Methods for Supporting Vectoring when Multiple Service Providers Share the Cabinet Area

Andrea Colmegna[(1)], Stefano Galli[(2)], Marc Goldburg[(2)]

[(1)] FASTWEB S.p.A., Milan, Italy – andrea.colmegna@fastweb.it
[(2)] ASSIA, Inc., Redwood City, CA, USA – {sgalli, mgoldburg}@assia-inc.com

***Abstract* — Vectoring cancels crosstalk between multiple VDSL2 lines and can greatly improve rate reach performance of VDSL2, particularly on short loops. Vectoring is only performed on lines within a vector group so that vectored lines may experience un-cancelled crosstalk from nearby lines that are either non-vectored or in a separate vector group (mixed scenario). If nothing is done to mitigate the effects of un-cancelled crosstalk in mixed scenarios, then vectored lines may suffer performance degradation which in some cases may be substantial. This argument has been used to advocate for market competition restrictions that would only allow the incumbent service provider (SP) to manage all vectored lines in a single DSLAM in order to save the full benefits of vectoring. In this paper, we address the issue of compatibility between lines within a vectored group and outside of it, pointing out that there are various solutions that allow reaping the full benefits of vectoring even when multiple SPs are allowed to share the lines in a cabinet, which is the typical case where Sub-Loop Unbundling (SLU) is permitted by the regulator. We conclude that there are no technical reasons for arguing in favor of imposing restrictions on SLU, a very harsh limitation to competition that would severely damage competitive SPs and penalize subscribers.**

I. Introduction

Very-high rate Digital Subscriber Line (VDSL2) technology uses high frequencies and short loop lengths to transmit at speeds up to a few hundred Mbps [1]. VDSL2 uses frequency-division duplexing, upstream and downstream, to avoid near-end crosstalk. However, VDSL2 can still be limited by far-end crosstalk (FEXT) which causes VDSL data rates to drop in dense deployments.

Vectoring, as defined by the ITU-T G.993.5 standard, can greatly improve the performance of VDSL2 [2], [4]. Vectoring removes the FEXT created within a vectored group (*self*-FEXT) by performing precoding at the transmitter (downstream) and crosstalk cancellation at the receiver (upstream). Downstream vectoring is able to cancel only the crosstalk within a given vectored group of lines, but not the crosstalk generated by lines outside the vector group.

Copper is today the least expensive and most direct way of providing broadband services. Vectoring is the next step along the DSL roadmap, providing 100+ Mbps speeds in dense deployments and has allowed operators to reposition DSL as a technology for providing next generation broadband services at competitive costs. Reaping the full benefits of vectoring has become a strategic necessity for SPs that face increasing customer demand for higher data rates and also strong competition from cable companies.

In some countries, the incumbent only is allowed to access the very last mile of copper to deliver broadband services through the deployment of technologies such as VDSL2 or G.fast at a remote street cabinet. Adoption of vectoring when all the lines are controlled by a single operator is not straightforward as commonly believed. In fact, in addition to the DSL lines belonging to a specific vectored group, there may also be additional non-vectored lines. This is a practical scenario that is likely to arise in early phases of vectored deployment as the introduction of vectoring can only be gradual. Since the entire DSL plant cannot be replaced with vectored DSL overnight, operators must have tools to manage the simultaneous presence of vectored and non-vectored lines and also multiple vectoring groups as existing DSL equipment is replaced with vectored DSL. Thus, the belief that having a single operator controlling the entire copper infrastructure eliminates alien disturbers [5], [6] is erroneous since the presence of aliens is independent of whether SLU is allowed or not and will occur even when a single operator controls all the lines.

In other countries, multiple SPs are allowed to access the very last mile of copper in order to foster competition. This competition puts lines under the control of multiple operators, which naturally leads to a mixed scenario where lines in a vector group share a cable with additional non-vectored lines or additional lines belonging to other (disjoint) vectored groups.

Although vectoring is able to cancel self-FEXT, crosstalk from non-vectored lines or from lines belonging to other vectored groups within the same cable or binder (*alien*-FEXT) may not always be removed within a vectored group (see Sect. III for more details). Since this alien crosstalk may cause performance degradation to vectoring, the important question of whether it is possible to reap the benefits of vectoring in the presence of alien crosstalk naturally arises [5]-[8].

As is it not always possible to cancel crosstalk from disturbers outside of a vectored group, naive speculation has often led to the conclusion that vectoring benefits would be lost if multiple SPs were allowed to share the lines in a cabinet [5], [6]. This argument is today being used to argue in favor of market competition restrictions and a reversal of unbundling in favor of so-called "bitstream unbundling," where a single carrier controls the entire infrastructure. Thus, a better understanding of this technical issue is timely and important in order to avoid the imposition of unnecessary restrictions to SLU that would limit competition and, as a consequence, penalize the end user.

The literature on the topic of compatibility between lines within a vectored group and outside of it is still scarce and only a dozen scholarly papers are available [8]. Furthermore, a unitary framework for assessing the impact of alien crosstalk on vectoring is still lacking. Nevertheless, as recently reported [8], all published results actually confirm that there are indeed successful mitigation techniques that allow coexistence between vectored lines and alien disturbers, vectored or not.

Alien crosstalk is most detrimental to vectoring when the lines outside the vectored group are completely unmanaged. However, when non-vectored lines and vectored groups are properly managed, the actual impact of alien-FEXT on the vectored lines is limited and coexistence is indeed possible [7], [8]. For example, the authors of [8], conclude that vectored lines deployed along-side non-vectored lines can achieve excellent performances and retain most of the benefits of vectoring when an appropriate use of Dynamic Spectrum Management (DSM) [9], [10] is made.

In this paper, we continue the work started in [8] providing additional details on the benefits of DSM as well as discussing the advantages ensured by increasing the vector group size from Board Level Vectoring (BLV, vectoring of all lines in a line card) to System Level Vectoring (SLV, vectoring across line cards) which are already a technical reality. Furthermore, we will look at the benefits of performing vectoring across multiple DSLAMs (cross-DSLAM vectoring), i.e. when the vectored group size is increased to include lines terminating on different DSLAMs. In all cases, we demonstrate that success of vectoring in an SLU environment is technically feasible.

II. Overview of the Advantages of Vectoring

VDSL2 is a next generation access technology [1] that allows SPs to support very high speeds and deliver 100 Mbps to all users within about 500 meter of radius for a lightly loaded cable (see Figure 1 which assumes a 17 MHz band plan). Higher data rates are possible with the wider 30 MHz band plan. In the case of dense deployments, i.e., when a high fraction of the lines in a cable is used for VDSL2, unmanaged crosstalk becomes the primary source of performance degradation and the median radius for 100 Mbps VDSL2 can be reduced 150 meters. Furthermore, as Figure 1 confirms, the data rate range that VDSL2 is able to deliver widens when unmanaged crosstalk becomes dominant.

The recently defined Vectored-VDSL2 standard [2]-[4] defines a technology which is an extension to VDSL2 and that is capable of removing self-FEXT by performing DSM Level 3, i.e. real-time crosstalk management [10]. Removing crosstalk greatly improves the performance of VDSL2 so that SPs can deliver 100 Mbps to all users within about 550 meter of radius regardless of the density of their deployment. Note that DSL provides dedicated bandwidth to each user, unlike the case of other broadband access technologies where the available bandwidth is shared among all users, e.g. coax, wireless, and power line communications (PLC). This higher and dedicated data rate at longer ranges enables SPs to finally meet consumers' increasing hunger for higher speeds and also to keep pace with the ever-increasing speed of wireless and wired home networks.

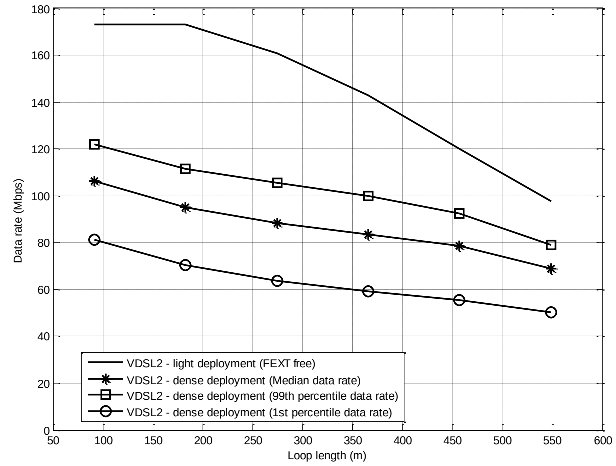

Figure 1 - Rate-reach plot for a group of 48 non-vectored VDSL2 lines for the two cases of light and dense (48% cable fill) deployment. For simulation details, see the Appendix.

Another important advantage of vectoring is that it substantially reduces the wide variability in rate experienced by non-vectored VDSL2 lines. As the effects of crosstalk are not uniform across DSL lines and depend on many factors, non-vectored lines experience a wide variation in performance – especially on shorter lines where crosstalk is stronger. Vectoring removes FEXT and thus vectored lines experience a much smaller performance variation from line to line, allowing SPs to offer higher data rates to a much larger percentage of users. Basically, the same FEXT-free performance is ensured on all lines, although in practice non-crosstalk noises like impulse noise, narrowband interference (NBI), interference due to PLC, etc., are not cancelled by vectoring alone and would have to be independently managed and mitigated. Thus, the actual performance of vectored lines is less than in the ideal FEXT-free case.

Power savings are another important consequence of vectoring. When vectoring is enabled, transmitter power can be reduced on all lines in a group because the receiver experiences a self-FEXT free higher Signal-to-Noise Ratio. This makes vectored lines more "polite" and, thus, less crosstalk is created into neighboring alien lines. As discussed in the following sections, alien-FEXT mitigation techniques applied to alien lines also tend to reduce transmitter power. Overall, general power optimization across all lines of a node can be realized in concert with a deployment of vectoring, regardless of a mixed scenario or not.

Vectoring is different from past DSL technologies and requires adapting management practices, especially those related to line diagnostics. A general review of the best management practices for vectored VDSL2 is given in [11].

III. Mitigation and Cancellation of Alien-Crosstalk

In vectoring, the downstream and upstream transmissions of all the lines in a vectored group are synchronized to a common clock, which allows transmitters to cooperate in the removal of self-FEXT [13]. In the downstream, transmitters collocated at the DSLAM cooperate to eliminate crosstalk by performing pre-subtraction of the crosstalk 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 crosstalk and here there are more degree of freedom for mitigating alien-FEXT and non-crosstalk noises like impulse noise, NBI, PLC interference, etc.

As mentioned in the Introduction, vectoring performs cancellation of self-FEXT in both downstream and upstream directions but it cannot always cancel alien-FEXT, i.e. the crosstalk generated by lines that are either non-vectored or that belong to other vectored groups. Furthermore, vectoring does not cancel non-crosstalk noises like NBI, impulse noise, interference due to PLC, etc., which must be addressed by other management techniques.

The capabilities for cancellation and mitigation of alien-FEXT as well as other noises depend on whether we consider the downstream and the upstream, so we will address these two cases separately.

A. *Downstream*

In the downstream, other noises present at the receiver (alien-FEXT, non-crosstalk noises) cannot be pre-subtracted at the DSLAM and must be necessarily mitigated at the receiver or through proper selection of DSL profile parameters.

There are techniques for limiting the performance degradation due to non-crosstalk noise, for example via DSM Level 1 where more impulse noise protection is introduced on the vectored links (stronger forward error correction coding) at the expense of throughput. However, nothing can be done today against alien-FEXT at the receiver so that the only effective technique is limiting it at its source. This implies reducing the transmit power or transmit spectrum of the non-vectored lines using for example DSM Level 1 and Level 2 techniques as suggested in [8]. This means that the sources of alien-FEXT must be "managed" so that they behave as politely as possible and thus cause the least amount of crosstalk possible to the vectored lines, consistent with the service objectives for the alien lines. The benefits of management apply not only to the case of coexistence of vectored and non-vectored lines but also to the case of coexistence of multiple vector groups.

B. *Upstream*

In the upstream, receivers collocated at the DSLAM cooperate to cancel self-FEXT but can also cooperate to cancel other kinds of noise. Therefore, unlike the downstream case, in the upstream direction vectoring can cancel alien crosstalk. We remark that this capability does not require the availability of multiple pairs per subscriber as erroneously reported in [6].

In the absence of other noises, it has been shown that simple linear Zero-Forcing (ZF) cancellers are quasi optimal in cancelling upstream self-FEXT due to the column-wise diagonal dominant nature of the DSL channel matrix [14]. However, when significant alien-FEXT and other non-crosstalk noises are present, a ZF linear receiver does not perform well because the information contained in the noise covariance matrix is not exploited. In this case, linear Minimum Mean Square Error (MMSE) receivers should be used as they perform better than ZF receivers [18]. The best alien-FEXT cancellation is provided by more complex non-linear receivers that feed back the already detected symbols. For example, several schemes have been proposed in the literature: the Generalized Decision Feedback Equalizer (GDFE) [12]-[18], Turbo-MMSE receivers [19], or others [20]-[24]. Techniques like tone selection, line selection, and grouping become important for lowering the computational cost of non-linear receivers for upstream crosstalk cancellation [18], [25].

Cancellation of alien-FEXT and non-crosstalk noises is most effective when the spatial correlation of these types of noise is high, which is the case when there is a small number of noise sources. As the number of noise sources goes up, cancellation is less complete and can be poor when the number of (dominant) sources is higher than the number of vectored lines. In this case, management of alien lines to ensure they behave politely becomes very important if mitigation of the effects of alien crosstalk on the vectored system is sought.

It has also been reported that the use of UPBO for upstream vectoring benefits mixed deployments. Forouzan et al. conclude that it is generally safer to always enable the UPBO unless one is sure that all loops will be managed by the same Spectrum Management Center (SMC) [17] (see also Sect.VI.A). If this is not the case, the achievable bit rates for the long (and managed) loops when UPBO is disabled are significantly smaller than when UPBO is enabled due to higher crosstalk originating from the unmanaged loops. As discussed in Sect. III.B of [8], dynamically adjusting the UPBO of non-vectored lines can strongly improve the performance of vectored lines compared to using static UPBO, since the optimal setting of UPBO parameters for vectoring depends (among other things) on the FEXT cancellation capability of the vectored system. The capability of dynamically adjusting the UPBO has been acknowledged as very important also for non-vectored systems and has been recently incorporated in the new UK Access Network Frequency Plan (ANFP) ratified in Sep. 2011. According to the new ANFP, an operator may either use a static configuration for the UPBO or a dynamic configuration that also allows exceeding the mask under some average constraint (see UK NICC ND 1602 [35], Sect. C.3).

IV. Improving Compatibility by Adding Functionalities at the Access Node

One way to exploit full vectoring gain even when multiple service providers share the cabinet area is to create very large vector groups so that the number of possible alien disturbers goes down and the effects of alien-FEXT become smaller. A first step in this direction was the introduction of Board Level

Vectoring (BLV), where vectoring is performed over all lines in a line card. However, BLV is generally not sufficient and provides only limited scalability so that increasing the vector group size requires solutions that provide a truly modular architecture to vectoring equipment as well as additional functionalities at the DSLAM level. Along these lines, System Level Vectoring (SLV, vectoring across line cards) is the first step towards a truly scalable solution for avoiding alien crosstalk. Several companies have already commercial SLV equipment in their product portfolio.

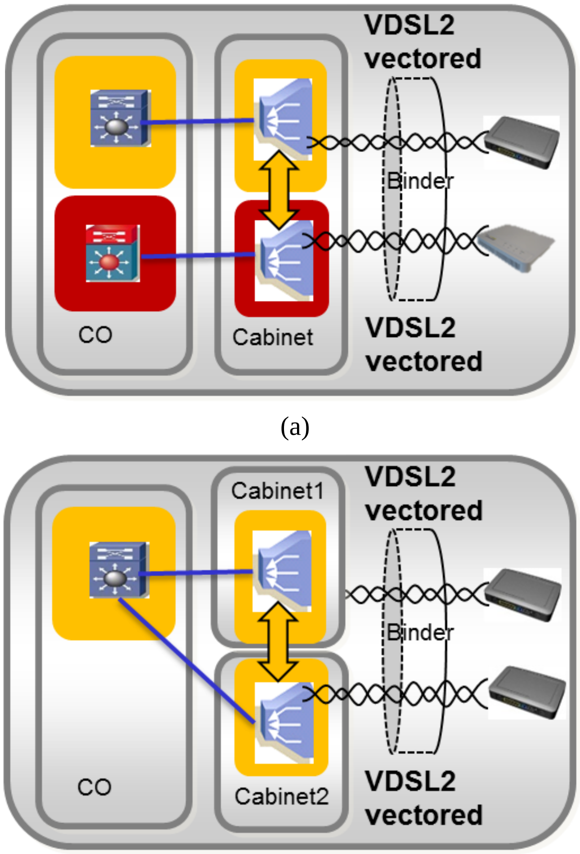

(a)

(b)

Figure 2 – An example of two DSLAMs connected via xDLV to create a larger vector group. The DSLAMs are owned by either (a) two different operators sharing the same cabinet area or (b) by the same operator.

A natural way to achieve chassis expansion for normal growth of capacity is cross-DSLAM level vectoring (xDLV) – see Figure 2. Such a solution makes it possible to maintain vectoring performance even when two or more SPs operate in the same cabinet provided that: a) the second operator deploying vectoring is willing to adopt a technology compatible to that chosen by the first SP and connects to the first DSLAM to achieve vectoring coordination; b) a shared vector processor or multiple coordinated processors manage the noise cancellation for all lines in a neutral way. Standards efforts for xDLV data exchange protocols are not yet underway so that initial xDLV implementations will likely be among DSLAMs of a single vendor. This approach seems a natural solution when considering large access nodes with hundreds of lines where there is a higher probability of finding legacy (non-vectored) lines.

Such an evolution may be desired by incumbents that already have VDSL line cards deployed in last-mile areas and that, for historical and environmental reason, are spread in multiple small chassis in small distributed street cabinets. If an incumbent SP starts delivering VDSL2 from the street cabinet area, in a city where there are hundreds of cabinets, small-size DSLAMs are initially chosen to minimize initial cost and environmental impact. This practice also depends on power and space constraints that only allow for hosting up to four line cards slots to be sustainable. As customer demand grows, the incumbent populates the DSLAM until the chassis is full. Due to size and power constraints it is not possible to expand the existing chassis in the cabinet. At this point, the incumbent ready to provide connectivity to additional customers can consider an additional cabinet. When a new cabinet is installed all new customers are wired and terminated on the new chassis. Some years later a SP willing to boost its lines with vectoring to meet its digital agenda may configure each chassis with disjoint vectoring groups, but in that case the crosstalk from lines connected to a DSLAM acts as alien crosstalk and would not be cancellable by the other DSLAM. Since rewiring cable would not be desirable in large developments, one must either mitigate the crosstalk between multiple vector groups via DSM Level 2 or completely eliminate it by increasing the vector size to include lines across multiple DSLAMs, i.e. xDLV.

With xDLV, the incumbent can solve the problem by connecting the two chassis through a proper cable to expand the size of the vectoring group. The natural growth in computational power that VCEs can handle, enables vectoring more and more lines together without having to rewire copper. This scenario assumes that the current limitations of existing SLV vectoring products to span only across line cards in a single chassis are removed. As the vectoring customer base grows, vectoring coordination across multiple DSLAMs, chassis or even cabinets may be needed regardless of their physical location.

An xDLV solution may also be desired by competitive operators since it enables the coexistence of multiple providers while maintaining enhanced vectoring performance. In this case, the incumbent and competitive providers create a single but larger vectored group by tying together their DSLAMs and letting the xDLV system process all the ports of the various DSLAMs. In this way, each operator does its own vectoring, retains control of its own lines but can also cancel via xDLV the crosstalk created by all other operators that affect its own lines. One or more spectrum management centers could provide further coordination with respect to power allocation and identification of inter-DSLAM interfering DSLs.

With SLV or xDLV, power optimization intrinsic to vectoring is achieved across all lines in a node. The possible reduction of power consumption at the node level emphasizes the importance of assessing the achievable performance of xDLV systems, followed by defining standardized interfaces and protocols for xDLV.

A. *Natural evolution from SLV to xDLV*

Cross-DSLAM vectoring is included in the near-term roadmap of some major DSL technology suppliers. Pre-commercial solutions were already available for demo at the BBWF'11 [26], at least with the constraint that all DSLAMs are from the same vendor. Ikanos was the first chipset vendor to announce a solution for supporting large vector groups (NodeScale vectoring) [27]. Huawei is developing a solution under the name of Node Level Vectoring (NLV) where two SLV units over chassis are vectored together [28]. Also Alcatel-Lucent foresees an extension of SLV allowing multiple chassis to share a common vectoring processor. Such an extension is a "multi-system vectoring" solution that allows vectoring "across multiple line cards in collocated access nodes" – for more details see [29]. "Multi-system Vectoring" is also described in [6] as a possible solution for multiple service providers even if with some operational impact. The available data seems to suggest that there is a growing attention to xDLV. Some cautious optimism can be expressed for its feasibility, at least when the different DSLAMs are physically located close to each other since xDLV opens some technical challenges that need industrial solution as described below.

There are several challenges related to cross-DSLAM vectoring that still deserve some attention: a huge amount of information needs to be communicated between line cards and the VCE; symbol timing must be distributed to the multiple chassis; the cable that delivers both clock and crosstalk information between the chassis must be reliable; if the architecture encompasses multiple VCEs, the computational load should be balanced among the VCEs.

There seems to be a reasonable path towards a solution for several of these challenges, at least for DSLAMs in close proximity, thanks also to the development efforts toward the realization of SLV products. High bandwidth channels between line cards and VCE are available today. The distribution of symbol timing can be addressed via clock over Ethernet and industry standard IEEE QSFP 40Gb/s cables can be used for both sync timing distribution and crosstalk information data flow between two DSLAMs provided that the chassis are quite near. A reduction in the data exchange and computational requirements for supporting xDLV can also be achieved by cancelling only the strongest disturbers and a subset of tones (partial cancellation).

The possibility of sharing computational load across multiple VCE brains is strictly dependent on the implementation of the cancellation algorithm. In a centralized approach, a shared vector processor can cancel the worst disturbers for all lines in the vectored group while remaining neutral with respect to which providers the cancelled disturbers belong to. Other parallel architectures employ multiple VCEs and resort to load partitioning to reduce the computational load per VCE.

Vectoring across multiple small and closely located DSLAMs seems to be close to becoming a commercial reality, but more studies on frame loss, latency and packet jitters from using Ethernet or other protocol over these cables at some distance are still needed. Furthermore, the effects that these non-idealities have over real time crosstalk cancellation need to be analyzed for assessing the distance limitations constraints for xDLV.

B. *Effectiveness of SLV and xDLV*

Vectoring is beneficial if the lines in the vectored group cause significant crosstalk to each other. As the size of the vectored group increases, there are an increasing number of lines that in practice cause little or no crosstalk to other lines in the vector group. This is the case when the vector group expands to include lines that span different binders and cables. This may not often be the case with BLV since all ports supported by the same line card are likely to belong to the same binder. However, it can happen more frequently with SLV and, of course, with xDLV. In the case of large vectoring groups, a mechanism for individuating the subsets of lines that cause strong crosstalk to each line in the vector group is helpful for lowering system complexity.

If inter-binder crosstalk is negligible or if all lines in the binder terminate on a single DSLAM, then there is no advantage in terms of data rate performance in doing xDLV and, from a complexity point of view, it is better to have separate vectoring groups – one per DSLAM. The case that is usually found in the field is that indeed inter-binder crosstalk is present. In this case, a single vectoring group can support higher data rates than separate groups because, in the latter case, the lines in different binders that have strong coupling with each other and belong to different vector groups create alien crosstalk to each other. On the other hand, not many pairs create substantial inter-binder crosstalk and, if it happens, it does not happen for all tones. In this case, the VCE of an xDLV system can spare resources and relax data exchange requirements as crosstalk cancellation complexity depends on fewer lines and is performed on fewer tones than in the case of full vectoring across all the lines supported by the multiple DSLAMs.

Note that if all lines in a binder do not terminate on a single line card or DSLAM, then it is possible in a SLV/xDLV system to identify which pairs are terminating on which line card/DSLAM thus guiding operators on where to perform binder management so that lines creating strong crosstalk to each other can be grouped together on the same line card/DSLAM. However, binder management is labor intensive or requires expensive automatic cross-connects, so that the knowledge of which are the strongest interferers across line cards/DSLAMs can be better exploited to perform partial cross-vectoring among the most severe inter-line cards and inter-DSLAM crosstalkers. This greatly reduces the requirements on data transfer to/from the VCEs and the computational complexity required to perform crosstalk cancellation via SLV or xDLV.

Practical limitations on complexity and data transfer speeds from/to VCEs may allow only partial vectoring via SLV/xDLV thus leaving some lines in an access node excluded from the large vector group created via xDLV – and this could happen even if all the lines in the access node are managed by a single SP. These possibly excluded lines could then become the source of alien crosstalk that couples into the vectored group, thus confirming that alien crosstalk can be present even if only a single SP is present in an access node. Going forward complexity limitations may disappear allowing full cross-vectoring, but we point out that even today there are tools that allow multiple operators to share the cabinet area and enjoy vectoring benefits even if not all lines are fully cross-vectored. This provides an evolutionary path and also allows for a gradual deployment of vectoring while still allowing competition in the local loop.

V. Dynamic Spectrum Management (DSM)

Alien crosstalk is likely to be initially present regardless of whether there is a single or multiple SPs that have access to the lines in a node, regardless of whether vectoring is introduced gradually or large vector groups are created via SLV or xDLV. As pointed out in the previous section, resorting to SLV or, when possible, to xDLV, can limit the generation of alien crosstalk to some extent, but does not completely eliminate it even if a single operator is present in an access node. Thus, in practice, limitations to unbundling do not eliminate the need for additional tools aimed at mitigating the effects of alien crosstalk on vectored lines.

Among these tools, DSM has been shown to be particularly effective at mitigating alien crosstalk and allowing vectored lines to experience close to ideal performance. As recently reported in [8], there is a wide consensus on the effectiveness of DSM in ensuring a good degree of compatibility between vectored and non-vectored lines.

A Spectrum Management Center (SMC) can use DSM techniques to control the impact of alien crosstalk. In unbundled environments, each operator may have their own SMC and, if each of those SMCs independently follows certain mutual "politeness" rules, then it is still possible for all SPs to enjoy the majority of the benefits of vectoring and mitigate alien crosstalk. Even better results can be achieved if the SMCs of the various operators can exchange information, or if the DSL access infrastructure is shared among the various providers and is managed by a single management system that has a full view of the network.

Good results have been reported on mitigation techniques that are based on DSM Level 1 (rate limiting, flat power back-off) and DSM Level 2 (spectrum balancing). DSM-based techniques can be used for mixed deployments in both upstream and downstream directions and do not require adding any complexity in the transceivers or additional functionalities at the access node as SLV and xDLV entail. Simulation results were reported in [7], [8], where also the effects of simple DSM Level 1 mitigation were investigated for the case of downstream in a mixed scenario.

Figure 3 shows the performance of 24 vectored lines when 24 non-vectored lines are present in the same cable. The transmit power of the non-vectored lines is managed by limiting their transmission rate at 25 Mbps and with a 6 dB margin. The simulation confirms that applying DSM Level 1 to the alien disturbers allowed the vectored lines to achieve an almost ideal performance in a mixed scenario with a 48% cable fill (dense deployment): 100 Mbps downstream out to 500m for 99% of subscribers. As the cap imposed on the non-vectored lines increases, the effects of alien crosstalk are more noticeable and some of the vectored lines perform worse.

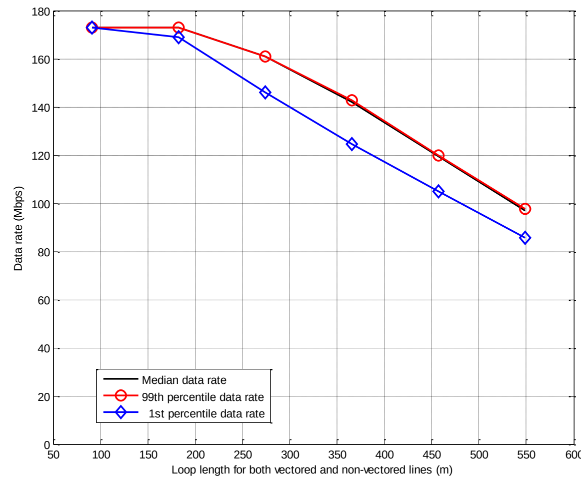

Figure 3 - Rate-reach plot for a group of 24 vectored lines when 24 non-vectored lines are also present in the same cable and their data rates are capped at 25 Mbps. For simulation details, see the Appendix.

Figure 4 displays the achievable vectored data rates when the rate limit on the non-vectored lines is increased to 45 Mbps. The Figure confirms that, in in a mixed scenario with a 48% cable fill (dense deployment), vectored lines can still achieve 100 Mbps downstream out to 300m for 99% of subscribers if the non-vectored lines are rate limited to 45 Mbps and their powers managed accordingly - which is a speed higher than many current VDSL2 deployments. These results confirm that, even in the presence of alien crosstalk, vectoring offers substantial benefits as compared to non-vectored VDSL2 where, in the case of dense deployments, 100 Mbps downstream can only be achieved out to 150m for 50% of subscribers (see Figure 1). More details on the effectiveness of DSM Level 1 and 2 in supporting vectoring are in [7], [8].

Some arguments have been made that non-vectored lines "harm" vectored lines [5], [6]. This is only meaningful if one states "how much harm" is made. The results in Figure 3 and Figure 4 as well as the results reported in [7], [8], quantify the effect that non-vectored lines have on vectoring. When properly managed via DSM, non-vectored lines whose data rate is capped at 25 Mbps only reduce the speed of vectored lines by 5-10% in 99% of cases. If the data rate of non-vectored lines is capped at 45 Mbps, then the speed of vectored lines decreases by less than 5% in 50% of cases and by 30% only in the 1% worst case. We remark that these results are for a high cable fill (48%) and that a speed of 45 Mbps for non-vectored VDSL2 is a rather high data rate that 99% of VDSL2 customers within 700 meters would be able to get in a dense VDSL2-only deployment. As additional VDSL2 lines become vectored over time (gradual deployment) and experience a decrease of transmitted power, the overall crosstalk affecting the remaining non-vectored line goes down and the rate reach of the remaining non-vectored lines will go up. These results as well as the results reported in the available literature [8] have been ignored in the analysis reported in [5], [6].

If the management system has a full view of the DSL system (see also Sect. VI.A), the trade-offs between the data rates of the vectored and non-vectored lines can be negotiated and optimized automatically between service providers. Similar performances could be achieved when the various SPs use separate management systems since the above results have been obtained using DSM Level 1 which performs independent processing of each line. Operators can independently exercise politeness rules, something that does not require communication between SMCs. A good example of this is the new dynamic PSD mask specification for upstream VDSL2 in the latest (2011) version of the UK Access Network Frequency Plan (ANFP) (see [35], Sect. C.3). Of course, if operators agree to exchange information between their SMCs, they would experience an increase in performance beyond that available with independent politeness rules.

We also remark that DSM Level 1 and Level 2 provide benefits in a DSM Level 3 setting for reasons that go beyond compatibility with alien disturbers. Even in a fully vectored scenario (no aliens), non-crosstalk noises like impulsive noise or time-varying PLC interference become even more disruptive as the partial masking provided by crosstalk vectoring is enabled and the crosstalk levels are reduced. A management system that judiciously applies DSM can compensate for all impairments in the network.

The effectiveness of DSM has also been recognized in various standardization committees. For example, ATIS has issued a first report on DSM in 2007 [30] and a new one (Issue 2) is expected to be released by the end of 2012 [31]. The new DSM Report Issue 2 also contains details on the use of DSM Level 2 and Level 3, the control parameters, and the gains that can be achieved in typical scenarios.

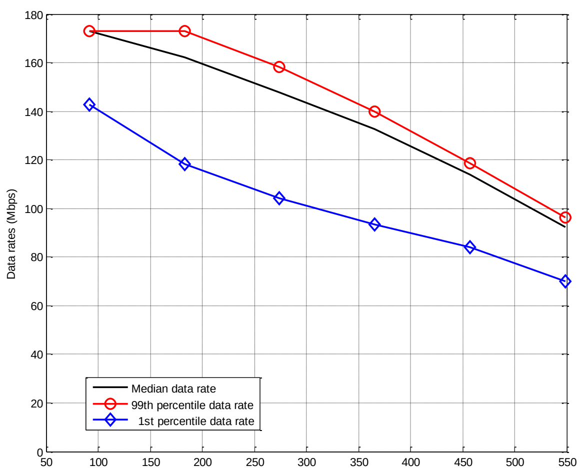

Figure 4 – Same as Figure 3, but for non-vectored data rates capped at 45 Mbps.

Another body that has shown significant interest in DSM and vectored DSL is the Network Interoperability Consultative Committee's (NICC) DSL Working Group that has recently produced a report on DSM methods in the UK access network [32]. This document describes the benefits of DSM Level 1 techniques, by showing simulation results for a number of different scenarios and concluding that "*DSM Level 1 is undoubtedly a useful capability for UK DSL network operators*" and that "*Improved politeness and consequential reduction in transmit power and crosstalk levels may bring benefits to UK DSL operators and their end-users*." As a consequence, the use of "Virtual Noise" should be avoided where vectoring is deployed either alone or in a mixed scenario as it makes DSL lines behave impolitely. In fact, as argued in [32], Virtual Noise reduces the attainable speed while possibly increasing power consumption for that line and increases the crosstalk into neighboring lines in the same cable/binder. Caution against the use of Virtual Noise for similar reasons is also expressed in [1]. Follow-up work on DSM Level 2 [33] and DSM Level 3 [34] is currently under study in NICC.

VI. ENABLING EFFECTIVE USE OF DSM

On the basis of the literature survey presented in [8] and of the arguments made here, we conclude that alien crosstalk harms vectoring only when the alien lines are left completely unmanaged. If the non-vectored lines or the other vectored groups are properly managed, the impact on the vectored lines is limited and predictable, and coexistence is indeed possible. Thus, vectored lines deployed along-side non-vectored lines can achieve excellent performance and retain most of the benefits of vectoring when an appropriate use of DSM techniques is made. Furthermore, when resorting to SLV/xDLV, operators may have to leverage DSM as complexity constraints do not exclude the possibility of having alien crosstalk – even in the case when a single operator is present at the cabinet.

A DSL management system must support an efficient and automated way to apply the following process across all lines in a DSL network [11]:

- *Collecting* operational and performance parameters from all DSLs in the network on a daily basis, and storing these parameters for long periods of time (days to months).
- *Analyzing* the stored parameters to either diagnose faults (e.g., copper impairment, DSL equipment fault. etc.), or to obtain performance projections, such as identifying lines that are eligible for upgrade. These analyses can then be provided to other operations support systems, or to customer care agents requiring such information.
- *Reconfiguring* DSL lines (also known as 'reprofiling') to meet coexistence objectives, satisfy quality of service requirements for each line, and maximize data rate based on the lines' service requirements. Only those lines that are not meeting coexistence or service objectives — usually defined in terms of rate and stability — are reprofiled.

Meaningful operational benefits are obtained only when the steps above are performed regularly on all lines in the network, preferably on a daily basis. The step of collecting management data daily from the DSL access network is followed by the diagnostics phase, which is performed for all lines, and by the reprofiling phase, which is performed only for those lines that do not meet their service objectives.

DSL standards, and in particular Recommendation ITU-T G.997.1 (PLOAM), require that a range of diagnostic and control parameters be made available to DSL management centers. Furthermore, SPs willing to leverage DSM need to require from vendors support for the following:

- Quantities defined in the standard should be exposed by DSLAM vendors, e.g., via SNMP or other management protocols, and any additional quantities necessary to effectively manage vectoring should be exposed (including profile and chassis management).
- A consistent, standards-compliant management interface across all DSLAMs to support management of mixed vectored/non-vectored and multi-vendor deployments.
- Direct access to the DSLAM for DSM and diagnostic purposes, as mediation systems generally filter data and are sometimes under-dimensioned.
- Adequate processing power and memory for the DSlAM management agent (e.g., SNMP agent) to enable all 15 minute performance counters to be retrieved for each line every day. Support for batch information retrieval should also be present.
- Compliance with the relevant open standards: RFC2662/RFC3440 (ADSL), RFC4706 (ADSL2) and, in particular, RFC5650 for VDSL2[1].

Ongoing industry performance and interoperability efforts such as the Broadband Forum's "DSL Quality Management" (TR198) should also be supported.

A. *DSM for vectoring in unbundled environments*

An interesting aspect of DSM is that its benefits can be enjoyed whether unbundling is allowed or not by the regulatory frameworks.

In Figure 5, 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. This can occur either because the provider has a geographical monopoly, or because the provider may only need to resell services to other providers as a bitstream, instead of providing full access to the copper plant. In this case, all lines are controlled by a single entity and the vectored DSL access infrastructure is owned by the same entity. The management system for vectored DSL 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.

Another case is when requirements for physical unbundling are present and two SPs must share the copper twisted pairs, and choose to use separate DSLAMs and separate management systems. The separate management systems do not have a full view of the entirety of twisted pairs in the network, and thus have limited diagnostics and reprofiling capabilities with respect to the lines controlled by the other management system, when compared to the case of a single system managing all lines in the network (see Figure 6). This is the case when certain types of advanced diagnostics processing consider multiple pairs at once in their calculations.

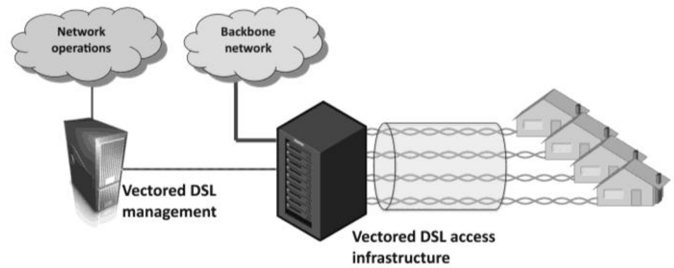

Figure 5 – Vectored DSL access infrastructure and management owned by a single provider.

As a result, some performance loss is expected for vectored DSL in this case, because of a reduced ability to coordinate the vectored DSL systems of provider A and non-vectored DSL systems of provider B, or the two vector groups that belong to A and B. 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. Such reporting allows provider B to request that provider A corrects the situation.

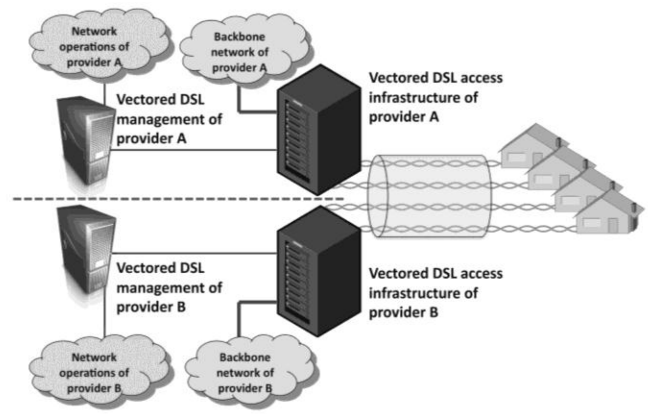

Figure 6 – Vectored DSL access infrastructure and management separately owned by multiple providers.

A second unbundled architecture that achieves competition in the local loop with vectored DSL, but which results in better efficiencies than in the previous example, is when the vectored DSL access infrastructure is shared among the providers and is managed by a single management system with a full view of all vectored DSL lines. As shown in Figure 7, each of the providers sharing the cabinet area still has access to the management functions but with appropriate restrictions to prevent disclosure of proprietary 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. At the same time, each provider can define and manage its services independently, subject to overall rules on fairness. In this case, coupling DSM with SLV and xDLV in access nodes where multiple SPs are

[1] Plans to consolidate all those separate RFCs in a single DSL RFC are being made in the BBF.

present provides the same performance benefits of vectoring as when only one single SP has control of the entire copper infrastructure, while permitting each SP to define its own differentiated offerings.

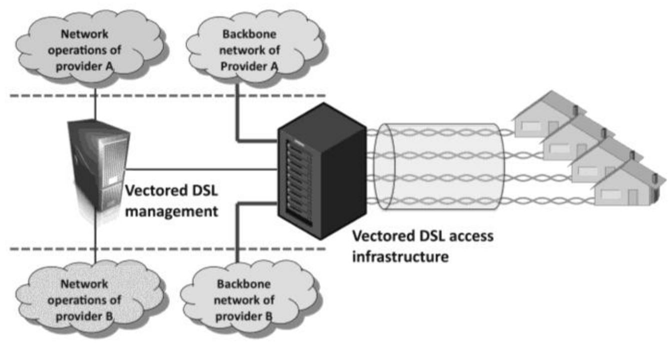

Figure 7 – Vectored DSL access infrastructure and management shared by multiple SPs. The DSLAM icon may represent a single DSLAM or also multiple ones (e.g., one or more DSLAMs per SP) and the DSLAMs of the various SPs can be cross-vectored via xDLV.

VII. Additional Management Considerations for the Support of Vectoring

DSM provides benefits to vectored deployments that go beyond ensuring compatibility with alien disturbers. Vectored systems can be managed to allow improved diagnostics based on the knowledge of crosstalk couplings between pairs and improved mitigation of the effects of impulse noise and of other time-varying noise sources [11].

The deployment of Vectored DSL systems has important implications on Operations and Maintenance practices. The vectoring benefits discussed in the previous sections are maximized when such practices are enhanced to account for the new effects and capabilities introduced by vectored DSL. An overview of how these practices can be enhanced follows.

A. *Fault Detection*

Vectored DSL systems can report through the management interface the actual crosstalk coupling among pairs, both for upstream and downstream. The availability of this quantity (also known as XLIN) leads to many capabilities in diagnostics, troubleshooting, management, and planning of VDSL2, cellular backhaul, and femtocell deployments.

XLIN makes it possible to identify lines that create excessive crosstalk. Typically, such lines are characterized by faults (e.g., poor balance) that lead to poor performance and that are prime targets for maintenance actions. It is possible that such pairs also generate crosstalk into non-vectored lines (which cannot be eliminated) so that the identification of such extreme crosstalkers and subsequent actions based on such information can improve the performance of the copper network over time. This can substantially reduce the time spent by technicians in trying to identify an offensive noisy line. XLIN can also indicate anomalies and potential root causes, such as determining if crosstalk cancellation is limiting or inhibiting performance or whether there are problems with excess crosstalk in the vectored group.

B. *Service Prediction*

The knowledge of XLIN also allows for more accurate prediction of the expected gains from FEXT cancellation. Knowing the expected gains provides essential guidance for choosing the line priorities, and to determine whether a service upgrade is feasible. This allows tuning of the DSLs to the demand for data rate and services among the customers.

C. *Line Prioritization*

The signal processing operations of vectoring are such that it is possible to direct the performance benefits towards certain lines and certain frequencies. The management interface allows the SP indirect control of the computational resources and of the signal processing operations of vectoring, for example to favor a line that is being offered a higher-end service. The configuration parameters that can be manipulated for this purpose include the power, minimum data rate, target data rate, maximum data rate and the line priority. The best practice is to configure these quantities individually for each line, based on both the line's capabilities and the service requirements.

D. *Management of Non-Crosstalk Noise Sources*

Non-crosstalk noise sources can be of two types, where each type has a different effect on vectored DSL systems: noise sources that become dominant after crosstalk is removed; noise sources that cannot be mitigated through vectoring. Although treatment of non-crosstalk noises is critical to achieving the full benefits of vectoring, very little has been said about this compared to the attention that the industry is given today to the more limited matter of handling alien crosstalk.

One example of the first type of noise is impulsive noise that is normally masked by crosstalk and that starts affecting performance only after vectoring is enabled. A second example is time-varying noise like PLC interference that may lead to line instability. Both examples occur even in a fully vectored scenario (no aliens) where a single operator controls all DSL lines. The treatment of such disruptions requires management algorithms that appropriately configure each line for impulse noise protection, or for coping with abrupt noise changes. From this point of view, we emphasize that DSM Level 1 and Level 2 become fundamentally important for the support of vectoring because, when vectoring is enabled, non-crosstalk noises become the dominant source of impairment.

In some cases of downstream transmission, it is not possible to cancel certain noise sources through vectoring. Such noise sources may include AM noise, various types of Radio-Frequency-Interference (RFI), or crosstalk from legacy DSL systems or other vector groups. Typically, such sources affect a specific set of frequencies. If such a set of frequencies is found to be affected from this kind of interference, the proper management action is to instruct the vectored DSL system to disable vectoring over those frequencies so that vectoring computational resources are directed to more productive uses.

E. *Fairness in Unbundled environments*

Depending on the specific regulatory environment, knowledge of upstream XLIN can become vital information to ensure fairness among several SPs. For example, in cases where bit-stream unbundling for VDSL is carried out with one provider per line, the availability of XLIN becomes essential for management of the upstream and downstream cable interests between different operators.

VIII. Conclusions

In this paper, we have addressed the issue of supporting vectoring in the presence of alien-crosstalk – a situation that arises from a variety of conditions including the case when only one operator is allowed to manage the copper plant. In fact, since there is no economic argument for the complete and simultaneous replacement of the existing plant with a vectored plant, the deployment of vectoring will be gradual. We have also argued that, although full cancellation of alien crosstalk is not always possible, mitigation techniques for minimizing the impact of alien crosstalk are available and various methods have been discussed. These conclusions are also supported by the existing literature.

This paper confirms that coexistence between lines within a vector group and outside of it is indeed possible. Various alternative solutions are viable and available to SPs: ensuring the coexistence of multiple SP by either eliminating alien crosstalk (xDLV) or by mitigating it (DSM and management), or a combination of both. We have discussed the desirability of a modular architecture for vectoring equipment so that hierarchical structures like SLV and xDLV can be supported to handle large vector groups. Furthermore, we have discussed how also DSM Level 1 and 2 are all very useful for the effective support of DSM Level 3 (vectoring) regardless of whether the copper plant is administered by a single or multiple operators.

The common denominator for the effective support of mitigation techniques is the availability of open, efficient, standardized interfaces that enable the best possible use of the data available at the DSLAM. We have also seen that coexistence of multiple SPs may be achievable with xDLV and that standard interfaces that allow the support of xDLV across DSLAMs from different vendors are desirable.

On the basis of these considerations, we conclude that:

- coexistence between vectored and non-vectored lines as well as between multiple vector groups owned by multiple SPs is not only technically feasible but would also entail limited performance loss on vectored lines; and
- large vector groups architectures can include lines from multiple SPs.

If vector groups are too large for full cross-vectoring due to complexity limitations or cost, alien crosstalk resulting from partial vectoring can be managed. Similarly, interference from non-vectored lines, or additional vector groups, can all be handled today with DSM – regardless of whether all the lines are managed by a single SP or not.

We conclude confirming that the present status of technology does not justify any restriction to local unbundling and that competition among SPs in the local loop does not void the benefits that vectoring can ensure. SPs have today various tools at their disposal for providing next generation broadband services in a competitive environment and at competitive costs. Therefore, any restriction to unbundling seems at this point inappropriate and, ultimately, damaging to the end user.

Appendix – Simulation Assumptions

For the simulations presented here, we have used downstream VDSL2 profile 17a with PSD limit mask 998ADE17-M2x-B (B8-12), with crosstalk modeled according to the ATIS model for MIMO channels in ATIS-0600024. 48 DSL pairs were selected randomly from a cable with 100 pairs and four binders, and the gauge of the pairs was AWG26. For the mixed vectored/non-vectored simulations, 24 pairs were assigned randomly to a single vector group and the other 24 randomly chosen pairs were non-vectored.

For each loop length, a set of pair-to-pair couplings was randomly generated from the ATIS model and 500 random pair selections were made to generate statistics representing different assignments of vectored and non-vectored lines. Dynamic Power Back-off (DPBO) was not enabled and the background noise was assumed to be equal to -140 dBm/Hz.

When the downstream data rate of non-vectored lines was capped, their SNR margin was set to 6 dB, consistent with the application of DSM Level 1 for INP management, and power management was also applied.

References

[1] *Very High Speed Digital Subscriber Line Transceivers* 2 (VDSL2), ITU-T Rec. G.993.2, Feb. 2006.

[2] G. Ginis, J.M. Cioffi, "Vectored Transmission for Digital Subscriber Line Systems," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 5, June 2002.

[3] *Self-FEXT Cancellation (Vectoring) for Use with VDSL2 Transceivers*, ITU-T Rec. G.993.5 (G.vector), Jan. 2010.

[4] V. Oksman, H.Schenk, A. Clausen, J. M. Cioffi, M. Mohseni, G. Ginis, C. Nuzman, J. Maes, M. Peeters, K. Fisher, P.-E. Eriksson, "The ITU-T's New G.vector Standard Proliferates 100 Mb/s DSL," *IEEE Commun. Mag.*, vol. 48, no. 10, Oct. 2010.

[5] M. Guenach, J. Maes, M. Timmers, O. Lamparter, J.-C. Bischoff, M. Peeters, "Vectoring in DSL Systems: Practices and Challenges," *IEEE Globecom Conf.*, Houston, TX, Dec. 5–9, 2011.

[6] S. Vanhastel, P. Spruyt, "VDSL2 Vectoring in a Multi-operator Environment – Separating Fact from Fiction," Alcatel-Lucent TechZine, March 21st, 2012. [Online]. Available: http://www2.alcatel-lucent.com/blogs/techzine/2012/vdsl2-vectoring-in-a-multi-operator-environment-separating-fact-from-fiction/

[7] M. Mohseni, G. Ginis, J. Cioffi, "Dynamic Spectrum Management for Mixtures of Vectored and Non-vectored DSL Systems," *IEEE Conf. on Inform. Sciences and Syst.* (CISS), Princeton, NJ, Mar. 17–19, 2010.

[8] 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.

[9] K. Song, S. Chung, G. Ginis, J. Cioffi, "Dynamic spectrum management for next-generation DSL systems," *IEEE Commun. Mag.*, vol. 40, no. 10, Oct. 2002.

[10] K.J. Kerpez, D.L. Waring, S. Galli, J. Dixon, P. Madon, "Advanced DSL management," *IEEE Commun. Mag.*, vol. 41, no. 9, Sep. 2003.

[11] G. Ginis, M. Goldburg, and J. M. Cioffi, "The effects of vectored DSL on network operations," *Journal of Telecommunications Management*, Vol. 3, No. 2, July 2010.

[12] M.K. Varanasi, "Decision Feedback Multiuser Detection: A Systematic Approach," *IEEE. Trans. Inform. Theory*, vol. 45, no. 1, Jan. 1999.

[13] G. Ginis, J.M. Cioffi, "Vectored Transmission for Digital Subscriber Line Systems," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 5, June 2002.

[14] R. Cendrillon, G. Ginis, V. den Bogaert, M. Moonen, "A Near Optimal Linear Crosstalk Canceler for Upstream VDSL," *IEEE Trans. Signal Process.*, vol. 54, no. 8, Aug. 2006.

[15] G. Ginis, C.-N. Peng "Alien Crosstalk Cancellation for Multipair Digital Subscriber Line Systems," *EURASIP J. on Applied Signal Processing*, vol. 2006, Article ID 16828, 2006.

[16] C.-Y. Chen, K. Seong, R. Zhang, J.M. Cioffi, "Optimized Resource Allocation for Upstream Vectored DSL Systems With Zero-Forcing Generalized Decision Feedback Equalizer," *IEEE J. Sel. Topics Signal Process*, vol. 1, no. 4, Dec. 2007.

[17] A. Forouzan, M. Moonen, J. Maes, M. Guenach, "Joint Level 2 and 3 Dynamic Spectrum Management for Upstream VDSL," *IEEE Trans. Commun.*, vol. 59, no. 10, Oct. 2011.

[18] P.K. Pandey, M. Moonen, L. Deneire, "MMSE-based partial crosstalk cancellation for upstream VDSL," *Signal Processing (Elsevier)*, vol. 92, no. 7, July 2012.

[19] I. Wahibi, M. Ouzzif, J. LeMasson, S. Saoudi. "Stationary Interference Cancellation in Upstream Coordinated DSL Using a Turbo-MMSE Receiver," *Intern. J. of Digital Multimedia Broadcasting*, vol. 2008, Article ID 428037, 2008.

[20] G. Ginis, J.M Cioffi, "On the relation between V-BLAST and the GDFE," *IEEE Commun. Lett.*, vol.5, no.9, Sept. 2001.

[21] A. Forouzan, L. Garth, "Computationally Efficient Partial Crosstalk Cancellation in Fast Time-Varying DSL Crosstalk Environments," *EURASIP J. on Advances in Signal Processing*, vol. 2007, Article ID 72041, 2007.

[22] P. Biyani, A. Mahadevan, P. Duvaut, S. Singh, "Cooperative MIMO for alien noise cancellation in upstream VDSL," *IEEE Intern. Conf. on Acoustics, Speech and Signal Processing* (ICASSP), Taipei, Taiwan, April 19-24, 2009.

[23] L. Youming, H. Guanhua, S. Wei, "Blind adaptive crosstalk cancellation for multi-user VDSL systems," *Conf. on Commun. and Mobile Computing (CMC)*, Shenzhen, China, Apr. 12–14, 2010.

[24] S. Zafaruddin, S. Prakriya, S. Prasad, "Iterative Receiver Based on SAGE Algorithm for Crosstalk Cancellation in Upstream Vectored VDSL," *ISRN Commun. and Networking*, vol. 2011, Article ID 586574, 2011.

[25] R. Cendrillon, M. Moonen, G. Ginis, K. Van Acker, T. Bostoen, P. Vandaele, "Partial Crosstalk Cancellation for Upstream VDSL," *EURASIP J. on Applied Signal Processing*, vol. 2004, no. 10, 2004.

[26] Huawei FAN at BBWF, Heavy Reading, Oct. 2011. [Online]. Available: http://downloads.lightreading.com/wplib/huawei/Research_note_Huawei_FAN.pdf

[27] Ikanos website, "Ikanos Unveils NodeScale Vectoring to Deliver Broadband at 100 Mbps and Beyond for One-Tenth the Cost of Fiber" [Online]. Available: http://www.ikanos.com/press-releases/ikanos-unveils-nodescale-vectoring-to-deliver-broadband-at-100-mbps-and-beyond-for-one-tenth-the-cost-of-fiber/

[28] Huawei website, "Huawei Develops World's First Node Level Vectoring Prototype" [Online]. Available: http://www.huawei.com/en/about-huawei/newsroom/press-release/hw-093919-nlv-prototype-announce.htm

[29] Alcatel-Lucent website, "Turbocharge copper networks to extract more value" [Online]. Available: http://www.alcatel-lucent.com/wireline/copper-access.html.

[30] *Dynamic Spectrum Management*, ATIS Std. ATIS-0600007, May 2007.

[31] *Dynamic Spectrum Management - Issue 2*, ATIS Std. Under ballot.

[32] *Report on Dynamic Spectrum Management (DSM) Methods in the UK Access Network*, UK NICC ND 1513 (2010-01).

[33] *Dynamic Spectrum Management* (DSM), NICC Work Item 308 – under study, 2012.

[34] *Vectored VDSL performance in the UK*, NICC Work Item 309 – under study, 2012.

[35] *Specification of the Access Network Frequency Plan applicable to transmission systems connected to the BT Access Network*, UK NICC ND 1602 (2011-09).

Andrea Colmegna received his full-score MSc in Telecommunication Engineering, an advanced university master on advanced transmission systems and a master on telecommunication hardware technologies from Politecnico University in Milan. He joined FASTWEB in 2000 and presently is the IP Backbone Technologies Manager. Since 2006 he is directly responsible for FASTWEB IP backbone being involved in most of the new technologies tendering activities and focusing on network topology evolution and Quality of Service technologies. Since 2005 he has been involved in Standardization activities becoming the lead FASTWEB representative at BroadBand Forum (ex-DSLForum), BroadBandHome, HGI. He was editor of BBF TR69a4 where the project stream PD-199 on object addressing extensions entered the TR069 protocol for gateway management. He received a special "strategy and innovation" prize for a FASTWEB NGN project in innovation contest 2012.

Stefano Galli received his M.S. and Ph.D. in Electrical Engineering at the University of Rome "La Sapienza" (Italy). Currently, he is the Director of Technology Strategy at ASSIA, Inc. where he leads the company's overall standardization strategy and contributes to the company's efforts in the area of wired/wireless access and home area networking. Prior to this position, he held the role of Director of Energy Solutions R&D for Panasonic Corporation and was a Senior Scientist at Bellcore. He is a Fellow of the IEEE and an elected Member-at-Large of the IEEE Communications Society (ComSoc) Board of Governors. He has received the 2011 IEEE ComSoc Donald W. McLellan Meritorious Service Award, the 2011 Outstanding Service Award from the IEEE ComSoc Technical Committee on PLC, and the 2010 IEEE ISPLC Best Paper Award. He holds twenty issued/pending patents, has published over 90 peer-reviewed papers, has co-authored two book chapters, and has made numerous standards contributions.

Marc Goldburg received a Ph.D. from Stanford, an MSEE from the University of Washington and a BSE from Princeton University, all in Electrical Engineering. He is currently CTO & EVP of ASSIA, Inc., a leading developer of automated management and diagnostics tools for broadband networks. Marc is a technical founder and former CTO of ArrayComm, LLC, whose spatial processing technology is deployed worldwide in more than 300,000 wireless base stations. Later, as a consultant, he guided companies in the mobile satellite, digital television, and WiMAX carrier industries. His prior positions include a staff position at MIT Lincoln Laboratory and a research position at Stanford University. Marc is a Fellow of IEEE and was recognized by Scientific American as its Communications Researcher of the Year in 2002.