MANAGING UNVECTORED LINES IN A VECTORED GROUP

STRATEGIC WHITE PAPER

Vectoring technology supports the removal of far-end crosstalk (FEXT) from coordinated lines. However, challenges arise when vectored lines are mixed with legacy VDSL2 lines in the same binder. This paper examines a dynamic line management (DLM) technique supported by Alcatel-Lucent's Motive Network Analyzer – Copper product. This technique is based on scalar parameters that control the maximum transmit power spectral density (PSD), signal-to-noise-ratio margin (SNRM) and data rate on the legacy lines.

In addition, the paper shows that the DLM technique can allow operators to recover some of the performance lost due to the presence of legacy lines when the data rate on these legacy lines is significantly capped. This is a practical solution in cases where all lines are served by a single operator. When more than one operator serves lines in the same access network, the DLM technique does not support fair competition between the operators without also conceding some of the performance improvements gained through vectoring. Virtual unbundling may be preferred in this case, as it promotes a competitive landscape.

TABLE OF CONTENTS

1. INTRODUCTION

The twisted copper wiring telephony network is widely used to deliver broadband access with the VDSL2 technology. However, the performance of a VDSL2 access network is limited by electromagnetic coupling between wire pairs within a binder. Recently, digital signal processing means have been standardized in the International Telecommunication Union (ITU) G.vector recommendation. As described in reference [1], the G.vector standard will help reduce the impact of far-end crosstalk (FEXT) through a multi-line vectoring operation at the access node.

The performance of vectored lines is similar to that of a line experiencing no FEXT, provided that all mutually cross-talking lines are connected to the same vectoring group and that all customer premises equipment (CPE) connected to these lines supports the required vectoring functionalities [2]. The challenges addressed in this paper deal with scenarios in which these conditions are not fulfilled.

Much of the CPE in use today is firmware upgradeable to full G.vector or a vector-friendly mode that can cancel crosstalk into other lines, but that does not remove crosstalk on the vector-friendly line in question. Equipment that cannot be upgraded to either mode is called legacy VDSL2 CPE.

Several techniques can be used to estimate the crosstalk channels in the presence of legacy lines. One estimation method is based on legacy signal-to-noise (SNR) reporting capabilities and allows vectoring across legacy lines. See reference [3] for details. Another method exploits the robustness of sync symbols to add a low-power pilot sequence to downstream legacy lines. This method is described in reference [4].

These estimation techniques require that all legacy lines be controlled within the same vectoring group. This paper distinguishes between two scenarios in which this condition is not achieved. In the first scenario, the vectored and legacy lines are managed by the same operator, but physically connected to different access node equipment. In the second scenario, two operators each control a fraction of the lines in a binder through a process called sub-loop unbundling (SLU). The second operator may deploy legacy lines, or vectored lines that are not part of a single vectoring group. The paper considers dynamic spectrum management (DSM) techniques that (co-)optimize the transmit power spectral density (PSD) of all lines under imposed constraints. For an overview of these performance-optimizing DSM techniques, see reference [5].

References [6] and [7] provide results for the single-operator scenario and discuss an extrapolation of these results to the multi-operator scenario. In the single-operator scenario, DSM significantly improves the performance of the vectored lines at the expense of data rate degradation on the legacy lines. This is considered to be acceptable since the single operator may use the G.vector service for high-end customers and the legacy service for lower-end customers.

Simulations have been performed in which each legacy loop has an equal loop length. The results of these simulations overestimate the vectoring gain that can be attained in real-world deployments. Longer legacy loops are less able to perform PSD reduction, and will have higher crosstalk than is shown in equal loop length simulations. Section 2 of this paper provides performance results for a more realistic loop length distribution. It examines a dynamic line management (DLM) technique supported by Alcatel-Lucent's Motive Network Analyzer — Copper product. This technique is based on scalar parameters that control the maximum transmit PSD, SNRM and data rate on the legacy lines.

Reference [7] uses the results for the single-operator scenario to draw conclusions about the multi-operator scenario. This approach is only appropriate in cases where the second operator does not deploy any vectored lines and is willing to significantly reduce the data rates offered on its legacy lines. Such a situation would strongly favor the first operator and distort market competition. As a result, it is more relevant to consider multi-operator scenarios in which all operators are treated equally and fairly. These scenarios are discussed in section 3 of this paper.

The European Union (EU) has set forward a Digital Agenda that aims to make Internet speeds of 30 Mbps downstream available to all its citizens, and 100 Mbps to 50% of households, by 2020. This paper uses Digital Agenda targets as prime performance benchmarks and analyzes the sensitivity of the various results to these targets.

2. SINGLE-OPERATOR SCENARIO

2.1 Simulation environment

In the simulation environment, $K = 48$ lines, of which $K_V = 24$ vectored VDSL2 lines and $K_L = 24$ legacy VDSL2 lines. For each line, the loop length is drawn from a Gamma distribution with a mean of 604 m and a standard deviation of 426 m (Figure 1). This length distribution has been derived from analysis of actual distribution networks across the world.

Figure 1. The loop lengths are Gamma distributed with a mean of 604 m.

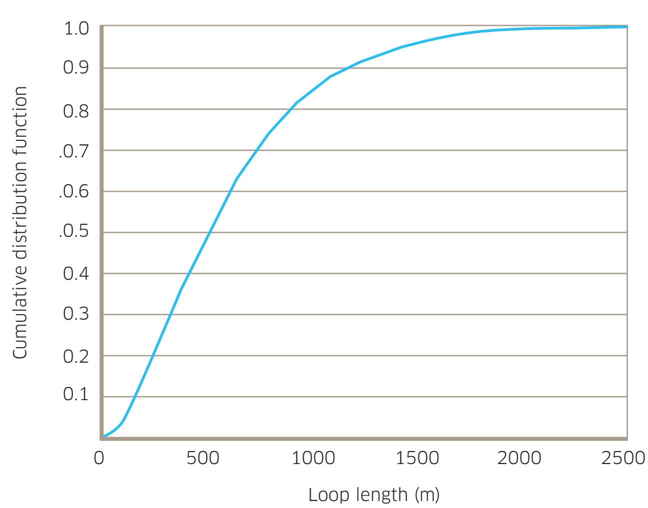

The direct channel response $|H(n)|^2$ at frequency bin n is modeled as 26 awg polyethylene insulated cable, as indicated in reference [8]. The crosstalk interference $I(n)$ is beta distributed according to the ANSI Multiple-Input Multiple-Output (MIMO) model described in references [9] and [10]. Here, each of the K lines is assigned a random index i from 1 to 100. Crosstalk coupling coefficients from and into the line with index i are taken from the ith column and row of the 100x100 reference matrix provided in reference [10]. The 998ADE17 band plan is taken with the M2 PSD mask. With this band plan, 71% of carriers are allocated to downstream transmission. PSD is augmented at the lower frequencies to increase performance on the longer loops. The equipment noise floor is taken at -130 dBm/Hz. The gap to capacity is Γ = 10.75 dB, which corresponds to a coding gain of 5 dB and a signal-to-noise ratio margin (SNRM) of 6 dB. The transmission efficiency of η = 78.5% accounts for cyclic extension, frame structure and Reed-Solomon coding overhead. The simulation parameters are summarized in Table 1.

Table 1. Simulation parameters

Parameter	Value
Binder type	26 awg (0.4 mm) polyethylene insulated
Band plan and mask	998ADE17 M2 over POTS
Carrier spacing f_c	4.3125 kHz
Noise floor σ_0^2	-130 dBm/Hz
X_{dB}	Beta distributed
Γ	10.75 dB
η	78.5%
b_{max}	15 bits

The net downstream data rate R is derived as:

$$R(k) = \eta f_c \sum_{n=DS} \min(\log_2(1 + \frac{|H(k,n)|^2 PSD(k,n)}{\Gamma(\sigma_0^2 + \sum_{l \neq k} I_{k,l}(n))}), b_{max})$$

PSD(k,n) is the PSD of user k at carrier n. Because of the statistical nature of the interference matrix $I_{k,l}$, loops of the same length may experience different interference levels. Figure 2 shows the performance of the vectored (green) and legacy lines (red). A data rate spread around the median (full curve) is observed for the legacy and vectored lines. The dash-dot line indicates the performance achieved by 99% of the lines. The dashed line indicates the performance achieved on 1% of the lines. The performance increase gained by perfectly cancelling crosstalk among the vectored lines is a fraction of the increase that could be achieved with full crosstalk cancellation in the absence of legacy lines (blue).

Figure 2. When all lines transmit at full power, vectoring gains are significantly compromised.

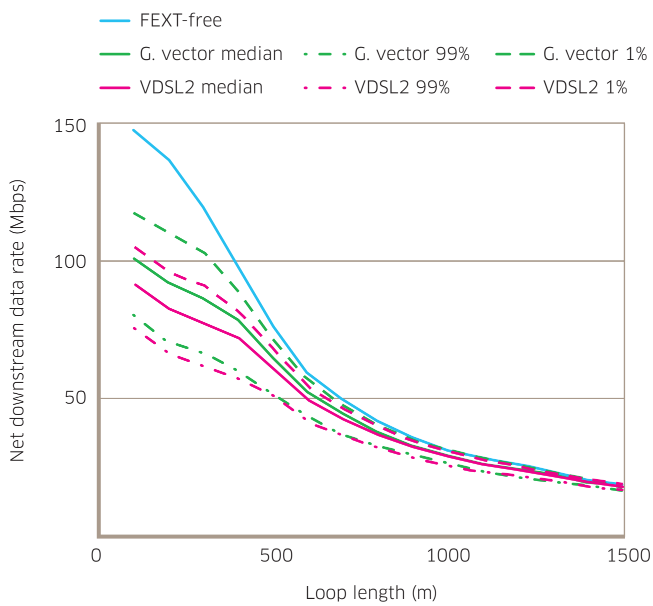

2.2 Restriction of PSD and SNRM

Crosstalk from legacy lines into vectored lines can be reduced by lowering the PSD on the legacy lines. This also reduces the attainable data rate and SNRM on the legacy lines. We can now manage interference between the lines by exploiting two scalar parameters per line *k: maxPSD(k)* and *maxSNRM(k)*. We assume that the data rate on the vectored lines is uncapped. This means that the lines are allowed to synchronize at their attainable data rate.

We can make vectored lines “polite” — and, therefore, less likely to crosstalk with other lines — by restricting their SNRM to *maxSNRM(k)* = 6 dB. This restriction reduces the PSD, primarily on lower frequencies where the bit-loading capacity exceeds the maximum allowed bit-loading of b_{max} = 15 bits. On the legacy lines, the maximum allowed PSD (maxPSD) is reduced so that the net data rate reaches a rate of R_L^{tar} = 30 Mbps, which is the target rate set by the EU Digital Agenda. A PSD reduction on one line will increase attainable data rate on other lines, so the procedure is repeated until the *maxPSD* and *maxSNRM* parameters have converged. Figure 3 outlines the algorithm used for this procedure.

Figure 3. Applying the maxPSD and maxSNRM parameters to restrict PSD and SNRM across legacy and vectored lines.

Algorithm 1 Apply maxPSD and maxSNRM

```
Unitl convergence
  For all legacy lines
    Select maxPSD(k) that minimizes |R(k)-R_L^tar|
    Apply maxSNRM
  End
  For all vectored lines
    Apply maxSNRM
  End
End
```

In practical implementations, the *maxPSD* and *maxSNRM* parameters are derived by a dynamic line management (DLM) system. This system selects optimal parameter values while ensuring stability under line-specific time-varying noise conditions, as described in reference [11]. In a steady noise environment where all modems are continuously on, the *maxPSD* can be controlled indirectly using the *maxSNRM* and rate capping parameters, provided that the modems obey a strict implementation of these parameters. Here, the *maxSNRM* parameter should be optimized per line, for example, by an automated management system.

The *maxPSD* parameter must be configured directly in transient noise conditions — for example, where end users switch their modems on or off based on their broadband usage patterns. If a given line (line A) is initialized while only a few other lines are active, imposing *maxSNRM* will ensure that line A will see a reduction in SNRM when more lines become active. This reduction will allow the line A to maintain the target data rate. Stability control by DLM system will ensure that lines maintain their data rate and stability at all times.

The results from Algorithm 1 are accumulated for 2000 times $K = 48$ random draws of loop length and crosstalk indices. The result is shown in Figure 4, where the legacy lines achieve the $R_L^{tar} = 30$ Mbps target up to 1200 m. The PSD reduction achieved through rate capping significantly benefits the vectored lines: The median performance is close to the FEXT-free performance, and 100 Mbps can be achieved on 99% of the vectored lines up to a loop length of 300 m.

Figure 4. Using maxPSRD and maxSNRM to reduce the legacy line target rate to 30 Mbps significantly increases the attainable data rate for the vectored lines.

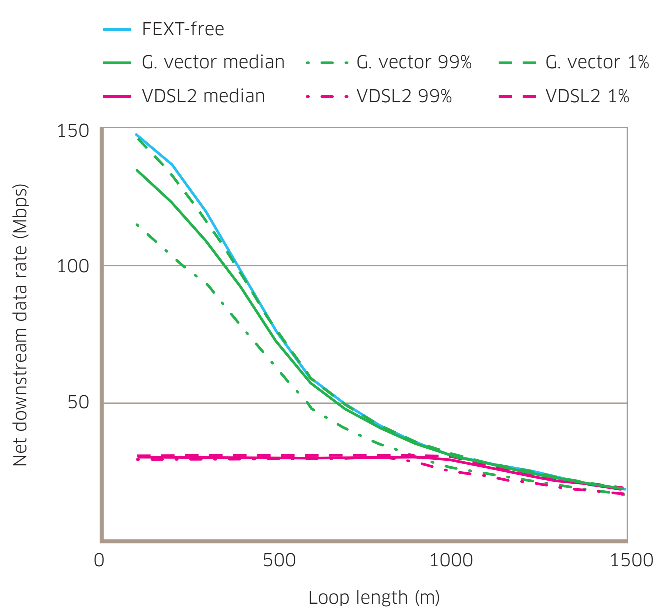

Figure 5 shows the percentiles of the *maxPSD* value obtained after optimization. Some vectored lines transmit at full power, but most vectored and legacy lines apply significant power back-off up to a loop length of 1200 m. For longer loops, attenuation is such that power back-off implies a reduction in attainable data rate. While long legacy lines transmit at full power, the median short vectored lines are able to reduce their maximum PSD by 10 dB. The spread in *maxPSD* values at any given loop length makes it clear that a generic loop length-dependent *maxPSD* configuration is insufficient for maintaining the target data rate.

A line- and binder-specific optimization is an appropriate alternative to this generic configuration. Optimization can be achieved in a network-scalable manner using the automated DLM feature provided by the Motive Network Analyzer. The spread in *maxPSD* values on the vectored lines may cause lines with low *maxPSD* values to experience high crosstalk amplitudes relative to their direct signal amplitude. It has been confirmed that, in a practical implementation, the residue of the crosstalk that remains after cancellation does not notably affect the vectored lines' ability to back off their power as compared to the simulated results shown in Figure 5.

Figure 5. Percentiles of the maxPSD parameter of the bundle containing 24 vectored and 24 legacy lines with rate capping on legacy lines (30 Mbps).

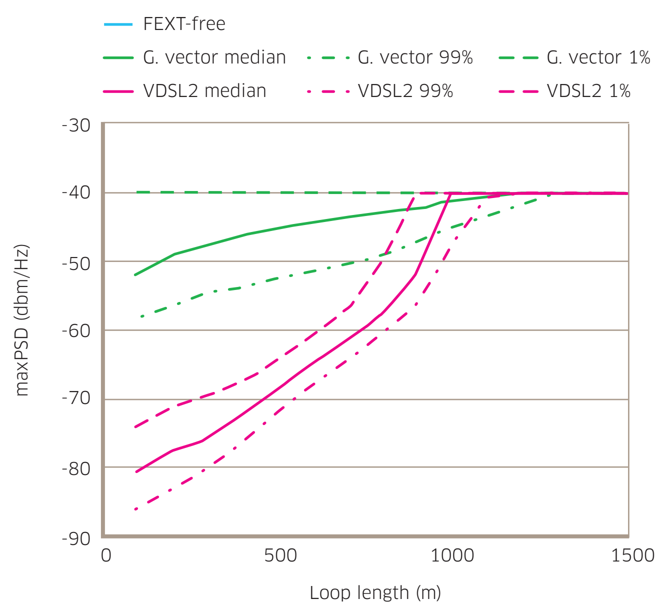

The net aggregate data rate R_V achieved on the vectored lines depends on the target data rate selected for the legacy lines. Figure 6 shows R_V relative to the FEXT-free rate R_{FF} for the 1st, 50th and 99th percentiles of the distribution at 500 m. The green reference curves show the R_V and R_{FF} references in cases where the data rate on the legacy lines is not capped, but where *maxSNRM* is applied to tones that achieve b_{max}. The figure shows that near-R_{FF} rates can be guaranteed on the vectored lines if the data rates on the legacy lines are capped at 10 Mbps or less. The attainable rates on the vectored lines drop quickly if the data rates on the legacy lines are allowed reach 20 Mbps or more.

Figure 6. Relative downstream data rate for the vectored line as a function of the target data rate for legacy lines.

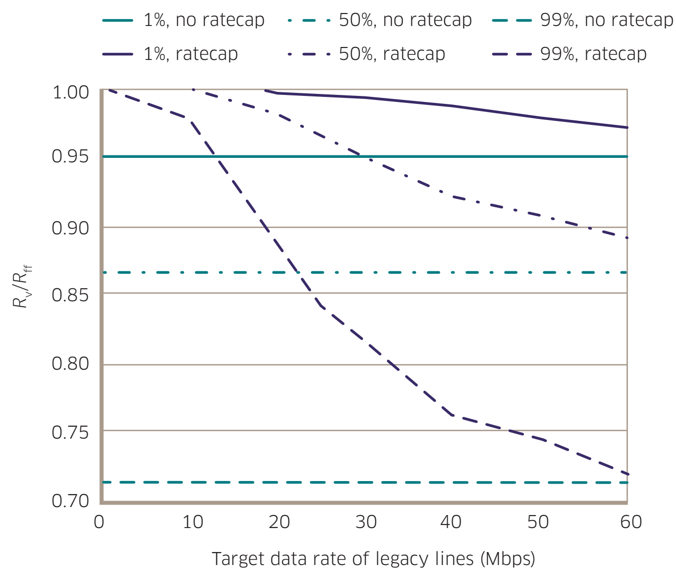

In a last step for the single-operator scenario, the data rate on the vectored lines is capped at 100 Mbps. This permits further reduction of PSD on the legacy lines. The resulting decrease in crosstalk coupling does not significantly increase the reach of the 100 Mbps service on the vectored lines, as shown in Figure 7. In alignment with reference [6], we confirm that negligible further improvement is obtained through the application of iterative spectrum balancing (ISB), a dynamic spectrum management level 2 technique

Figure 7. Rate capping does not significantly increase the reach of vectored lines (100 Mbps).

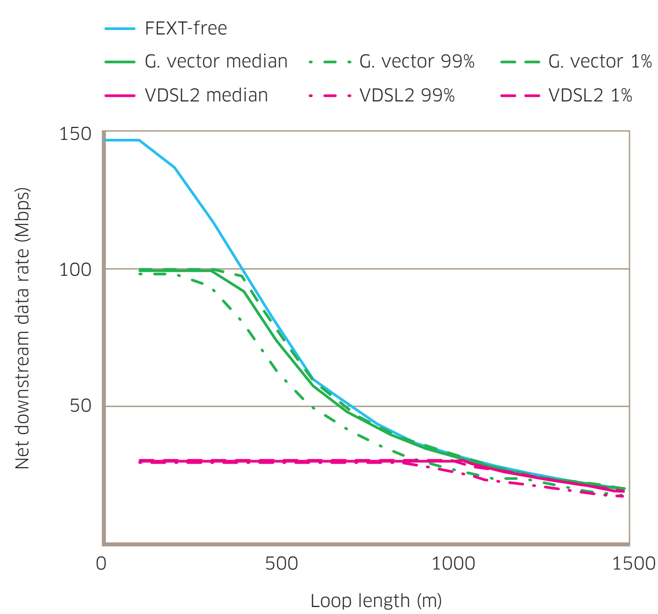

3. MULTI-OPERATOR SCENARIO

When several operators deploy access node equipment at the same location through a local-loop unbundling (LLU) process, their digital subscriber lines will experience mutual crosstalk. This section describes two distinct multi-operator scenarios, one that combines vectoring technology with legacy VDSL2 technology and one that involves vectoring technology only.

In the first scenario, operator A deploys vectoring technology and operator B deploys legacy VDSL2 technology. The data rate performance achieved through this scenario corresponds to that achieved through the single-operator scenario described in section 2. Operator A can benefit significantly from vectoring gains only if operator B is willing to sacrifice data rate. This is an unfair competitive landscape, since operator A can offer much higher data rates than operator B.

In the second scenario, two operators deploy vectoring technology. If both operators allow their lines to transmit at full power, the rate–reach performance for both operators is that indicated by the green curves in Figure 2. The red curves in Figure 2 indicate the performance increase — 5 to 10% for loops of 0 to 500 m — provided by crosstalk cancellation. For these loop lengths, the performance degradation compared to full coordination is still 50 to 100% for the 99th percentile.

We can evaluate the ability of Algorithm 1 (shown in Figure 3) to increase performance for the 99th percentile. We do this by applying Algorithm 1 to two vectored groups that we wish to cap at 75 Mbps. No legacy lines are present in this scenario. Algorithm 1 is adapted to apply a *maxPSD* (line 3 of Algorithm 1) to all vectored lines. In an alternative procedure, a frequency-selective power back-off has been achieved through the near-optimal iterative spectrum balancing (ISB) algorithm described in references [5] and [12]. Figure 8 shows the rate-reach obtained through ISB. These results can also be obtained with Algorithm 1, with tolerances below 1 Mbps.

Figure 8. Rate–reach of a bundle containing two uncoordinated groups of 24 vectored lines with and without rate capping at 75 Mbps.

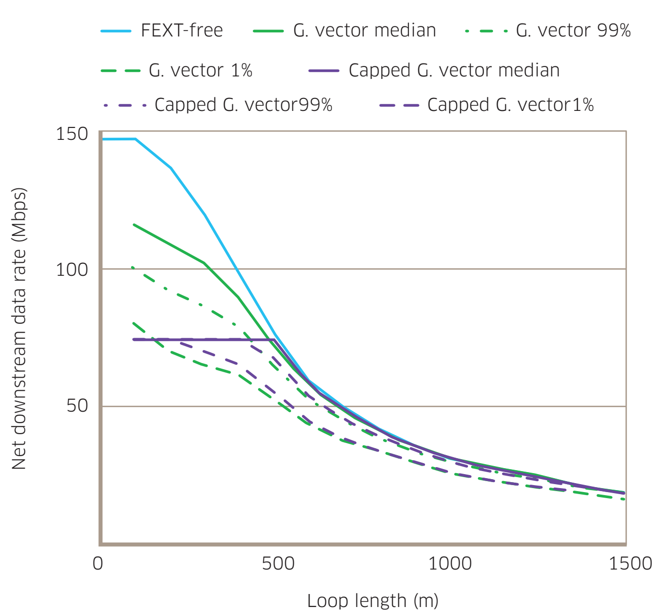

In a scenario with two vectored groups, the ISB approach achieves the same median rates as the *maxPSD* approach used by Algorithm 1. The difference between the median rate achieved through *maxPSD* and ISB is below 0.2 Mbps at every rate–reach point. In addition, the ISB approach does not guarantee more reach than the *maxPSD* approach. At the 99th percentile, the difference in the median rate achieved through *maxPSD* and ISB remains below 1 Mbps for every rate–reach point.

Neither ISB nor Algorithm 1 increases the reach for the 75 Mbps service. However, both increase the median data rate by up to 2 Mbps for loops that do not attain the 75 Mbps target (Figure 8). The data rate increase is most pronounced for the 99th percentile. Even there, however, the data rate gain remains below 5 Mbps.

Example spectra obtained through ISB (Figure 9) show that, although the spectra are not flat, there is a higher back-off at lower frequency. For instance, compared to the mask, the black curve (100 m loop) shows a 25 dB back-off in downstream 1 and a 15 dB back-off in downstream 3. In general, shorter loops can perform a higher back-off than longer loops. However, the two 100 m loops shown (black and green curves) reveal large variations between loops of the same length. These variations highlight the need for per-line optimization. Even with ISB, however, the difference between attainable rates (black curves in Figure 8) and FEXT-free rates (blue curve in Figure 8) shows that much of the vectoring gain is sacrificed in cases where lines are located in different vectoring groups.

Figure 9. Example transmit PSDs obtained through ISB for the case with two vectored groups with rate capping at 75 Mbps.

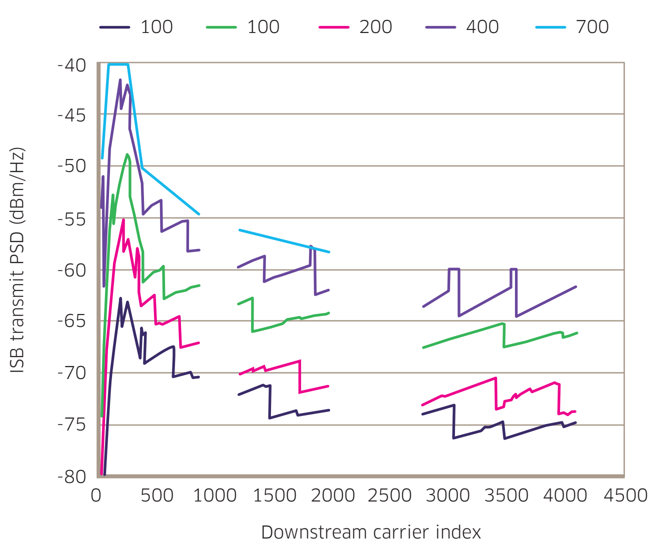

The same trend is seen in Figure 3 of reference [13]. The performance gap to full coordination remains at 40% if a zero-forcing precoder is combined with optimal spectrum balancing (OSB), and at 30% if OSB is combined with an optimal linear precoder. The latter combination demands accurate knowledge of the crosstalk channel state information between the two groups of lines. It is impractical with state-of-the-art technology.

A significant gap to crosstalk-free performance remains in cases where vectoring technology is applied in an LLU context. Some alternatives to LLU can retain the full benefits of vectoring. These include bit-stream unbundling, wholesale access and virtual unbundled local access (VULA). All of these alternatives are variants of the concept in which one operator owns the access equipment, and multiple service providers offer services over this access equipment. In each case, the data of the multiple service providers is (de)-multiplexed at different locations in the network (access, edge or core). Service providers can have lesser or greater degrees of freedom to manage their lines.

4. CONCLUSIONS

When legacy and vectored lines are combined in the same cable binder, their mutual crosstalk reduces the data rate gain provided by vectoring technology. In an environment featuring lines served by multiple operators, virtual unbundling is the preferred choice, as it promotes a competitive landscape. The spectrum management technique is less suited to an unbundled environment, as it can give one operator an advantage over the others or reduce the vectoring gains to below 10%.

In cases where all lines are served by the same operator, a DLM technique based on scalar parameters can create an effective trade-off between the crosstalk cancellation gain on vectored lines and the data rate on legacy lines. Alcatel-Lucent's Motive Network Analyzer — Copper product supports the DLM technique as part of a rich suite of diagnostics. With the DLM capabilities provided by the Motive Network Analyzer, operators can manage legacy and vectored lines in real time by observing crosstalk levels and fine-tuning Tx power levels to ensure optimal performance across the binder.

5. REFERENCES

[1] V. Oksman et al., "The ITU-T's new G.vector standard proliferates 100 Mb/s DSL," IEEE Communications Magazine, October 2010.

[2] M. Guenach, J. Maes, M. Timmers, O. Lamparter, J.-C. Bischoff and M. Peeters, "Vectoring in DSL systems: Practices and Challenges," in *Proc. IEEE Globecom 2011*, Houston, TX, December 2011.

[3] M. Guenach, J. Louveaux, L. Vandendorpe, P. Whiting, J. Maes and M. Peeters, "On the signal to noise ratio assisted crosstalk channel estimation in downstream DSL systems," in *IEEE Trans. Signal Proc.*, 2010.

[4] J. Maes, C. Nuzman, A. v. Wijngaarden and D. V. Bruyssel, "Pilot-based Crosstalk Channel Estimation for Vector-enabled VDSL Systems," in *Proc. IEEE Conference on Information Sciences and Systems (CISS2010)*, Princeton, NJ, 2010.

[5] S. Huberman, C. Leung and T. Le-Ngoc, "Dynamic Spectrum Management (DSM) Algorithms for Multi-User xDSL," *IEEE Commun. Surveys Tuts.*, Vol. 14, No. 1, pp. 109-130, February 2012.

[6] M. Mohseni, G. Ginis and J. Cioffi, "Dynamic Spectrum Management for Mixtures of Vectored and Non-vectored DSL Systems," in *Proceedings of 44th Conference on Information Sciences and Systems*, March 2010.

[7] K. Kerpez, J. Cioffi, S. Galli, G. Ginis, M. Goldburg, M. Mohseni and A. Chowdhery, "Compatibility of Vectored and Non-Vectored VDSL2," in *Proceedings of 46th Conference on Information Sciences and Systems*, March 2012.

[8] American National Standard Institute (ANSI) T1.413, "*Network and Customer Installation Interfaces*," 1998.

[9] J. Maes, M. Guenach and M. Peeters, "Statistical Channel Model for Gain Quantification of DSL Crosstalk Mitigation Techniques," in *Proc. IEEE Int'l Conf. Commun. (ICC09)*, Dresden, 2009.

[10] ATIS 0600024, "Multiple-Input Multiple-Output Crosstalk Channel Model," 2009.

[11] J. Maes, M. Guenach and M. Peeters, "Autonomous Dynamic Optimization for Digital Subscriber Line Networks," *Bell Labs Technical Journal, Self-Organizing and Self-Optimizing Networks issue*, vol. Vol. 15, No. 3, pp. 119-130, December 2010.

[12] R.Cendrillon and M. Moonen, "Iterative spectrum balancing for digital subscriber lines," in Proc. *IEEE Int. Conf. Commun.*, May 2005.

[13] A. Forouzan, M. Moonen, J. Maes and M. Guenach, "Joint Level 2 and 3 Dynamic Spectrum Management for Downstream DSL," accepted for IEEE Transactions on Communications, 2012.

6. ABBREVIATIONS

ANSI	American National Standards Institute
CPE	customer premises equipment
DLM	dynamic line management
DSM	dynamic spectrum management
EU	European Union
FEXT	far-end crosstalk
ISB	iterative spectrum balancing
ITU	International Telecommunication Union
LLU	local-loop unbundling
MIMO	multiple-input multiple-output
POTS	plain old telephone service
PSD	power spectral density
SLU	sub-loop unbundling
SNRM	signal-to-noise ratio margin
VDSL2	very-high-speed digital subscriber line 2
VULA	virtual unbundled local access

Alcatel·Lucent
Enterprise