. higher rate targets for shorter lines and lower rate targets for longer lines.

This is evident from the fact that the dotted lines are diagonal instead of straight in Figure III- 4. The figure shows that although the simultaneously achievable rate tuples (represented by the blue dotted lines) change, but the region around the diagonal and the outermost curve remain almost the same as that in Figure III- 3. This shows that this method of finding and plotting subsets of the rate-region for multiple, distributed-length, binder scenarios can provide a good representation of overall behaviour of all lines in the binder, and varies slightly with the exact rate targets assigned to any group of lines. The figure also plots collocated OSB and IWF curves for 12 vectored and 12 non-vectored 500m lines for comparison. These two curves are plotted in most of the following figures, for the sake of comparison.

Figure III-4 - Vectored and non-vectored upstream bit rates with IWF, distributed end-points, d0 = 300m, d = 25m, linearly varying rate targets.

Figure III-4 implies for instance that that vectored lines can achieve roughly 10 times the upstream data rates of non-vectored lines if those non-vectored lines operate politely at 3 Mbps or less. If the 3 longest non-vectored lines run politely at 5 Mbps or less, and the remaining upstream non-vectored lines run at 8 Mbps (or more as indicated with rate increasing with length), then all the vectored lines will run roughly 2.5 times faster than non-vectored lines at the respective same lengths.

Figure III-3and Figure III- 4 thus illustrate how management can be used with politeness to offer very large vectoring improvements while still maintaining current non-vectored VDSL2 data rates upstream.

Figure III-5 shows similar results for d0 = 500m for linearly varying rate targets while Figure III-6 shows similar results for d0 = 1000m.

The results show the trade-off between rates of vectored and non-vectored lines for different lengths. These plots also show how the rates vary with line-length within the same group (vectored or non-vectored). It is evident from the results that crosstalk within the same group for non-vectored lines causes rates of the long non-vectored lines to be considerably affected by the crosstalk from the short lines in the binder. This is where techniques such as OSB or any of its approximations like MLWF or ISB can be used to increase the rates of longer lines at the cost of lowering the short lines' data rates. The ANFP mandates the use of UPBO to protect the long lines from the harmful crosstalk of short in such a case. The next section therefore plots the UPBO curves for comparison.

Figure III-5 - and non-vectored upstream bit rates with IWF, distributed end-points, d0 = 500m, d = 25m, equal rate targets.
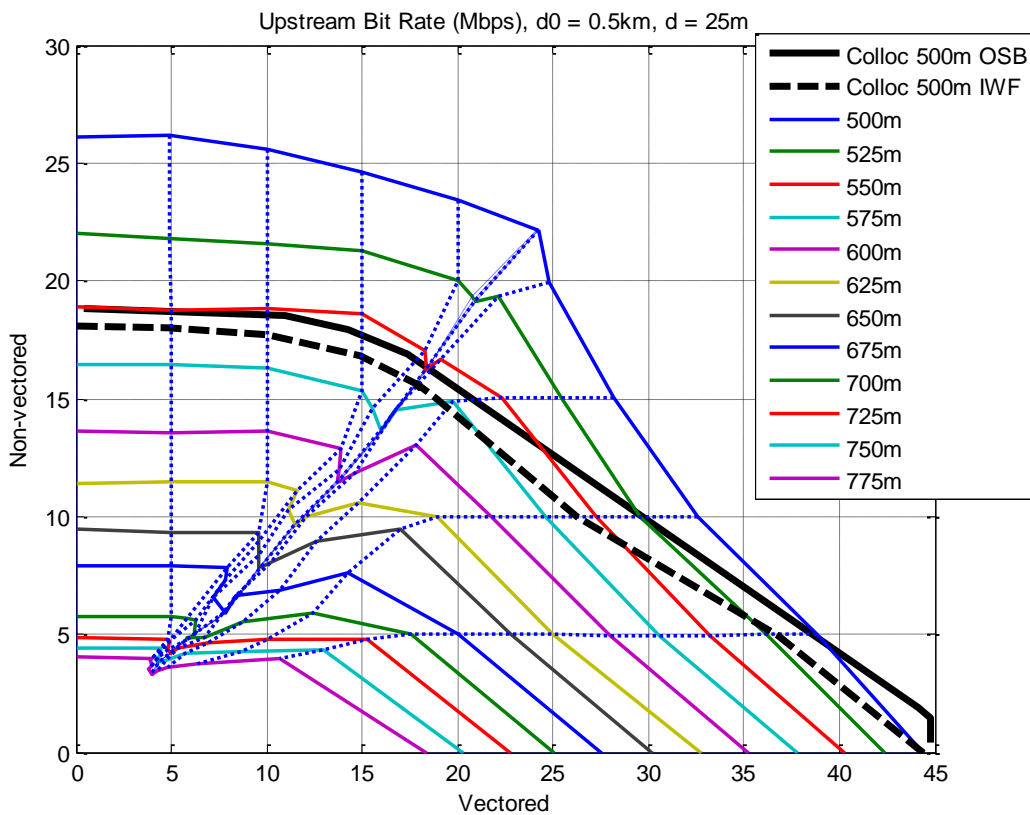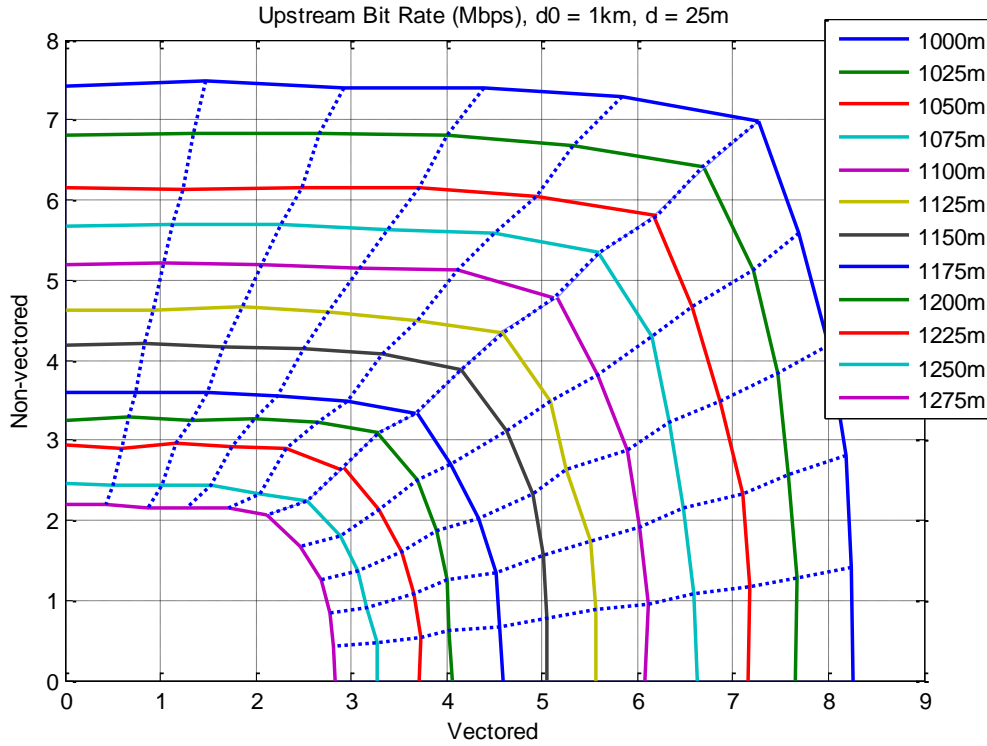
Figure III-6 - Vectored and non-vectored upstream bit rates with IWF, distributed end-points, d0 = 1000m, d = 25m, linearly varying rate targets.

Figure III-5 and Figure III-6 on the dotted blue lines show again that vectored VDSL2 can still offer significantly higher data rates than current non-vectored VDSL2 data rates if non-vectored VDSL2 uses simple distributed polite IWF at rates no greater than what those systems run at today without vectoring.

III.7   Downstream Distributed Results

Figure III-7 and onward show the rate region for the downstream case for similar scenarios as the upstream results presented earlier. Data rates for various values of d0 are shown with d = 25m. These results use equal rate targets for the fixed rate lines.

Figure III-7 - Vectored and non-vectored downstream bit rates with IWF, distributed end-points, d0 = 300m, d = 25m, equal rate targets.

Figure III-7 illustrates that best case highest speed offering for all lines vectored is 80 Mbps (the 575m line). This speed is 3x current non-vectored VDSL2 speeds of roughly 25 Mbps. However, if these non-vectored 12 existing unbundled (separate service provider lines) continue to run politely at 25 Mbps, then the vectored lines will still all get at least 70 Mbps, still roughly 2.7x the non-vectored speed. Again unbundling allows large gain with vectoring if the system is well managed.
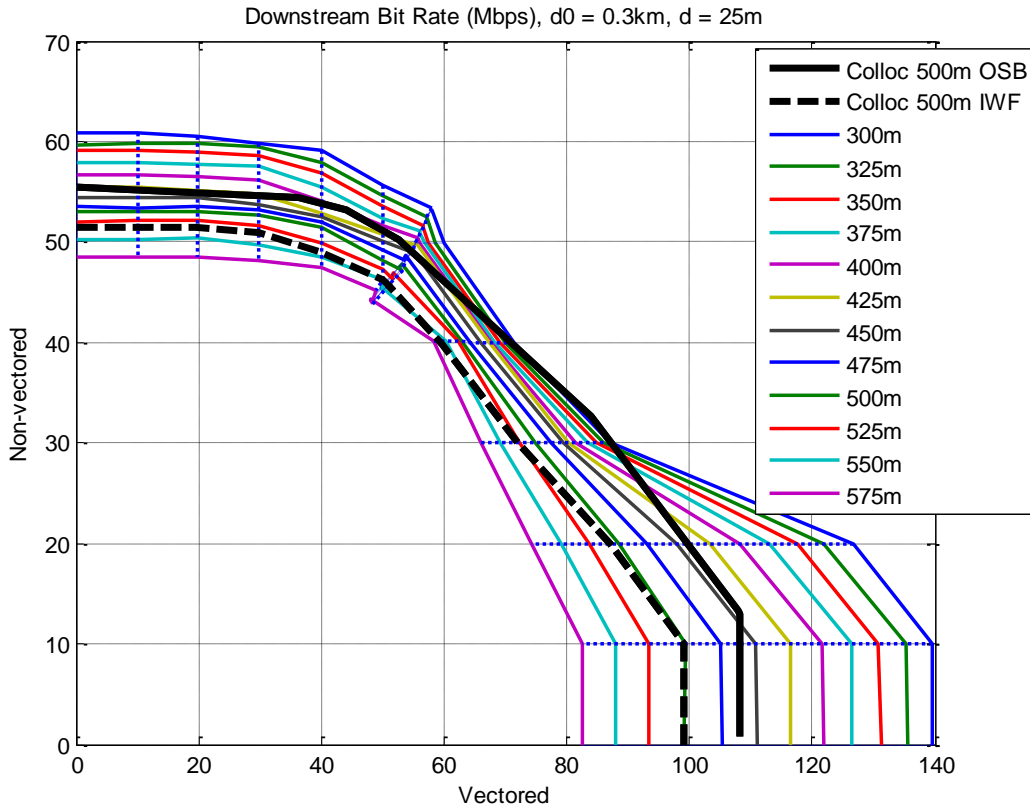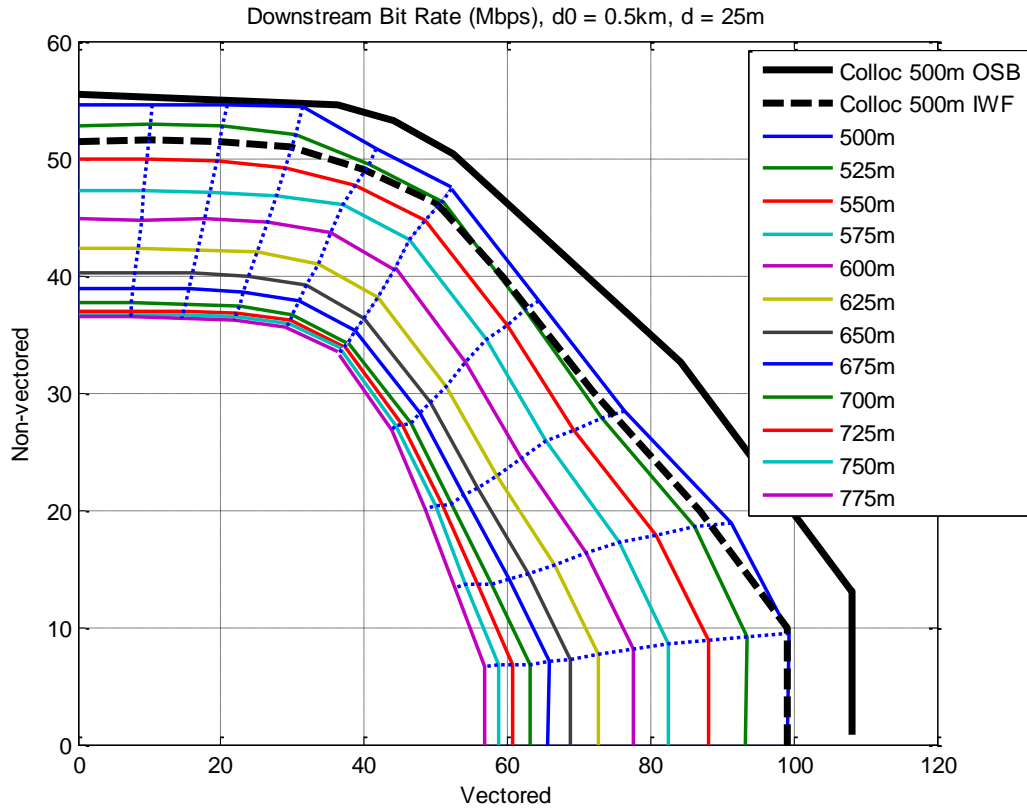
Figure III-8 - Vectored and non-vectored upstream bit rates with IWF, distributed end-points, d0 = 500m, d = 25m, linearly varying rate targets.

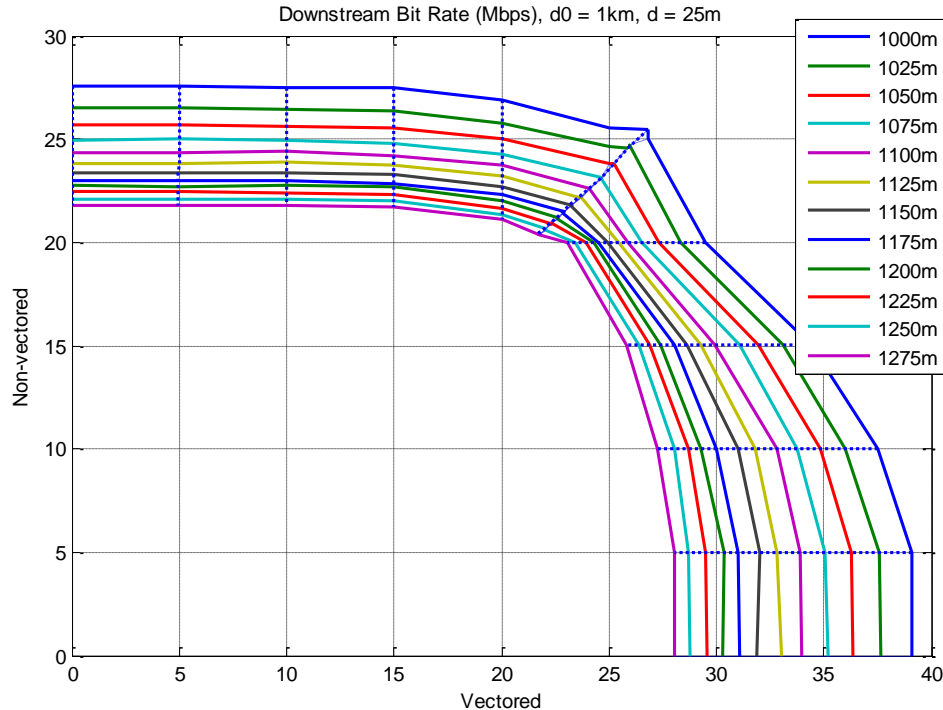**Figure III-9 - Vectored and non-vectored upstream bit rates with IWF, distributed end-points, d0 = 1000m, d = 25m, equal rate targets.**

The results show that the non-vectored lines can get their current rates while only affecting the rates of the vectored lines to a small extent. However, certain politeness rules need to be followed by each line as a matter of common politeness, which is run at no more than current speeds and limit maximum margin to a reasonable value (say 6-10 dB) in loading with water-filling.

III.8   Conclusions

This appendix showed achievable rate regions for some cases of deployments where a vectored group and a set of non-vectored lines share the same cable binder. This appendix simulated a realistic case of distributed users representing a typical street layout.

Worst-case 24-disturber crosstalk, low background noise, and perfect vectoring cancellation were simulated, accentuating the impact of the 12 non-vectored worst-case crosstalkers on the vectored lines and the impact of the 23 worst case crosstalkers on the non-vectored lines. This appendix uses an easily implementable, distributed DSM technique. The rate regions show good trade-offs, and in particular show that the non-vectored lines can have little impact on the vectored group even under worst case conditions, if the non-vectored lines are properly managed. Rate region trade-offs with distributed users are very similar to trade-offs shown for collocated users [See Appendix I].
The overall conclusion is that it is technically feasible for mixed deployments of vectored and non-vectored lines to successfully coexist, and it was shown here that this can be implemented via Dynamic Spectrum Management (DSM).

Appendix IV. DSM Iterative Waterfilling (IWF) and Flat Power Back-Off (PBO) Simulation Results - Mix of Vectored and Non-vectored Lines

IV.1   Introduction

This appendix explores mixed deployments of vectored and non-vectored VDSL2. Some simulations use the DSM spectral optimization technique known as Iterative Waterfilling (IWF) [9],[6], and some simulations use flat Power Back-Off (PBO). The simulations here are for cases of collocated users. Simulation conditions are the same for IWF and PBO so that performance of the two algorithms can be directly compared.

This appendix builds on previous results of [See Appendix I and III] by adding results using just flat PBO, which is easy to implement. The simulations determine and use the minimum transmit power that allows a given line to attain a given target bit rate with 6 dB margin, etc. The target bit rates are varied to trace out an achievable rate region.

Results show that the performance of simple PBO practically identical to the performance of IWF.

IV.2   VDSL2 Bit Rate Calculations

The simulations here calculate VDSL2 downstream and upstream bit rates in mixed deployments of vectored VDSL2 and non-Vectored VDSL2.

The simulations use the models and parameters as used previously, including the BT loop model, which are generally found in the NICC DSM technical report . Downstream VDSL2 Profile 17a is simulated. The transmit PSD is at most 3.5 dB below the VDSL2 profile 998ADE17-M2x-A PSD limit mask defined in Annex B of G.993.2 [1]. Additionally, downstream transmit PSDs are limited to at most -55 dBm/Hz and upstream transmit PSDs are limited to at most -44 dBm/Hz to meet the G.993.2 average power constraints.

The margin is 6 dB and the total coding gain is 3 dB. Bit rates are calculated by summing the capacity calculation of each 4.3125 kHz tone with a 9.75 dB SNR gap, with bits per sub-carrier lower limited to at least one bit and upper limited to 14 bits per sub-carrier. As defined in T1.417, for VDSL2 6% guard-band is imposed at the edge of each passband below 12 MHz and these guard-bands carry no bits. Above 12 MHz, a guard band of 175 kHz is assumed at the edge of each passband. Loops are all 0.4 mm / 26 AWG. AWGN at -140 dBm/Hz is added to each receiver.

All lines are cabinet-based with collocated customers, so all non-vectored lines and vectored lines originate at the same location and terminate at the same distances. There is no UPBO. FEXT is generated by the ATIS MIMO FEXT coupling model in ATIS-0600024 [11], and only same-binder FEXT couplings are used here. The drawn couplings are increased by 1 dB to align the ATIS MIMO model 1% worst-case with the ATIS 1% worst-case crosstalk model for multiple disturbers. The transmit PSDs of all VDSL2 lines are either shaped by IWF, or lowered by the same amount of dBs at every frequency with PBO. It is assumed that self-FEXT cancellation is perfect, so there is no residual crosstalk between the vectored lines.
The level of IWF or PBO is found that allows either non-vectored lines or vectored lines to achieve a target bit rate with 6 dB margin. The target bit rates are varied, and each target bit rate creates a

point on the plots presented here of the achievable rate regions. One thousand cases of different randomly-drawn crosstalk couplings are run, and for each case the bit rates of all lines are determined. Plots show the 1% worst-case bit rates and the median bit rates. Cases are shown with 2 vectored and 2 non-vectored lines (2x2) in the same binder, and with 12 vectored and 12 non-vectored lines (2x2) in the same binder.

IV.3   Downstream Results

Figure IV-1 - Downstream 1% worst-case bit rates on 500m loops, 2 vectored and 2 non-vectored lines.
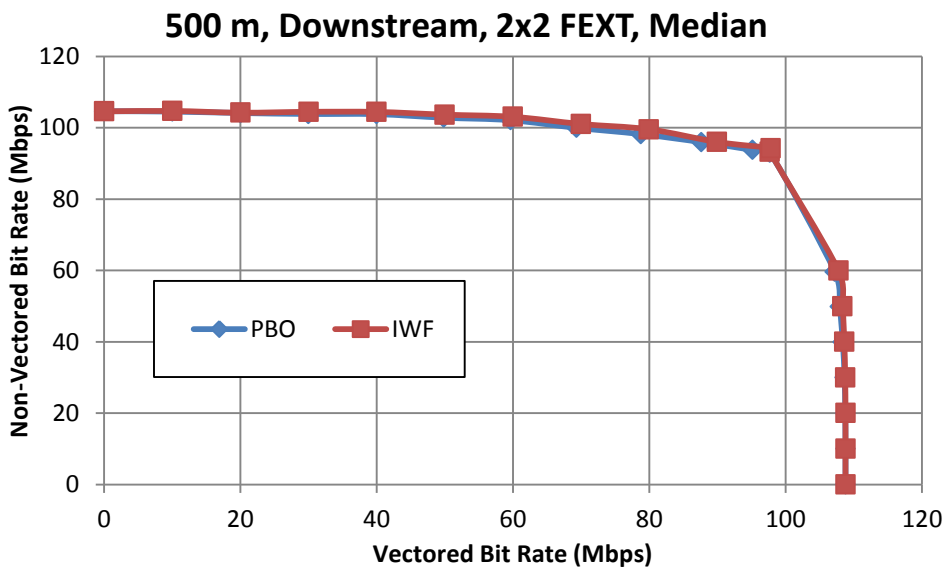
Figure IV-2 - Downstream Median bit rates on 500m loops, 2 vectored and 2 non-vectored lines.
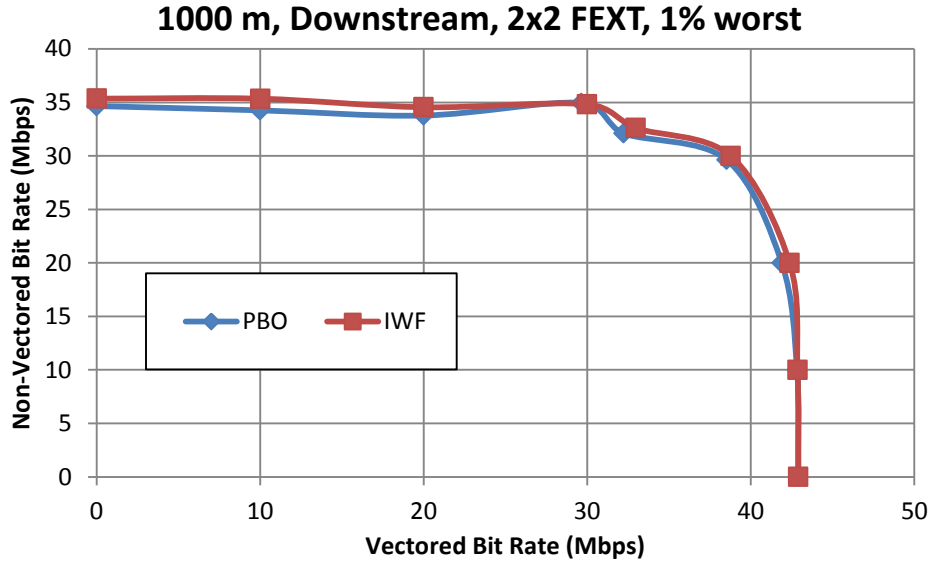
Figure IV-3 - Downstream 1% worst-case bit rates on 1000m loops, 2 vectored and 2 non-vectored lines.

Figure IV-4 - Downstream 1% worst-case bit rates on 500m loops, 12 vectored and 12 non-vectored lines.

500 m, Downstream, 12x12 FEXT, Median

Figure IV-5 - Downstream Median bit rates on 500m loops, 12 vectored and 12 non-vectored lines.

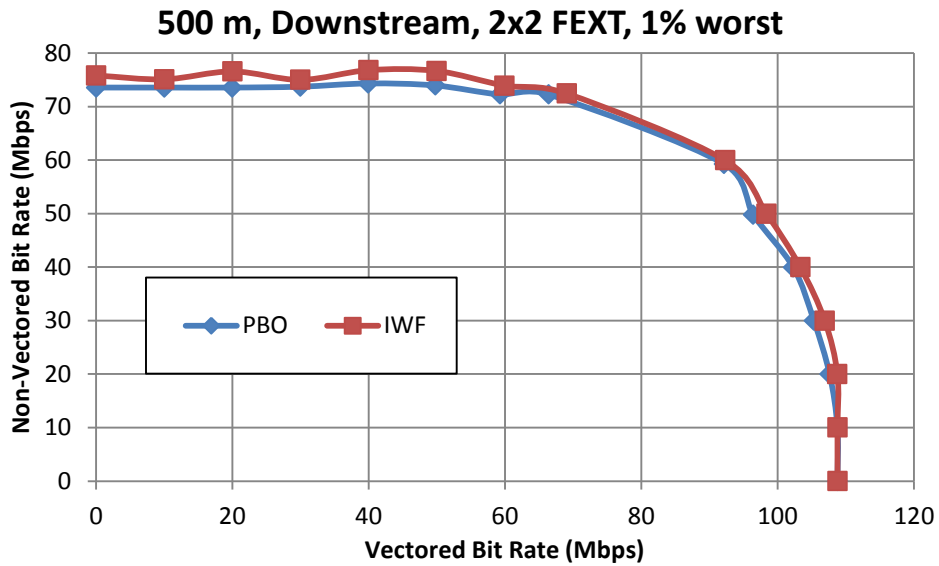## IV.4   Upstream Results

300 m, Upstream, 2x2 FEXT, 1% worst
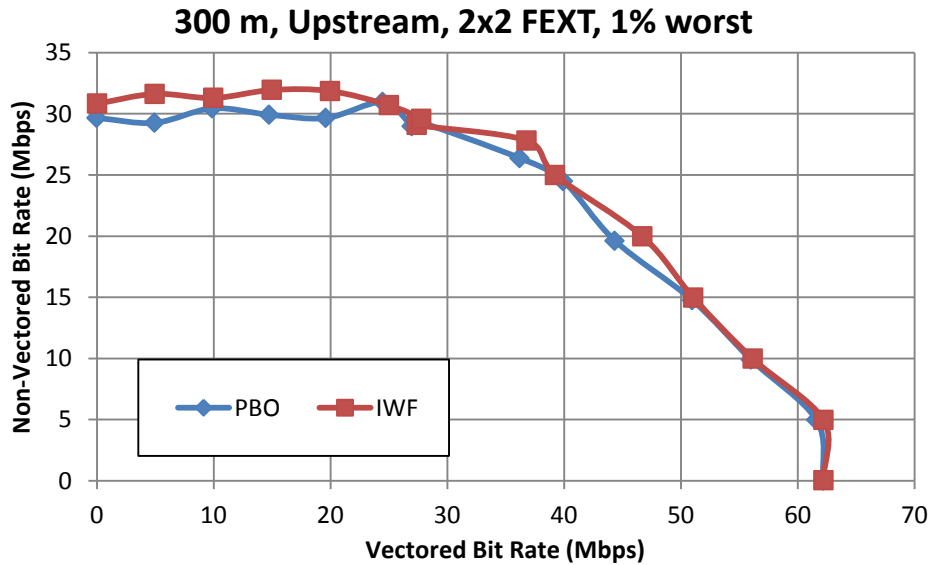
Figure IV-6 - Upstream 1% worst-case bit rates on 300m loops, 2 vectored and 2 non-vectored lines.

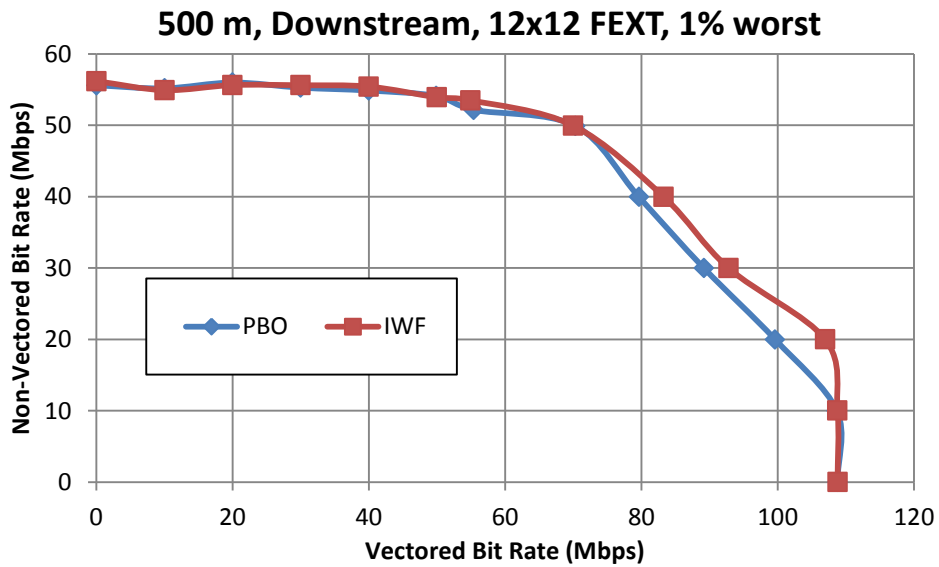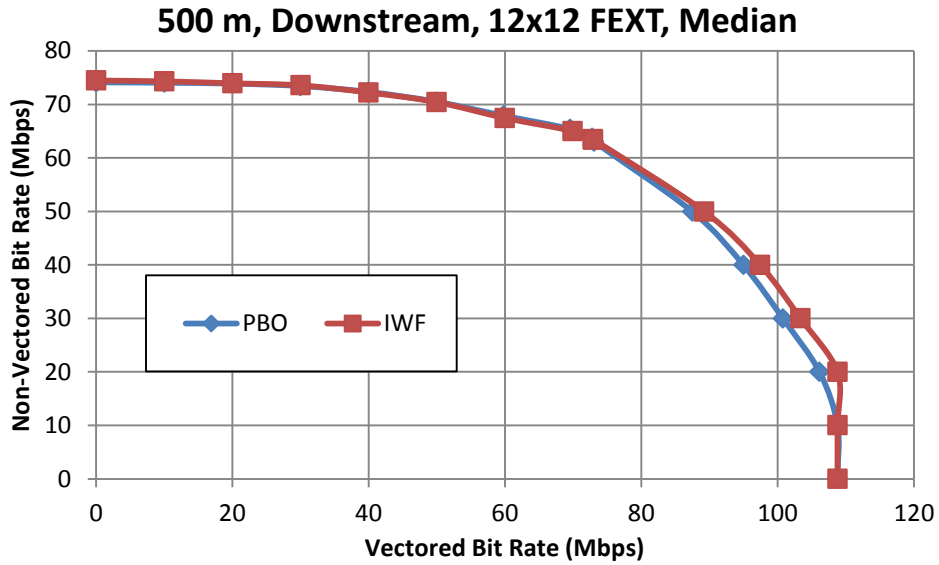**700 m, Upstream, 2x2 FEXT, 1% worst**
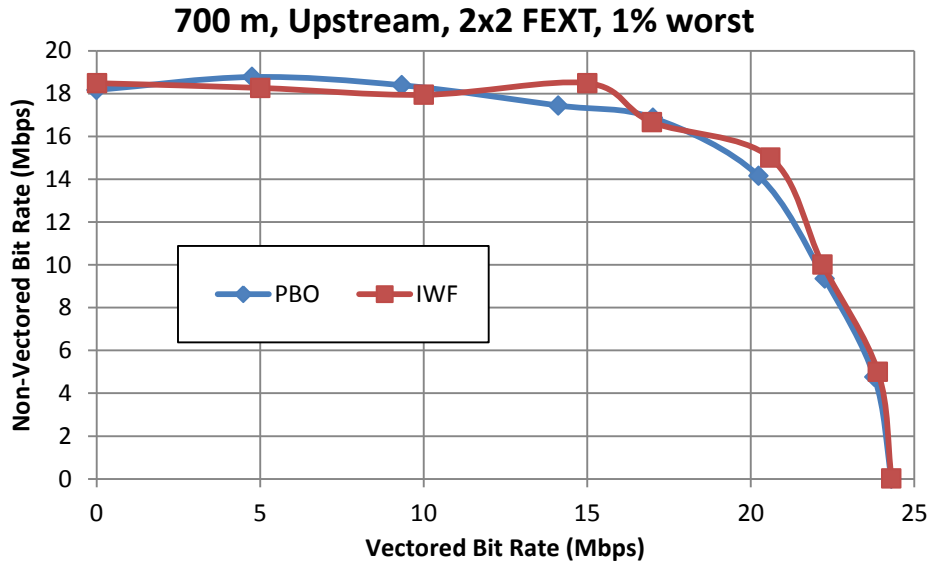
Figure IV-7 - Upstream 1% worst-case bit rates on 700m loops, 2 vectored and 2 non-vectored lines.

500 m, Upstream, 12x12 FEXT, Median
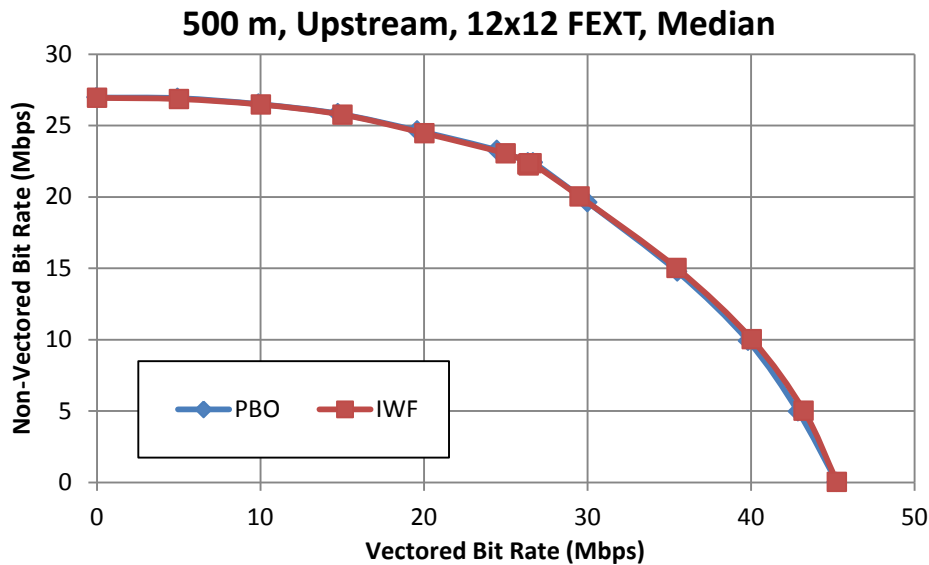
Figure IV-8 - Upstream Median bit rates on 500m loops, 12 vectored and 12 non-vectored lines.

IV.5   Conclusions

This appendix showed that the performance of Iterative Water-Filling (IWF) and flat Power Back-Off (PBO) are essentially identical. Therefore, the spectral compatibility trade-offs between between vectored and non-vectored VDSL2 shown in previous appendixes [See Appendix I and III] can be readily implemented with simple flat PBO.

 Both IWF and flat PBO are both simple DSM techniques that can be implemented by modems autonomously and do not require a centralized Spectrum Management Center (SMC) nor message passing between users.

Appendix V.   DSM Iterative Waterfilling (IWF) Simulation Results - Mix of Vectored and Non-vectored Lines (Collocated Case)

V.1   Introduction

This appendix explores mixed deployments of vectored and non-vectored VDSL2. Simulations are run using the DSM level-2 spectral optimization technique known as Iterative Waterfilling (IWF) [13]. These simulations assume that all customers are collocated. The achievable rate region is found, showing the trade-offs between vectored lines and non-vectored lines that can be achieved by spectral optimization with ideal assumptions and also when some non-idealities are present (higher noise level, imperfect crosstalk cancellation).

V.2   VDSL2 Bit Rate Calculations

The simulations here calculate VDSL2 downstream bit rates with both vectored VDSL2 and non-Vectored VDSL2.

Downstream VDSL2 Profile 17a is simulated. The transmit PSD is at most 3.5 dB below the VDSL2 profile 998ADE17-M2x-A PSD limit mask defined in Annex B of G.993.2 [1]. The margin is 6 dB and the total coding gain is 3 dB. Bit rates are calculated by summing the capacity calculation of each tone with a 9.75 dB SNR gap, with bits per Hz per sub-carrier upper limited to 14 bits per Hz per sub-carrier. As defined in T1.417 [8], for VDSL2 6% guard-band is imposed at the edge of each passband below 12 MHz and these guard-bands carry no bits. Above 12 MHz, a guard band of 175 kHz is assumed at the edge of each passband. Loops are all 0.4 mm/26 AWG.

Simulations have 12 vectored VDSL2 and 12 non-vectored VDSL2 lines in the same binder. All lines are of the same length, and different loop lengths have been simulated. For each length case, simulation results with different noise and crosstalk assumptions have been presented.

The following loops lengths have been simulated: 150m, 300m, 500m, 700m..

The following two background noise cases have been considered: -140dBm/Hz, -130dBm/Hz.

The ATIS MIMO FEXT model has been simulated , using only same binder FEXT couplings. The ATIS MIMO model simulations use a full binder of 12 vectored and 12 non-vectored lines.

The following two cases of vectoring performance have been considered:
- Perfect self-FEXT cancellation on all vectored lines.
- Self-FEXT cancellation of 20dB on all vectored lines, which lowers each self-FEXT source by 20 dB.

Starting with the classic case of -140dBm/Hz, worst case crosstalk model and perfect self-FEXT cancellation between vectored lines as simulated in previous appendixes, the simulation assumptions made here progressively become more realistic by:
- Using the MIMO crosstalk model instead of the worst case model
- Increasing the noise floor
- Considering imperfect self-FEXT cancellation between vectored lines

All simulation results calculate achievable rates using the Iterative Water-filling (IWF) procedure. IWF is a distributed DSM technique where DSL modems operate autonomously and that does not require a centralized Spectrum Management Center (SMC) nor any message passing between users [13].
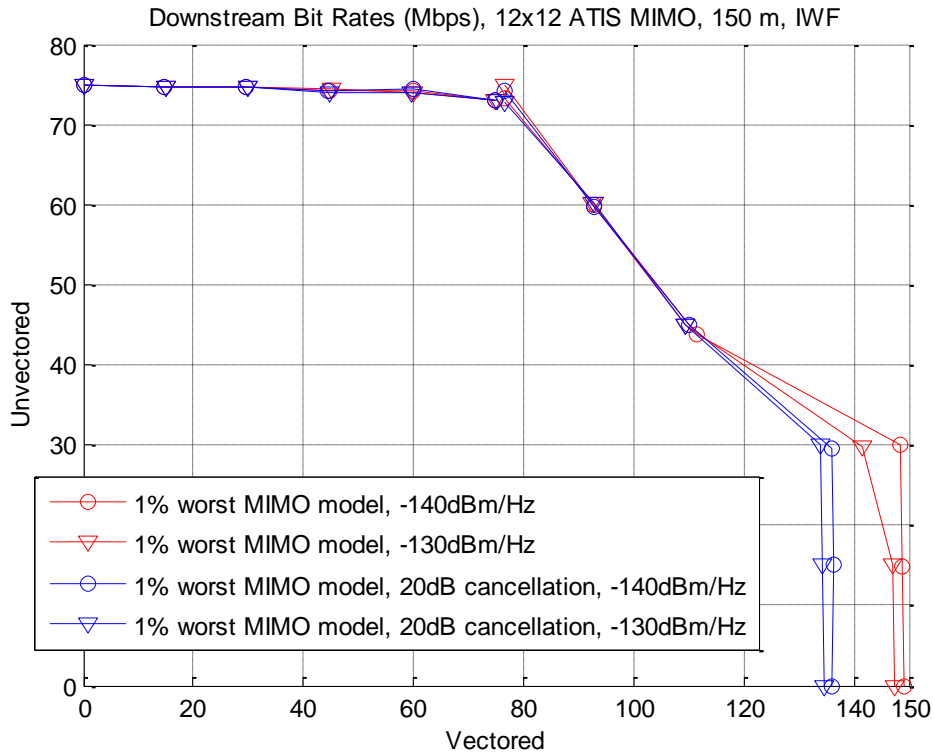
V.3   Rate Regions

Figures V-1 to V-5 show the rate region plots for the different loop lengths considered. The rate regions are calculated by fixing the rates of either the non-vectored or the vectored lines and letting the other group of lines do Rate Adaptive Water-filling. All lines adjust their spectra iteratively using the IWF technique. For the MIMO model, a total of 50 different crosstalk permutations are run giving 600 samples of vectored data rates and 600 samples of non-vectored data rates. Out of these, the 1% worst data rates are picked and plotted.

The figures show, as expected, that performance degrades for both vectored and non-vectored lines when a higher noise PSD is assumed, except for the case of the shortest 150m lines, where the performance in almost similar. A comparison of the more realistic -130dBm/Hz results with -140dBm/Hz results, shows how the trade-off for vectoring lines coexisting with non-vectored lines, is actually better than it appears with the use of a lower noise PSD.

The results also show comparison between perfect crosstalk cancellation and imperfect 20dB crosstalk cancellation. Except for the 150m and 300 m cases, there is little difference between perfect and imperfect cancellation which means that 20dB is sufficient to cancel out the crosstalk in cases of longer loop length. However, the 150m and 300m results indicate that when more realistic assumptions are considered, the tradeoff between vectored and non-vectored rates is better than otherwise expected.

March 2014   88 of 101

Figure V-1 - Rate Regions (Mbps) with IWF at 150m, for 12 Vectored and 12 Non-Vectored lines. The 1% worst data rates are plotted.
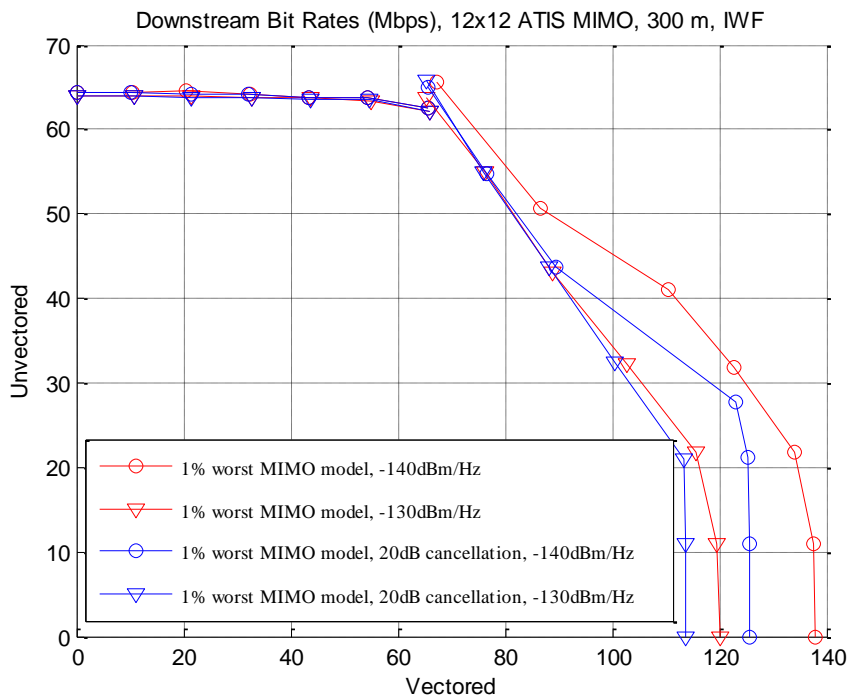
Figure V-2 - Rate Regions (Mbps) with IWF at 300m, for 12 Vectored and 12 Non-Vectored lines. The 1% worst data rates are plotted.
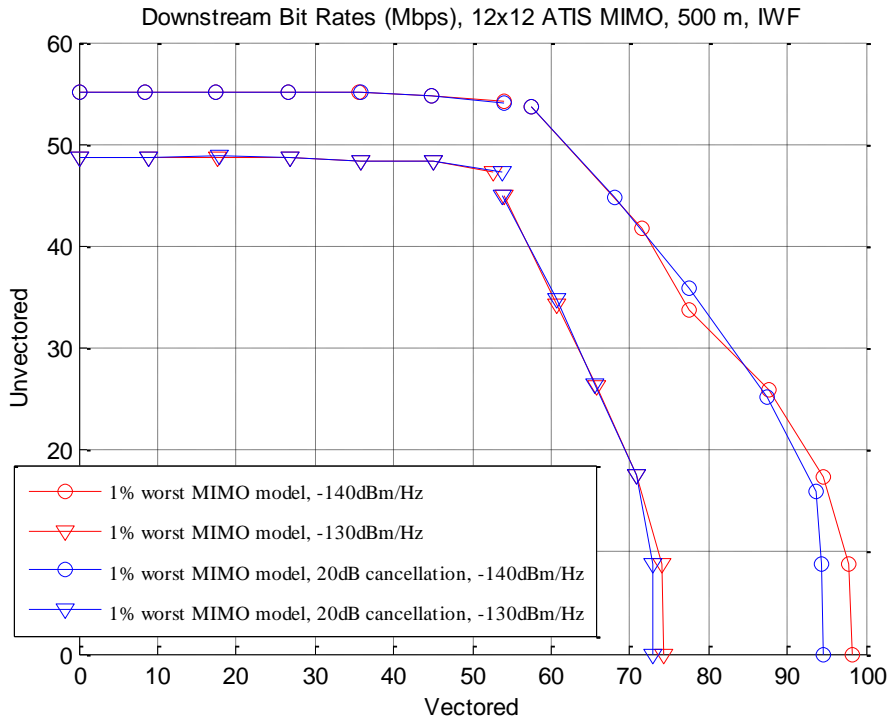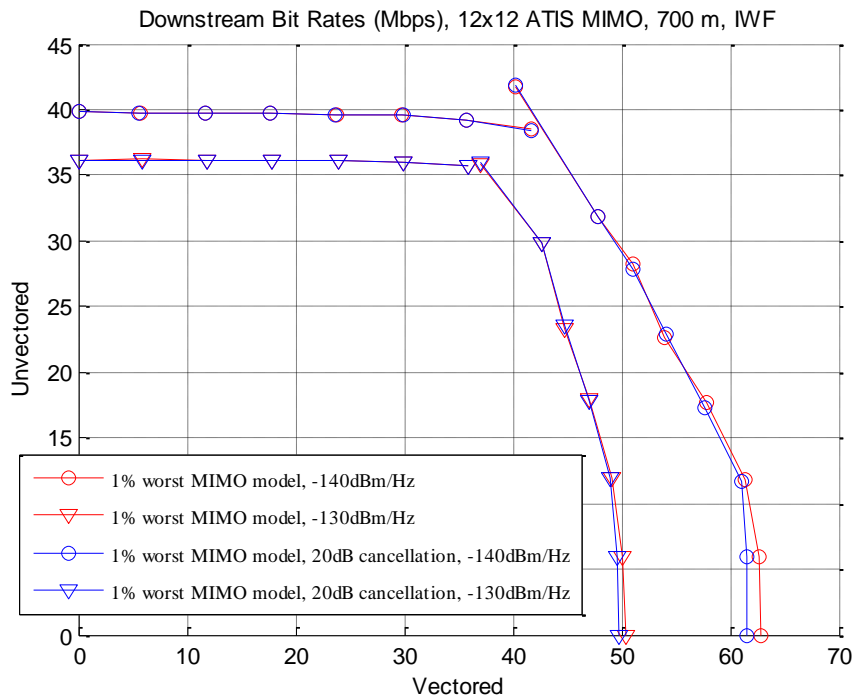
Figure V-3 - Rate Regions (Mbps) with IWF at 500m, for 12 Vectored and 12 Non-Vectored lines. The 1% worst data rates are plotted.

Figure V-4 - Rate Regions (Mbps) with IWF at 700m, for 12 Vectored and 12 Non-Vectored lines. The 1% worst data rates are plotted.

Assuming a 1% worst case target rate of 100 Mbps for the vectored lines, we can use the previous results to show that non-vector lines can operate at 1% worst case speeds that are in line with some new VDSL2 offerings and also much higher than what is generally offered today. This is shown in V-I, which reports the 1% worst case data rates achieved by the non-vectored lines when the vectored lines are able to achieve a 1% worst case data rate of 100 Mbps. Both ideal (perfect self-FEXT cancellation on all vectored lines, -140 dBm/Hz AWGN) and more realistic conditions (up to 20 dB self-FEXT cancellation on all vectored lines, -130 dBm/Hz AWGN) are considered.

Table V - 1 - Data rates of non-vectored lines when vectored lines achieve 100 Mbps. All rates are 1% worst case.

Distance	Ideal conditions	Realistic conditions
150m	54 Mbps	54 Mbps
300m	45 Mbps	33 Mbps

One can also calculate the degradation of the vectoring data rates from their maximum data rate achievable when there are no non-vectored lines (with a max of 100 Mbps). Tables 2-4 show that for the case of non-vectored lines running at 30, 40 and 50 Mbps, the degradation of the vectored lines is very small at short loop lengths and grows with loop length. All data rates are 1% worst case data rates.

Table V - 2 - Data rate loss of vectored lines from their maximum achievable rate (at most 100 Mbps) when non-vectored lines achieve 30 Mbps. All rates are 1% worst case.

Distance	Ideal conditions	Realistic conditions
150m	0%	0%
300m	0%	0%
500m	15.3% (98 → 83)	12.3% (73 → 64)
700m	20.6% (63 → 50)	28.0% (50 → 36)

Table V - 3 - rate loss of vectored lines from their maximum achievable rate (at most 100 Mbps) when non-vectored lines achieve 40 Mbps. All rates are 1% worst case.

Distance	Ideal conditions	Realistic conditions
150m	0%	0%
300m	0%	8.0% (100 → 92)
500m	25.5% (98 → 73)	19.4% (72 → 58)
700m	33.3% (63 → 42)	N/A

Table V - 4 - Data rate loss of vectored lines from their maximum achievable rate (at most 100 Mbps) when non-vectored lines achieve 50 Mbps. All rates are 1% worst case.

Distance	Ideal conditions	Realistic conditions

150m	0%	0%
300m	12.0% (100 → 88)	19% (100 → 81)
500m	N/A	36.7% (98 → 62)
700m	N/A	N/A

## V.4   Histogram Results

The figures below plot histograms for the rate distribution of the vectored lines as the non-vectored lines are assigned specific rate targets. Histograms are presented for four different points in the rate region where the non-vectored lines are given fixed rate targets and the vectored lines adapt their rates in response. Plots are shown for a realistic scenario (-130dBm/Hz noise, 20dB crosstalk cancellation on all vectored lines). Histograms are reported for loop lengths of 150m and 700m.

These histograms help to show the spread of data rates for vectored lines as the non-vectored lines are allowed to have higher rates and thus provide a better view of the trade-offs that will be seen as vectored and non-vectored lines coexist. For example, Figure V-4 shows how the spread of rates for the vectored lines is almost the same if the non-vectored lines achieve 15 Mbps or 30Mbps. The pdf shifts by about 30 Mbps when the un-vectored rates are doubled, from a mean of around 140Mbps to around 110Mbps.
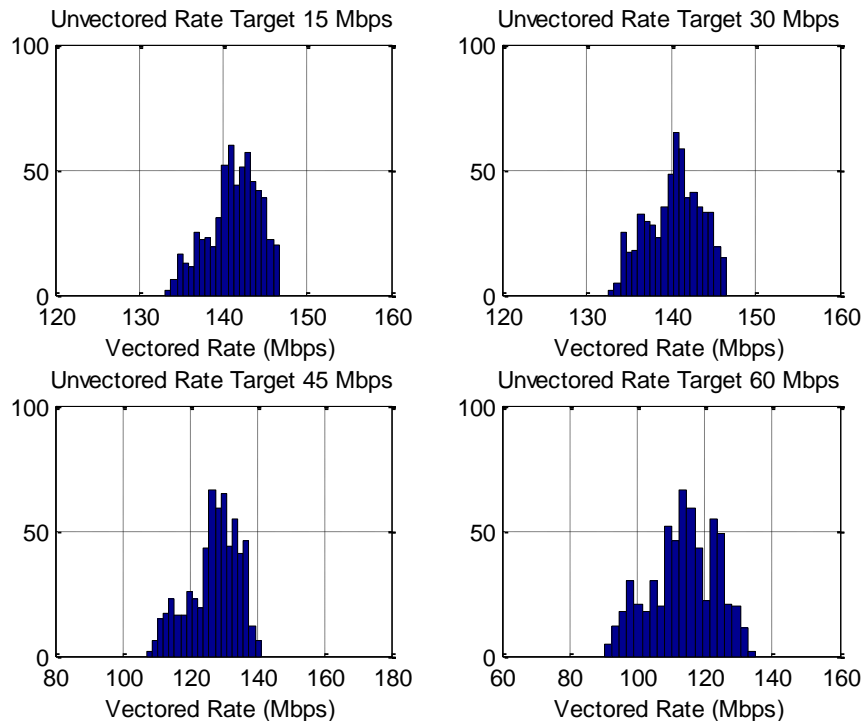
Figure V-5 - Imperfect cancellation, 150m, -130dBm/Hz.
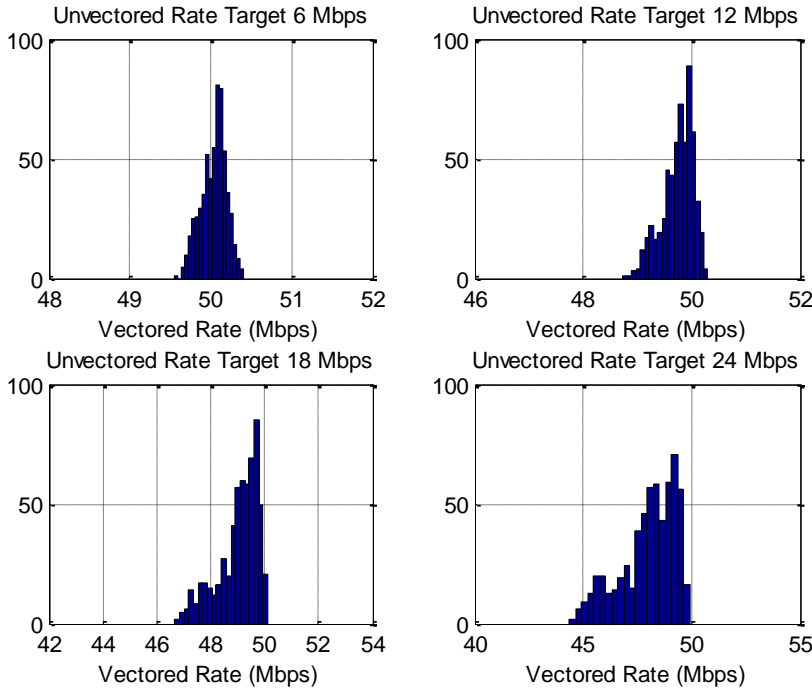
Figure V-6 - Imperfect cancellation, 700m, -130dBm/Hz.

V.5   Conclusions

This appendix showed achievable rate regions for some cases of deployments where a vectored group and a set of non-vectored lines share the same cable binder. These rate regions were calculated using IWF, a distributed DSL Level 2 algorithm that operates autonomously and does not require gaining centralized knowledge, i.e., neither a centralized Spectrum Management Center nor message exchange are needed for the algorithm execution. The simulations compared an ideal scenario (low noise at the CPE, perfect cancellation) with a more realistic one (higher noise at the CPE, imperfect crosstalk cancellation).

If non vectored lines are capped at 40 Mbps, a data rate much higher than most of today's offerings, the degradation that vectored lines would have from a target of 100 Mbps or whatever is their maximum achievable rate is 0% at 150m, 8% at 300m, and 19% at 500 m. Considering averages or higher percentiles would reduce the losses calculated here.

Statistical results were presented in the form of histograms that suggest that the assessment of the performance of vectored lines in the presence of non-vectored lines should be evaluated beyond the traditional 1% worst case data rate and that a better and more complete picture can be given when taking into consideration higher percentile points and the whole distribution of data rates.
A companion appendix [See Appendix IV] shows that rate results similar to the ones obtained with IWF can also be achieved by using a simple flat power back off scheme, i.e. a scheme that only requires controlling the transmit power of DSL modems.

The overall conclusion is that it is technically feasible for mixed deployments of vectored and non-vectored lines to coexist, and it was shown here that this can be implemented via a simple and distributed DSM Level 2 technique.

Appendix VI. DSM Iterative Waterfilling (IWF) Simulation Results - Mix of Vectored and Non-vectored Lines and Multiple Vector Groups

This Appendix explores mixed deployments of vectored and non-vectored VDSL2. Simulations are run using the DSM technique known as Iterative Waterfilling (IWF), a distributed DSM algorithm requiring little or no inter-SMC communication [13]. These simulations assume that all subscribers are collocated. Rate-reach plots are shown, showing worst case and average data rate trade-offs between vectored and non-vectored lines. Ideal (perfect crosstalk cancellation), realistic (imperfect crosstalk cancellation), and capped (max max data rate of 100 Mbps) conditions have been considered for the performance of the vectored lines

VI.1   VDSL2 Bit Rate Calculations

The simulations here calculate VDSL2 downstream bit rates with both vectored VDSL2 and non-Vectored VDSL2.

Downstream VDSL2 Profile 17a is simulated. The transmit PSD is at most 3.5 dB below the VDSL2 profile 998ADE17-M2x-A PSD limit mask defined in Annex B of G.993.2 [1]. The margin is 6 dB and the total coding gain is 3 dB. Bit rates are calculated by summing the capacity calculation of each tone with a 9.75 dB SNR gap, with bits per Hz per sub-carrier upper limited to 14 bits per Hz per sub-carrier. A 12 tone guard-band is assumed between each upstream/downstream passband and these guard-bands carry no bits. Loops are all 0.4 mm/26 AWG. Background noise is -140dBm/Hz.

Simulations have 12 vectored VDSL2 and 12 non-vectored VDSL2 lines in the same binder, for the mixed scenario case. For the two vector group case, there are two 12-line vectored groups. All lines are of the same length, and different loop lengths up to 1 km have been simulated.

The ATIS MIMO FEXT model [11], has been simulated, using only same-binder FEXT couplings, with a full binder of 24 lines.

The following three cases of vectoring performance have been considered:
 • Ideal full vectoring: perfect self-FEXT cancellation on all vectored lines.
 • Partial vectoring: self-FEXT cancellation of 20dB on all vectored lines, which lowers each self-FEXT source by 20 dB.
 • Capped ideal full vectoring: under perfect self-FEXT cancellation conditions, a maximum data rate of 100 Mbps is imposed on vectored lines.

All simulation results calculate achievable rates for vectored lines versus distance when non-vectored lines are capped to operate at a given data rate, and DSM is used to balance transmit spectra.

VI.2   The mixed vectored/non-vectored case

For the mixed vectored/non-vectored case we plotted results as a bar chart for three loop lengths (see Figure VI- 1) and as a rate reach plot for distances up to 1 km in 100 meter increment (see Figure VI-2 to Figure VI-5).

The bar chart Figure VI-1 shows the performance of various simulated scenarios (from top to bottom):
- A dense (full binder) VDSL2 only deployment
- The ideal performance achieved when all VDSL2 lines are vectored in a single vector group, showing gains of 2x or more compared to the non-vectored case.
- The more realistic performance achievable with vectoring when not all crosstalk is eliminated, still showing at least 2x gains with respect to the non-vectored case.
- When 12 non-vectored lines coexist in the same binder with 12 vectored lines and no remedial action is undertaken, vectored lines suffer a sharp loss in performance and the achievable data rates are very close to the VDSL2-only case (top case)
- When DSM is used to mitigate the effects of uncancelled crosstalk, vectored lines recover most of their gains, while non-vectored lines operate between 40-50 Mbps.
- When DSM is used to mitigate the effects of uncancelled crosstalk and vectored lines have a maximum bit rate of 100 Mbps, DSM allows vectored lines to continue to operate at 100 Mbps, while non-vectored lines operate between 45-95 Mbps.
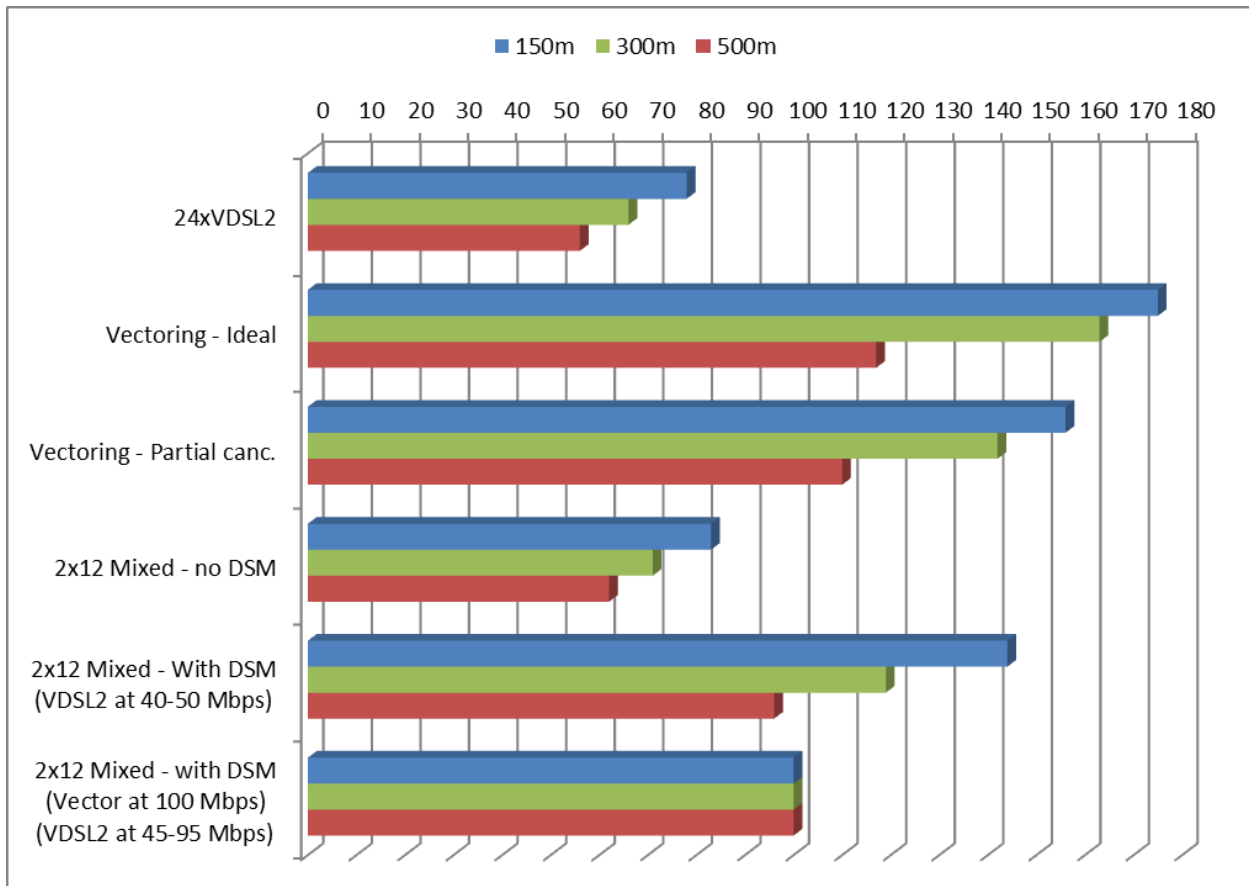
Figure VI - 1 - 1% worst case downstream vectored data rates for the mixed scenario.

VI.2.1   Uncapped vectored data rates: rate-reach plot

In this section the case is considered where vectored lines are allowed to achieve their maximum data rate while in the presence of non-vectored lines that are capped at a data rate that varies with

loop length as shown above the x-axis of the plots. The non-vector caps here were chosen to achieve appealing tradeoffs, but this could be varied somewhat to favor either the vectored or non-vectored lines.
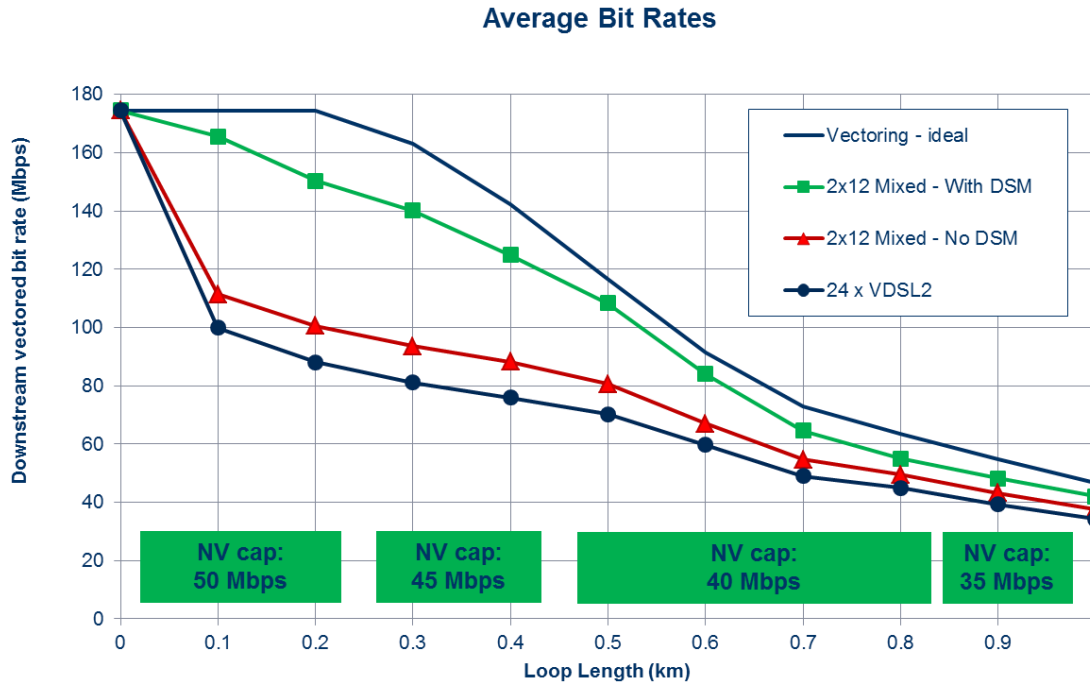
Figure VI-2 shows 1% worst case data rates, where "2x12 Mixed" is a mixture of 12 vectored and 12 non-vectored lines. It can be noted that:

- Partial vectoring entails a 9-13% data rate loss compared to ideal vectoring at short distances, while this loss disappears above 500 meters.
- In a mixed case, almost all vectoring gains are destroyed when no remedial action is undertaken. In fact, the performance of vectored lines is very close to the performance that would be achieved by 24 non-vectored VDSL2 lines.
- When DSM is used to mitigate the effect of uncancelled crosstalk, vectored lines suffer only an additional 15% data rate degradation from realistically achievable rates while non-vectored lines are capped at speeds ranging between 35-50 Mbps.

Figure VI-3 shows the same scenario of Figure VI-2, but average data rates are plotted instead of the 1% worst case. DSM allows vectored lines to achieve (on average) a ~9% data rate degradation from ideal vectoring while non-vectored lines are capped at speeds ranging between 35-50 Mbps.
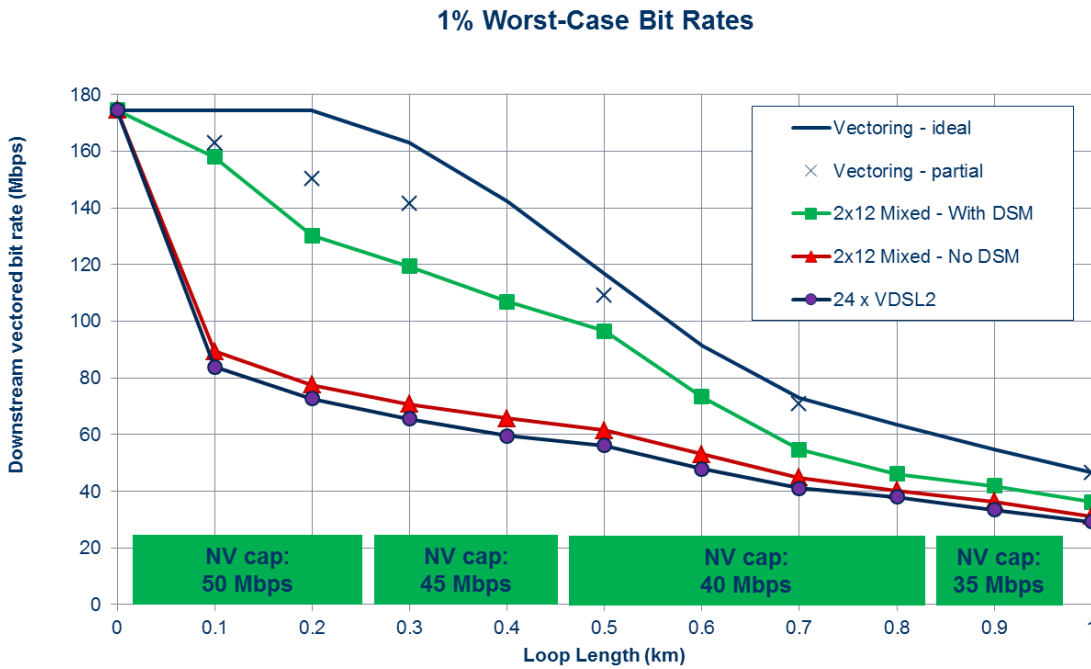
Figure VI - 2 - Uncapped downstream vectored data rates, 1% worst case. The data rate cap imposed on the non-vectored (NV) lines is shown at the bottom of the plot.

Average Bit Rates

Figure VI - 3 - Uncapped downstream vectored data rates, average data rates. The data rate cap imposed on the non-vectored (NV) lines is shown at the bottom of the plot.

VI.2.2   Vectored data rates capped at 100 Mbps – rate-reach plot

In this section the case is considered where vectored lines can at most achieve a maximum data rate of 100 Mbps. If this rate is not achievable, vectored lines are set to maximize their speed while in the presence of non-vectored lines that are capped at a certain data rate.

Figure VI-4 shows 1% worst-case data rates, where "2x12 Mixed" is a mixture of 12 vectored and 12 non-vectored lines. It can be noted that:

- In a mixed case, DSM allows vectored lines to maintain their 100 Mbps data rate target up to around 500 meters even in the presence of non-vectored lines whose data rate ranges between 45-95 Mbps.
- Above 500 meters, DSM effectiveness diminishes and vectored lines suffer a 20-30% degradation compared to what is achievable with ideal vectoring in the presence of non-vectored lines whose data rate ranges between 30-40 Mbps.

Figure VI-5 shows the same scenario of Figure VI-4, but average data rates are plotted instead of the 1% worst case. DSM allows vectored lines to achieve (on average) only a ~14% or less data rate degradation from ideal vectoring.
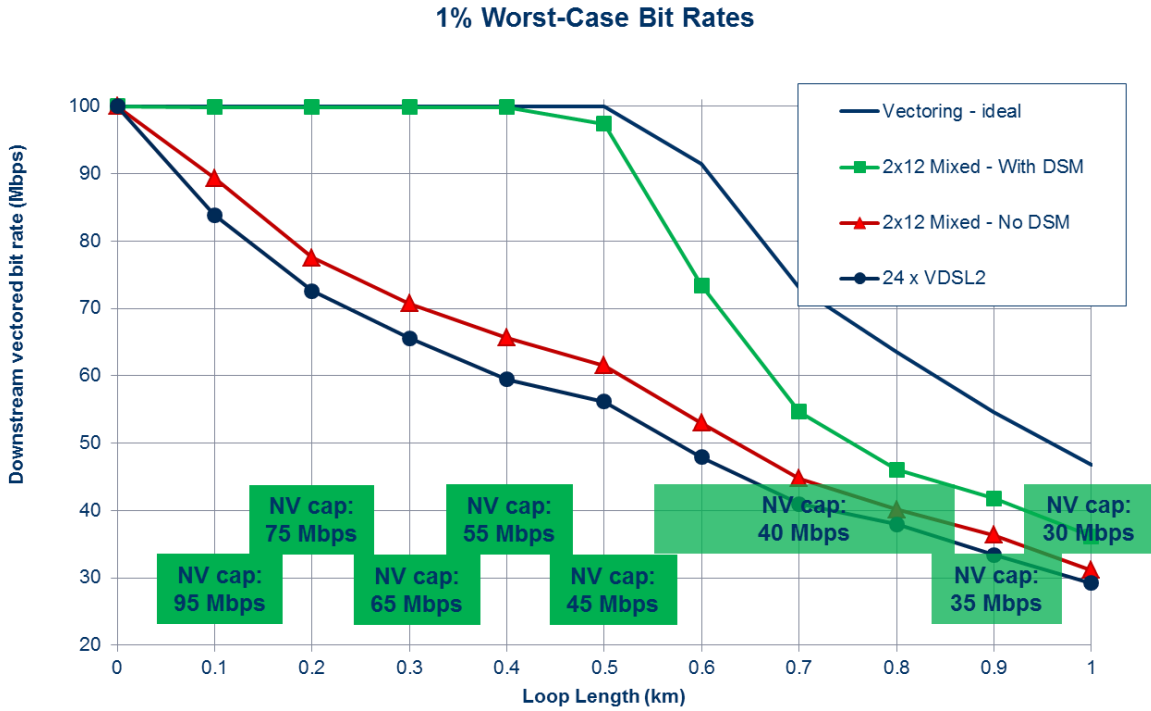
Figure VI - 4 - Downstream vectored data rates capped at 100 Mbps, 1% worst case. The data rate cap imposed on the non-vectored (NV) lines is shown at the bottom of the plot.
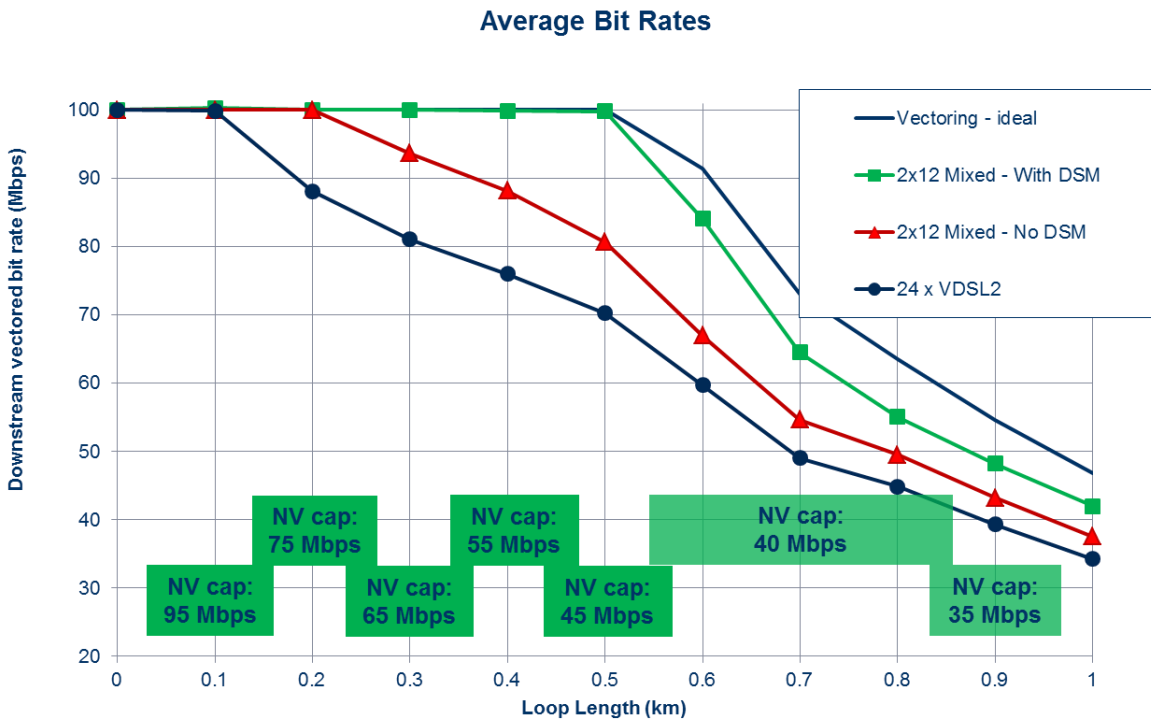

Figure VI - 5 - Downstream vectored data rates capped at 100 Mbps, average data rates. The data rate cap imposed on the non-vectored (NV) lines is shown at the bottom of the plot.

VI.3   The two vector group case

For the case of two vector groups, results are plotted as a bar chart for two loop lengths, 150 and 500 meters. The bar chart Figure VI- 6 shows the performance of various simulated scenarios (from top to bottom):

- A dense (full binder) VDSL2 only deployment
- The ideal performance achieved when all VDSL2 lines are vectored in a single vector group, showing gain of 2x or more compared to the non-vectored case.
- The more realistic performance achievable with vectoring when not all crosstalk is eliminated, still showing at least 2x gains with respect to the non-vectored case.
- When two vectored groups are present in the same binder and no remedial action is undertaken, vectored lines of the two groups suffer a sharp loss in performance and the achievable data rates are very close to the VDSL2-only case (top case)
- When DSM is used to mitigate the effects of uncancelled crosstalk, vectored lines recover some of the loss but still show degradation of around 35% compared to realistic vectored data rates. However, the two vector groups operate at speeds that are 30% higher than what is achievable in a dense VDSL2 deployment.
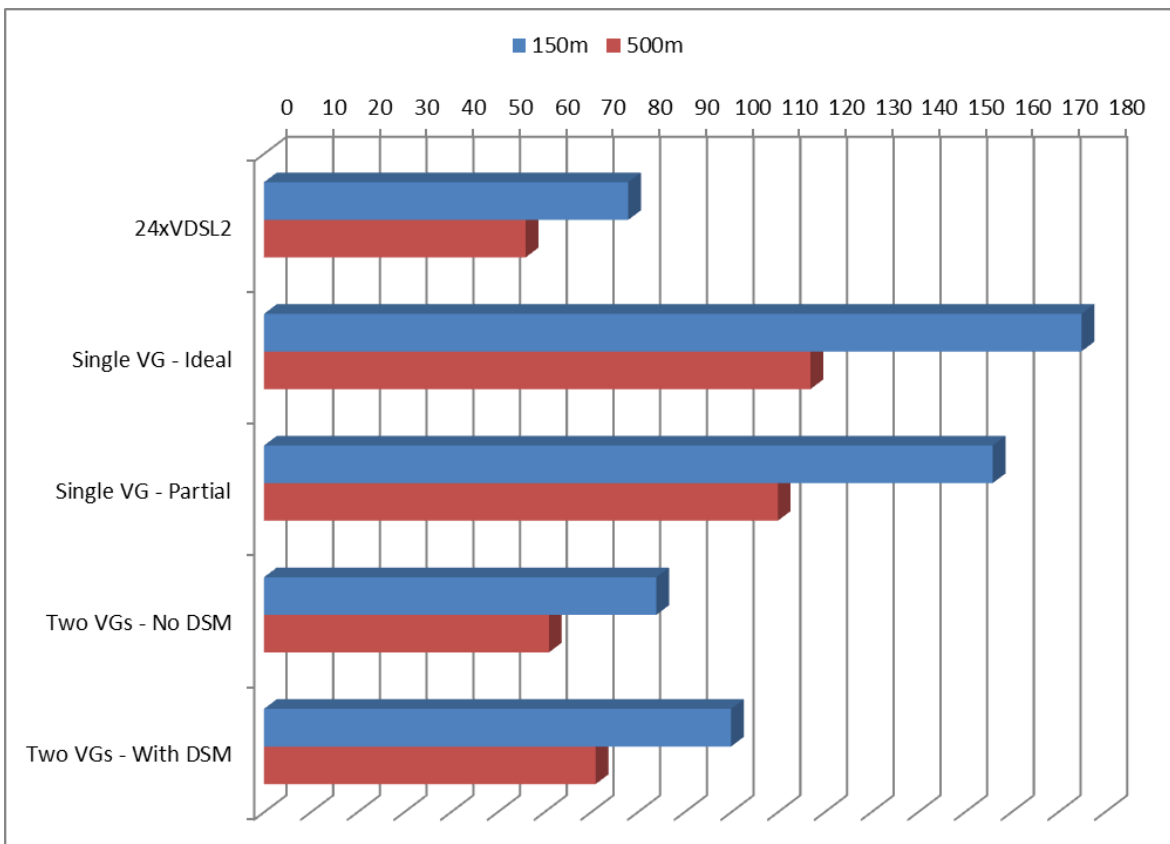

Figure VI - 6 - 1% worst case downstream vectored data rates for two vector groups.

VI.4   Conclusions

DSM provides a framework for preserving most of vectoring gains in mixed vectored/non-vectored scenarios:
  • Vectored lines suffer < 15% data rate degradation from realistically achievable rates, while non-vectored lines are capped at speeds even higher than those delivered today (40-50 Mbps)
  • If vectored lines operate at a 100 Mbps target, then vectored lines suffer no degradation at all up to 500 m while non-vectored lines can be capped at even higher speeds (40-95 Mbps)

Coexistence is more challenging for the two vector group case, and the results in Apprendix I are confirmed: only one of the two groups would be able to operate at vectoring speeds while the other group would have to be capped at VDSL2-like speeds.

If the data rate of both vector groups is maximized, then in scenarios with two vectored groups:
  • Vectoring gains suffer ~35% degradation from realistically achievable rates, while still maintaining ~30% or more improvement over dense VDSL2 deployments.
  • If all vectored lines are capped at 100 Mbps, then vectored lines suffer minimal degradation up to 200m and around 30% degradation on longer lines, while still maintaining ~30% improvement over dense VDSL2.

DSM cannot completely eliminate uncancelled crosstalk but it can mitigate its impact by implementing trade-offs between the speeds supported by the vectored and non-vectored lines. In a competitive unbundled environment, the acceptance of mitigation techniques depends on the definition of fairness, and DSM provides a flexible framework for enabling coexistence that can match various definitions of fairness. In the scenario of a single provider upgrading gradually legacy VDSL2 to vectored VDSL2, DSM similarly provides a flexible framework to achieve service level targets.

<div style="border:2px solid black; background:#d9d9d9; text-align:center; padding:40px;">
End of Broadband Forum Technical Report TR-320
</div>