ALCATEL-LUCENT SUBMISSION

ACCC Discussion Paper examining possible variation of the service declaration for the unconditioned local loop service

AUTHORS:
MICHAEL PEETERS,
JOCHEN MAES,
MAMOUN GUENACH
AND
JAN VERLINDEN

JULY, 2007

About Alcatel-Lucent

Alcatel-Lucent Australia Limited ("Alcatel-Lucent") is the global leader in broadband access technologies and has designed and deployed fixed and mobile broadband networks in most of the world's leading economies.

Alcatel-Lucent is proud to supply equipment and services to Australia's leading telecommunications incumbents and competitors. It has supplied the infrastructure for a significant portion of Australia's residential DSL community, making it a leader in helping Australians access the advantages of a digital lifestyle. Its solutions achieve advances in DSL, fibre optics, wireless and satellite access that help companies and individuals get maximum benefit from fast Internet services.

Alcatel-Lucent's commitment to Australia is not new. It has been part of the Australian telecommunications fabric since 1895.

Alcatel-Lucent's leadership in the development of Australia's communications infrastructure has included the country's first undersea cable network, the introduction of broadband Internet, the country's first 3G mobile network (m-Net) and the world's longest optical link, between Adelaide and Darwin.

Alcatel-Lucent wishes to continue to play a leading role in improving Australia's economic outlook and standard of living by ensuring that the community has access to a rich variety of broadband services, wherever they live.

Background to this submission

Telstra is major customer of Alcatel-Lucent and has engaged Alcatel-Lucent as a strategic partner. Alcatel-Lucent and Telstra have been collaborating to evaluate alternate candidate technologies for deployment as part of Telstra's proposed access transformation. The candidate DSL technologies include ADSL2+ and VDSL2.

The international community of DSL experts is currently engaged in the standardisation of a technique known as 'DSM' which promises to dramatically improve DSL performance in future generations of the technology. In July 2007, as part of the strategic partnership, Telstra has engaged Alcatel-Lucent to discuss and provide technical advice on the capabilities, characteristics and limitations of 'DSM' technology and its relevance to the roll out and development of sub-loop. This submission is its result and aims to introduce 'DSM' technology for an interested audience.

Alcatel-Lucent does not accept any liability in respect to the use of or reliance upon the advice. Alcatel-Lucent is pleased to offer this submission to the ACCC as an assistance for highlighting issues Alcatel-Lucent believes will be central in helping appropriately resolve sub-loop unbundling issues.

Local Loop Unbundling and DSM

Michael Peeters, Jochen Maes, Mamoun Guenach and Jan Verlinden,

Alcatel-Lucent DSL eXPert Team, Antwerp, Belgium

Introduction

Next generation services require higher broadband throughput than typically available from exchange based DSL. The deployment of fibre deeper into the access network is necessary to improve DSL broadband services for the majority of the community. This is the basis for the Fibre To The Node ("FTTN") architecture.

If too many DSL lines are incapable of supporting next generation applications, the service provider's addressable market will be limited, perhaps to the extent that investment is unviable. By taking current technological developments into consideration, which could lead to significant throughput improvements, the benefit from today's access transformation projects (FTTN) could be extended in terms of higher data speeds.

This paper focuses upon a specific enhancement which is being called 'Dynamic Spectrum Management' ("DSM"). The DSM approach compensates for DSL performance impairments introduced by cross-talk, the interference of one DSL service upon other DSL services sharing the same cable. Crosstalk has an effect upon most lines and is the limiting factor for FTTN performance.

Practically, management of power levels and spectral usage (DSM up to level 2) aims to improve the performance of the poorest lines while minimising the impact upon other lines. DSM level 3 goes one step further to process the transmitted signals in such a way that the crosstalk is effectively cancelled, boosting throughput for all FTTN lines, benefiting all users. Service improvements available through a DSM level 3 deployment promise to be significant. Compared against the incremental improvements achievable through the more basic DSM levels (0, 1 and 2), it is highly desirable to engineer today's FTTN deployments for compatibility with future DSM level 3 technologies.

Although FTTN can dramatically improve the broadband rates available to end users, longer FTTN lines remain intrinsically less capable of matching the throughput of shorter FTTN lines. In order for an application provider to offer a standardised service portfolio within a given market, the portfolio must fit within the capabilities of all lines. Thus the longest FTTN lines generally establish base-line service characteristics.

Specifically because of DSM improvement for the least capable lines, an access provider's minimum throughput guarantee can be increased and with an assurance of higher minimum throughput, a more capable basic service can be offered. The capability to offer improved services in the future will increase the intrinsic value of the FTTN investment and extend its overall life.

Key recommendation

Achieving the significant benefits offered by DSM necessitates a change in our approach to DSL deployment. In order to work efficiently, all lines that interact with each other need to be coordinated. To effectively deploy and benefit from the more advanced DSM level 3 techniques, an entire cable binder[1] needs to be controlled from the same line-card / DSLAM. Most of DSM's benefits are not accessible unless all lines in a cable are controlled by a single device.

Traditional DSL competition models rely upon loop unbundling where competitive providers deploy their own equipment to compete to offer broadband access services. Having multiple devices feeding the same cable binder is incompatible with DSM level 3 techniques because it eliminates the possibility of managing the medium to the degree required to successfully benefit from DSM3.

The authors are therefore of the opinion that the introduction of sub-loop unbundling in the access network will significantly constrain the future deployment of DSM. Sub-loop unbundling in fact takes away the possibility of coordinating the entire binder, as the individual pairs within the binder will be fragmented over different operators. Because of this, continued sub-loop unbundling would be a major impediment to the future deployment of DSM technology and its promise of greatly increased capacity.

Towards dynamic spectrum management

Telephone cables typically contain many individual 'pairs' grouped together in 'binders'. The arrangement of binders within cables can be pictured as 'smaller cables within a bigger cable'. In Australia, binders generally contain ten pairs and a separate pair is required for each broadband service.

The limiting factor in xDSL communications is crosstalk interference coming from other lines in the same binder which degrades the signal received by a 'victim' receiver[2]. Interference degrades the Signal To Noise Ratio ("SNR") thereby reducing the capacity.

Recent advances in signal processing technologies have created a possibility to significantly reduce the impairment due to crosstalk. By centrally and simultaneously processing the DSL signals of all the lines sharing the same cable binder, crosstalk can be avoided (or reduced) by either adapting the transmit Power Spectral Density ("PSD") or by pre- or post-processing the signals over multiple lines. These techniques are collectively called Dynamic Spectrum Management ("DSM"). They are classified according to the amount of coordination needed between different lines as either level 0, 1, 2 or 3.

It needs to be strongly noted that if the DSL signals dedicated to different lines in the same binder are not jointly processed, i.e. processed using one line card or DSLAM, the achievable SNR improvements will be limited and crosstalk noise will have similar impact to noise degradation from other sources.

DSM level 0

DSM level 0 is also called static spectrum management ("SSM"); the transmit PSDs cannot exceed the spectral masks as defined in the standard. Examples are margin adaptive ("MA") and

---

[1] 'Binder' is defined below.

[2] The inter-binder crosstalk (i.e. the cross-talk occurring between different binder groups within a cable) is much lower than the intra-binder (i.e. the cross-talk between individual pairs within a binder). It remains to be seen (and will very much depend on the cable type) if inter-binder cross-talk will result in similar problems to that of intra-binder cross-talk.

rate adaptive bit loading. In the MA case, all available power is used to maximise the noise margin. This increases the SNR on this line, but also increases the crosstalk noise on the other lines in the binder.

DSM level 1

DSM level 1 is referred to as autonomous power allocation. It consists of avoiding unnecessary crosstalk to neighbour lines, but without an exchange of information between lines. This is achieved by arbitrarily specifying a target throughput based upon calculated performance expectations and then having each DSL line individually limit its PSD to meet, but not exceed, this performance.

It has been shown that in the case of Central Office ("CO", i.e. 'exchange') fed ADSL scenarios, DSM level one based on iterative water filling ("IWF") does not significantly improve the reach of any of the tiers. Additional information on performance may be found in Alcatel Technology White Paper, "Dynamic Spectrum Management for Digital Subscriber Lines - Edition 2", June 2005, J. Verlinden, T. Bostoen, G.Ysebaert.

DSM level 2

DSM level 2 coordinates the (multiuser) power allocation over multiple lines. In contrast to level 1, level 2 also takes into account the line condition and service requirements of other lines in addition to its own line conditions. DSM level 2 by nature requires more cooperation between lines, and therefore can only practically be implemented in a centralized fashion. An example is optimal spectrum balancing ("OSB"), also known as optimal spectrum management ("OSM").

DSM level 2 yields substantially better performance than DSM level 1 as more coordination and information exchange is allowed between the active lines. In a mixed CO/RT deployment, OSB (DSM level 2) shapes the transmit spectra more intelligently than its counterpart IWF (DSM level 1).

The effectiveness of DSM level 2 in simultaneously protecting CO and RT services sharing the same binder is nevertheless limited compared with the gains of DSM level 3 applied from a single controller and further discussed in Alcatel Technology White Paper, "Dynamic Spectrum Management for Digital Subscriber Lines - Edition 2", June 2005, J. Verlinden, T. Bostoen, G.Ysebaert.

DSM Level 3: vectoring and cancellation

As previously mentioned, improving the SNR could be done by mitigating the self-crosstalk at the transmitter and the receiver for the downstream and the upstream respectively. In both cases, some analogue or digital signal processing is required and this is called vectoring. This vectoring does not optimise the transmit PSDs as with DSM level 1 and 2, but rather compensates the present self-crosstalk while transmitting at full power. In a way, DSM levels 1 and 2 reallocate spectra, reduce the overall power levels and come at marginal cost, while DSM level 3 increases the signal and processing power and will require a significant engineering development.

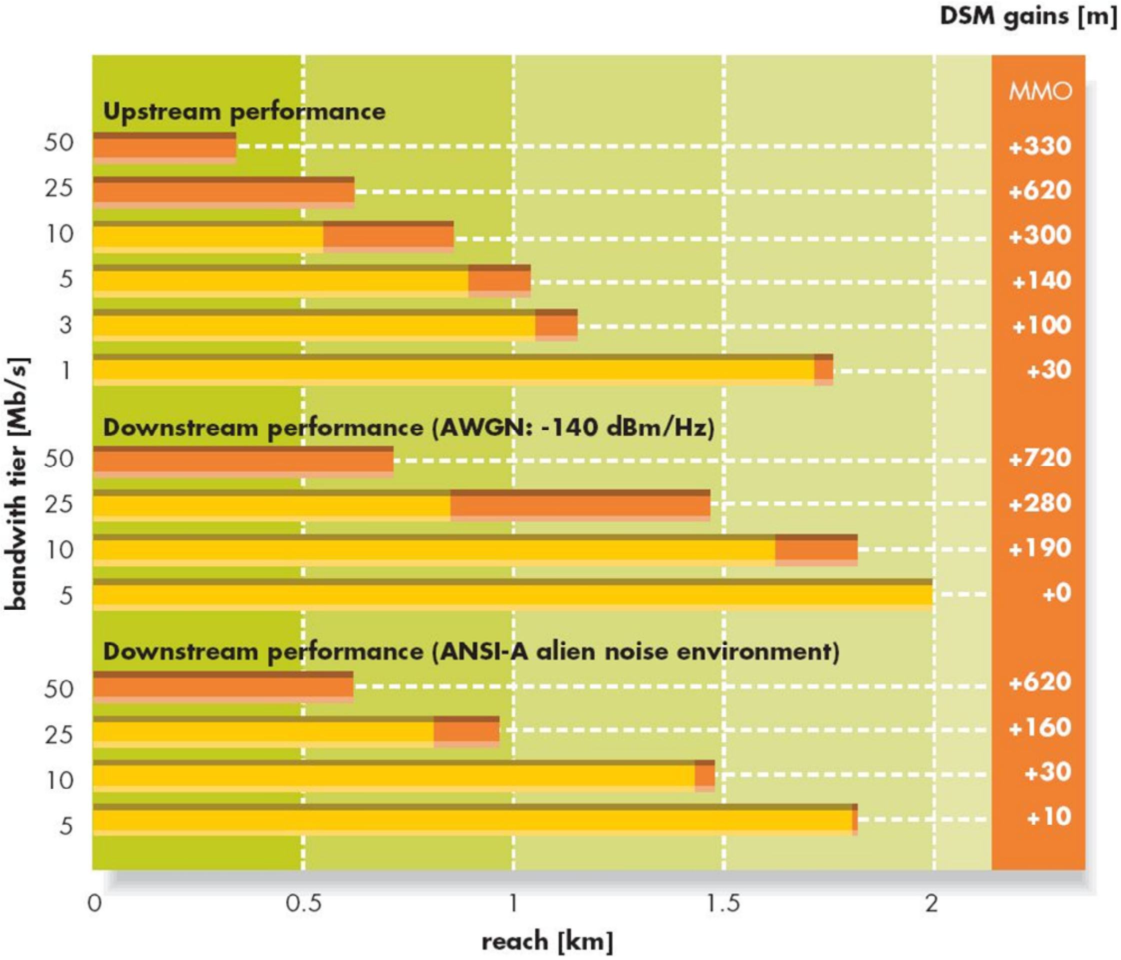

Figure 1 DSM level 3 performance gains (orange) for like-distance VDSL lines, all lines in a binder are part of a DSM3 MIMO noise cancellation system[3]

For residential customers with one twisted pair loop arriving at the premises, the signal processing has to be inside the DSLAM for both upstream and downstream, as this is the only location where joint processing at the symbol level is possible.

Upstream - Cancellation

In the upstream direction, the interference cancellation is performed over a set of strong interferers and no feedback is required from the customer premises equipment ("CPE") : the DSLAM cancels the crosstalk on each line which allows joint decoding of the data arriving on each of the lines. This cancellation requires an estimate of the crosstalk channel.

Downstream - Precoding

In the downstream direction, assuming that the characteristics of the crosstalk channels have been well estimated, one can predict and therefore pre-compensate the crosstalk of each line. However some feedback from the CPE is needed to estimate the crosstalk channel and therefore

[3] Chart referenced from Alcatel Technology White Paper, "Dynamic Spectrum Management for Digital Subscriber Lines - Edition 2", June 2005, J. Verlinden, T. Bostoen, G.Ysebaert.

results in a larger standardisation effort and a dependence on deployment of CPE that implements the required functionality. This CPE feedback can be used as an error signal to track the cross-channel needed for crosstalk precompensation.

Applicability and added value

Vectoring is mainly efficient in increasing the rate/reach of short loop VDSL systems deployed from the RT, such as Fibre To The Node systems. For example, at a loop length of around 800m (a typical 'long' loop in an Australian FTTN deployment scenario), the achievable bit rate can be extended by 40% from around 35 Mbps up to around 50 Mbps by removing 9 out of 10 crosstalkers for both 8d and 17a VDSL2 profiles. At 1 km, more than 25% rate increase is obtained.

It is also interesting to notice that on short (resp. medium) lines, very high bit rates can be obtained by using full vectoring on 17 MHz (resp. 8 MHz). On long loops the vectoring gain is relatively small because the loop attenuation causes most of the far-end crosstalk to drop below the receiver noise floor. As is indicated in Figure 2, we always take into account the presence of at least one crosstalker to allow for new lines coming up. The crosstalk from these lines will not necessarily be cancelled right away, so a provision has to be made for their presence.

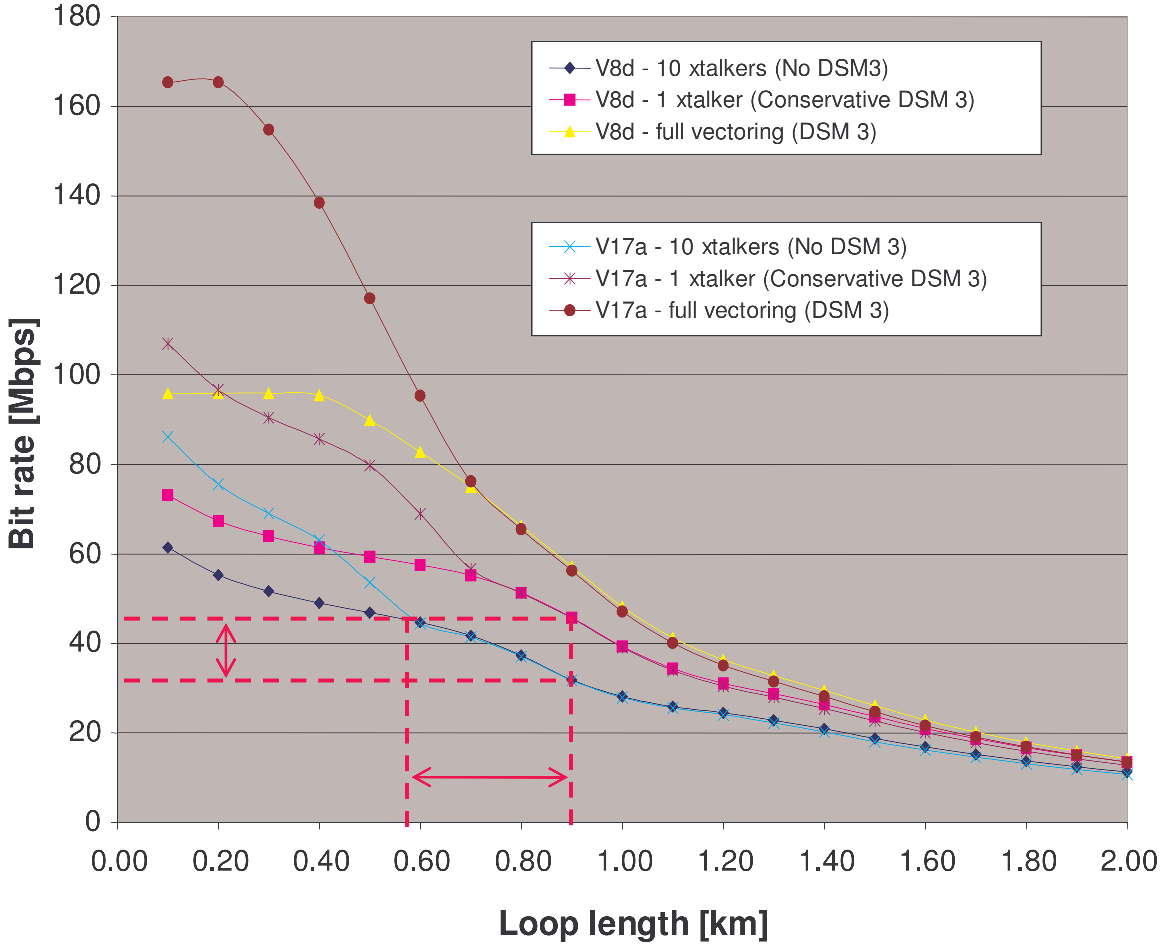

Figure 2 VDSL2 with and without DSM level 3 shows considerable gain

On the other hand, simulations show that vectoring when applied to ADSL2+ systems does not provide any noticeable gain over the alternative, deploying VDSL2, and therefore is not suited for such systems.

It is noteworthy that DSM level 2 and level 3 could be combined to improve the robustness to crosstalk. For example, in RT/CO deployment and assuming that we implement vectoring only on a subset of tones, one may apply DSM level 2 on tones on which no interference cancellation is implemented (e.g. below 2.2 MHz) and boost on the tones subject to vectoring (e.g. above 2.2 MHz).

DSM and Local Loop Unbundling

DSM levels 0 and 1 are, in principle, independent of binder fragmentation across operators. The gains are however very limited. DSM level 2, in the form of PSD-shaping to protect CO-fed lines from RT deployments, is already being put in place now, and many regulators in fact mandate some form of shaping to avoid unwarranted interference and disputes about the quality and noise levels.

The promising DSM level 3 gain comes as a cost in terms of complexity. It requires an in-depth knowledge of the access network topology and the crosstalk couplings to be set and controlled efficiently. Part of the DSM approach will need to be implemented in a distributed fashion (on for example the line-cards), but without a centralised control centre (for example when services originate from more than one DSLAM) the gains will be limited [4].

The implications for loop unbundling are clear: in order to remain viable, DSM level 3 should not be deployed in an environment where more than one DSLAM serves a binder. Depending upon the nature of a market's current broadband competition model, this may have important implications.

Whereas with DSM level 0, 1 and 2 individual loops can be unbundled (either in a shared or full way) at the physical layer, the implementation of DSM level 3 creates a de facto shared medium (i.e. the spectrum within the binder). Within this shared medium, the DSM level 3 operates as a kind of Medium Access Control ("MAC") layer, which is only effective when all lines obey the same rules. Indeed, when an access provider invests in Fibre To The Node systems relying upon DSM3 gains to meet service performance guarantees, the guarantees will be jeopardised by the presence of lines in the same binder served from an alternate DSLAM.

All lines within a so-called vectoring group need to be synchronized at the symbol level from a central clock on the line-card or in the DSLAM. The additional required digital signal processing results in a non-negligible processing and communication overhead. As an example, assuming only ten crosstalkers are cancelled for all tones and all users, approximately 40%-120% additional floating point operations per line are needed. The requisite bus speed for the communication between the chipsets within a line board can be estimated as 48 lines x 1000 tones x 4000 symbols/sec x 16 bits = 3 Gbps. The crosstalk channel coefficients that are intrinsic to precoding and cancellation technology also need to be stored and shared. These requirements immediately show that it is not realistic to expect cross-DSLAM DSM level 3 to ever be possible in a cost-effective way. Even across line-cards, the requirements are at this time not considered realistic: grooming of the line pairs is expected to be needed to make sure all strongly coupled pairs are fed from the same line-card to allow for precoding to be worthwhile.

---

[4] See also "AT&T and nSpired Design, Analysis of loop reach increase due to ordered FEXT cancellation", May 2007.

Table 2 in AT&T's contribution shows that when there are a number of lines that are not coordinated (in a so-called set S2), for example due to physical layer (sub-loop) unbundling, performance gains are very limited. However with full coordination (implying no sub-loop unbundling), much higher gains can be expected for DSM level 3.

The combination of the central control, the required knowledge of the access topology and the impossibility to jointly process lines not originating on the same line-card indicates that the unbundling for DSM 3 will only be possible logically, not physically (i.e. above the physical layer). This is also the case for example in coax cable access loops, where it is termed bitstream, packet or wholesale unbundling.

Standardisation and timelines

Because of the required feedback from the CPE, standardisation is required to make DSM level 3 and more in particular precoding a reality. Within standardisation, effort and studies have been going on for close to two years and have recently culminated in the decision of the ITU to start a new standard, G.vector. It will build on the existing VDSL2 standard but add the requirements for vectoring to be possible. It is however expected that the process will still take more than one year before a final agreement is reached.

Furthermore, there are still challenges that are not being addressed properly in the different standardisation vectoring proposals such as the transient behaviour and the stability. What happens if a new line is added? How will this affect the stability? What is the exact overhead in terms of capacity loss before the new line is added to the vectoring system? Stability is a key issue that has to be tackled properly if any of the vectoring technologies is deployed. Backwards compatibility with legacy CPE is another open topic.

The system integrators, although some are very active in the G.vector process, very much depend on the chipset vendors to offer appropriate chipsets before the design of new line-cards and access multiplexers can start. As the expectation is that some vendors will have prototype systems by the end of 2007, samples and especially volumes are not expected before the end of 2008, partly due to the standardisation effort still required. As such, we do not expect access multiplexers implementing DSM level 3 before 2010 at the earliest.

About the authors

Jan Verlinden, Ir., is a member of the Alcatel-Lucent DSL eXPerts R&D team which specializes in the DSL PHY layer. He has worked on all aspects of DSL technology with a recent focus on all levels of DSM. He represents Alcatel-Lucent Bell in ANSI standardisation activities.

Michael Peeters, Jochen Maes and Mamoun Guenach form the Research and Innovation core of the Alcatel-Lucent DSL eXPerts team within the CTO organisation. Holding PhDs in Applied Physics, Physics and Telecommunications respectively, their research covers the entire spectrum of bundle optimisation technologies, including DSM level 3, both theoretically and experimentally. As such, they are actively involved in the G.vector standardisation process.

Selected references

- G. Carty, Broadband Networking, chapter 4, MacGraw-Hill/Osborne, 2002.
- P. France (ed.), Local Access Network Technologies, chapter 7, IET, 2004.
- Developments in local loop unbundling, OECD report, September 2003.
- Spectrum Management for Loop Transmission Systems – Issue 2., American National Standard for Telecommunication, T1.417-2003.
- Dynamic Spectrum Management Technical Report, NIPP-NAI contribution 2007-038R3. http://contributions.atis.org/UPLOAD/NIPP/NAI/NIPP-NAI-2007-038R3.doc

- T. Bostoen, J. Verlinden, E. Van den Bogaert and R. Suciu, Alcatel Technology White Paper : Dynamic Spectrum Management in Practice – Edition 1, July 2004.
- J. Verlinden, T. Bostoen, G. Ysebaert, Alcatel Technology White Paper : Dynamic Spectrum Management for Digital Subscriber Lines – Edition 2, June 2005
- AT&T and nSpired Design, Analysis of loop reach increase due to ordered FEXT cancellation, May 2007, http://contributions.atis.org/UPLOAD/NIPP/NAI/NIPP-NAI-2007-053.doc