D B2.1 Static and Dynamic Spectral Management Analysis - Issue 2

Author 1:	Rob Kirkby (Editor)	BT	rob.kirkby@bt.com
Author 2:	Ignacio Grande Olalla	TID	igo22@tid.es
Author 3:	Manuel Sanchez Yanguela	TID	msy@tid.es
Author 4:	John Cook	BT	john.w2.cook@bt.com
Author 5:	Andrew Wallace	BT	andy.d.wallace@bt.com
Author 6:	Les Humphrey	BT	Les.humphrey@bt.com

BT = British Telecommunications plc

TID = Telefónica Investigación y Desarrollo Sociedad Anónima Unipersonal

Identifier:	Deliverable D B2.1
Class:	Report
Version:	2
Version Date:	06/01/2006
Distribution:	Public
Responsible Partner:	BT, with contributions from TID
Filename:	WPB2_0009_V06.0_DB2.1-i2.doc

DOCUMENT INFORMATION

Project ref. No.	IST-6thFP-507295
Project acronym	MUSE
Project full title	Multi-Service Access Everywhere
Security (distribution level)	Public
Contractual delivery date	M15
Actual delivery date	Issue 1: 22/06/2005 Issue 2: 06/01/2006
Deliverable number	D B2.1
Deliverable name	Static and Dynamic Spectral Management Analysis
Type	Report
Status & version	V2.0
Number of pages	74
WP / TF contributing	B2/TF2
WP / TF responsible	B2
Main contributors	BT, TID
Editor(s)	Rob Kirkby
EU Project Officer	Pertti Jauhiainen
Keywords	DSL, ADSL, DSM, DLCM, SSM
Abstract (for dissemination)	This deliverable considers improvements to broadband delivery by means of dynamic spectrum management and dynamic line code management. Results from simulations, laboratory experiments and field experience are presented.

DOCUMENT HISTORY

Version	Date	Comments and actions	Status
1.0	22/06/2005	Final version	Final
2.0	06/01/2006	Issue 2	Final

TABLE OF CONTENTS

LIST OF FIGURES AND TABLES

ABBREVIATIONS

AAL1	ATM Adaptation Layer - Type 1
ADLU-64R	All Digital Loop, 64 Upstream carriers
ADSL	Asymmetric Digital Subscriber Loop
AOC	ADSL Overhead Control Channel
ATM	Asynchronous Transfer Mode
ATU-C	ADSL Transceiver Unit – Central Office
ATU-R	ADSL Transceiver Unit – Remote
BER	Bit Error Ratio
bps	bits per second
BT	British Telecommunications plc
CES	Circuit Emulation Service
CO	Central Office
	Synonyms include “exchange” and “telephone exchange”
DLCM	Dynamic Line Code Management
DLM	Dynamic Line Management this includes DLCM and DSM
DMT	Discrete Multi-Tone
DSL	Digital Subscriber Loop
DSLAM	DSL Access Multiplexer
DSM	Dynamic Spectrum Management
	NB in MUSE ‘DSM’ does not include DLCM. DSM only concerns changing transmit signal spectra co-operatively . In some literature elsewhere ‘DSM’ is used as MUSE uses DLM.
DSPBO	Downstream Power Back Off
EOC	Embedded Operations Channel
EU	European Union
FDD	Frequency Division Duplex
FEC	Forward Error Correction
FEXT	Far-End Cross-Talk
FP	Framework Project
FTW	Forschungszentrum Telekommunikation Wien
	(the Austrian telecommunications university in Vienna)

IP	Internet Protocol
ISDN	Integrated Services Digital Network
IST	Information Society Technologies
ITU-T	International Telecommunications Union, Telecommunications Sector
IWF	Iterative Water Filling
LAN	Local Area Network
LLU	Local Loop Unbundling
MUF	Maximum Usable Frequency
MUSE	Multi-Service Access Everywhere
N-QAM	Quadrature Amplitude Modulation, N-point constellation
NRIA	normalised-rate iterative algorithm
POTS	Plain Old Telephony Service
PSD	Power Spectral Density
QAM	Quadrature Amplitude Modulation
REIN	Repetitive Electrical Impulse Noise
RS	Reed Solomon
SCM	Single Carrier Modulation
SNR	Signal-to-Noise Ratio
SSM	Static Spectrum Management
TCP	Transmission Control Protocol
TDEQ	Time Domain Equalizer
TF	Task Force
TID	Telefónica Investigación y Desarrollo Sociedad Anónima Unipersonal
TVoADSL	Television over ADSL
VDSL	Very High Speed DSL
WP	Work Project
xDSL	any type of DSL

REFERENCES

[1] Anders Uvliden, C.J.Pitt, and Vladimir Oksman, "TF2.2 – Position paper on Dynamic Spectrum Management (DSM) and Dynamic Line-Code Management (DLCM)" TF2_0008_v00_7_DSM position paper_February 2005.doc

[2] J.M. Cioffi, W. Yu, S.T. Chung, G. Ginis "Iterative Water-filling for DSM" Contribution T1E1.4/2001-200R5, Greensboro, NC, November 2001

[3] G. Ungerboeck, "Performance of SCM and DMT in the presence of impulse noise" Contribution, SG15/Q4 FI-81, Fiji Islands, February 2000

[4] D. Statovci and T. Nordström, "Adaptive subcarrier allocation, power control, and power allocation for multi-user FDD-DMT systems", Proc. ICC2004, Paris, France, Jun 2004, vol. 1, pp. 11-16.

[5] “Specification of the Access Network Frequency Plan applicable to transmission systems used on the BT Access Network”
Oftel Technical Requirement OTR004:2002 Issue 2
available at URL http://www.oftel.gov.uk/ind_groups/nicc/Public/anfp_2.pdf

EXECUTIVE SUMMARY

This paper considers improvements to broadband delivery by means of Dynamic Spectrum Management (DSM) and Dynamic Line code Management (DLCM). We use the nomenclature of Dynamic Line Management (DLM) developed in MUSE deliverable DTF 2.1 in order to clearly differentiate between the two elements (DSM and DLCM) grouped together in the North American Spectrum Management under the name of DSM.

Results from simulations, laboratory experiments and field experience are presented. From these we conclude:

- There is still considerable scope for further innovation in Static Spectrum Management (SSM) methods. They are likely to remain central to regulation for the foreseeable future.
- Embedding voice gives a capacity increase, which may be significant when considering universal service.
- A single voice channel is small enough that it can be guaranteed capacity; however this requires modems be engineered to protect this traffic stream despite service data-rate uncertainty. Moreover, there are many higher layer issues to be addressed regarding how service management interacts with the modem.
- DLCM is highly desirable to maintain a line's serviceability.
- DSM offers the prospect of capacity (when needed) with a higher limit than can be guaranteed (that means sometimes the need will be disappointed). Some services would need reengineering in higher layers to be able to exploit capacity as it appears and survive its loss as it disappears. Currently only data transfer is resilient.
- In an established static spectrum management regime, as is typical in Europe, the potential capacity increase due to DSM appears modest. Concurrent use of spectrum for downstream transmission by cabinet and exchange hosted DSL systems has been previously identified as a key application of DSM. However this requirement can be efficiently implemented by static spectrum management techniques as defined in the UK ANFP[5].
- DSM potentially offers increased capacity, but the capacity attained on any particular line will still depend on its length and the physical network topology. DSM does not enhance universality in any network with lines of a variety of lengths.
- Autonomous goal seeking DSM has immensely complicated implications, and a wide range of options that need a much larger study in order to provide confidence before network implementation could be contemplated.
- Common SSM will continue to be needed even if DSM is adopted. It will become more important for sustainable broadband as usage increases

1 INTRODUCTION

MUSE is aimed at universal provision of broadband in Europe, and accepts that initially this will mostly be delivered to peoples' houses over the existing telephone network, using DSL technology. Current deployments leave considerable scope for service improvements in data rate, reach, and reliability.

The main impairments to DSL performance are attenuation and crosstalk[1].

For practical purposes loop attenuation is fixed[2]. Crosstalk however may be managed by controlling how much transmit power is used by modems, and how this power is disposed in the frequency domain : this is 'spectrum management', and it is the keystone for building sustainable broadband services.

This document addresses different management optimisations aimed at such improvements, loosely within the scope of static and dynamic spectrum management. SSM and DSM are related to minimising the limitation of DSL performance due to mutual interference effects through crosstalk. Dynamic Line Code Management is also considered (DLCM). This aimed at minimising the impact of extrinsic noise.

NB In MUSE the areas covered by 'DSM' and 'DLCM' do not overlap, and together they form 'DLM'. In some other literature the term 'DSM' is taken to include DLCM.

The Potential Enhancements

An earlier position paper [1] outlines the potential enhancements, and hopes for a "substantial increase in capacity". Briefly, both DSM and DLCM are adaptive modulation methods aimed at exploiting the channel better by responding to changes in it:

DSM	proposes to adapt the transmitted signal's spectrum, refraining from transmission of unnecessary signal power in order to benefit the other systems sharing a cable by reducing their crosstalk.
DLCM	proposes to adapt the error control coding of the system in response to the line's actual noise, perhaps without changing its neighbours' noise environments.

There are two fundamental problems underlying DSM:

1. even if the crosstalk characteristics the cable plant were to be completely known, there is no objective way to decide how various services and customer groups should be prioritised
2. the crosstalk characteristics of the cable plant are not completely known, and it would be impractical to determine them completely

---

[1] induced environmental electrical noise and interference from radio signals are also significant on some lines; longer lines are particularly susceptible

[2] The usual exception being fault repairs.

The treatment of various customer groups and their service needs is a central issue of spectrum management and regulation. It is often driven by opposing inspirational and pragmatic public policy trends, and in general is served by ad-hoc compromises. These compromises are all temporary to some extent, but existing deployments require protection. In practice the ongoing mass deployment of ADSL services will increasingly serve to constrain the scope for changes in spectrum management practices.

Consequently it is already too late to search for ideal or general solutions. Rather we look for pragmatic approaches to mine benefits for one customer set from the unused potential on other lines – it is always clear that service priorities may change, for individual customers, for service providers, or the regulator.

The starting point for spectrum management is the body of self imposed deployment rules developed before the impact of local loop unbundling. Historically these rules were system based, rather than spectrum based. Static spectrum management (SSM) methods are a natural development to enable new systems to be deployed by analogy with legacy approaches.

Given that complete a-priori knowledge of the plant is out of reach, automatic behaviour is key to optimising the service set. In this sense static spectrum management need not be truly static. There are two obvious optimisations that can happen within an SSM environment: service bit rate can be maximised to use all the legal power on a particular line, or the transmit power can be reduced to use the minimum necessary to meet a service need. A third possibility keeps the data rate to the minimum necessary for the customer, and then maximises power to minimise susceptibility to extrinsic noise and maximise service stability.

The service provider must assume that the maximum power may used by all interferers[3], and so cannot rely on the benefits of power reduction in setting service rate commitments.

Two key features separate DSM from SSM. Firstly, user data rate is actively managed, and secondly, the benefits of reduced transmit power are relied upon to deliver services that could not be provided if the limits of an SSM rule set were to be arbitrarily exploited.

Services to be Carried

Of course the real interest is to carry everything, but to focus the mind MUSE concentrates on the Triple Play services:

voice

data

video

We note that each has different needs and has different abilities to use the adventitious capacity increases offered by adaptive modulation.

Voice uses a constant bit rate, needs low delay, can survive momentary noise bursts, cannot use extra capacity if offered, and will just stop working if its capacity is reduced. This is as typically implemented at present – one can imagine a voice coding method which degrades gracefully in the event of reduced capacity.

---

[3] Even those lines not currently using DSL but which may in future do so

Data transfer is usually taken to mean file transfer activity. While faster is better slower is tolerable. Real-time delay is not an issue (the quality metric is how quickly the transfer is completed). While errors in the delivered file are not tolerable, the transfer mechanism will detect line noise and recover by causing retransmission.

Video is transmitted in a highly compressed form and cannot tolerate error bursts in the delivered data. In principle the capacity needed varies with the dynamics of the video content. A video transmission system can be imagined where varying line capacity is exploited as varying picture quality – for adaptive modulation methods this would have to include indications as to which data may be discarded if the capacity decreases suddenly. Video which is for passive watching can tolerate large delays but cannot tolerate variations in delay – typically achieved by maintaining a large buffer which both smoothes out the dynamic capacity needs and provides time for error recovery retransmission[4]. Interactive video, such as gaming[5], needs small delays.

The DSL system will need its traffic to be appropriately labelled so it can treat each service differently; and, for video, to treat the vital and optional contents differently.

One obvious method to do this is to supply the traffic in ATM virtual circuits; this would be a device for communicating with the modems, for efficiency the data on the line would not be in literal ATM cells. For example to handle voice we would expect cells at a higher rate than the content needs, to keep latency low; hence the cells will be mostly empty. In a LAN this low efficiency would be tolerable. On the line we would only expect real content to be carried, and we would expect the modem to preserve this capacity when the line fluctuates.

Currently ATM sends idle cells when there is no traffic; for DSM this would be undesirable, as we should like the line signal to be substantially reduced when idle, so we must either ask the modems to regenerate the idle traffic without literally sending it to line or change the ATM protocol to remove the idle cells.

We note that if triple play is achieved then the line no longer needs splitter filters : there is no voice and the modems can use the whole bandwidth.

2 ENVIRONMENT

The Givens

Aspects of DSL we do not expect to change include:

- The cables.

Reusing the existing telephony cabling is the economic rationale of DSL, so it will not be changed on a large scale to provide better performance.

---

[4] In the limit one might download a video programme as a file and play it locally at some future time.

[5] Throughout this paper 'gaming' means playing games, particularly role-playing action/adventure games with or against other human players

- Multiple operators.

Local Loop Unbundling (LLU) will continue to be enforced and drive spectrum management policy. Getting agreement of the competitors for radical changes to the spectrum management regime will be hard to achieve, and it may be that the only acceptable compromise is co-existence of DSM with more conservative approaches, and strict limitations on adaptive behaviour to enable backwards compatibility.

- Static spectrum management.

Conventional spectrum management is still needed to support mass legacy deployments. Static spectrum management advances and exploitation of dynamic spectrum management will therefore be constrained.

- Universal service

The driver for a genuinely universal service (at a uniform price) is political rather than economic : the technical measures must include a range of methods, to provide the cheapest workable solution for any given customer. Should the ‘universal’ data rate be set relatively high, this implies relatively expensive options for longer lines, e.g. building new remote network nodes.

Real World Noise

Classical communication engineering usually assumes all noise is additive white Gaussian noise (AWGN), and that this assumption is safe because it is conservative[6]. However real world DSL experience has enriched that significantly.

At the time ADSL was being engineered it was known that the noise is in general not white (and neither is the telephone line channel). So ‘stationary’ noise became the norm, and that’s what DMT is so good against. At this time everything else was assumed to be ‘impulsive noise’, and was not well characterised. The literature was sparse [Mandelbrot], and ADSL modems were engineered to be able to cope with short noise bursts provided they were short enough and sufficiently well spaced apart.

Much later we found mains originated impulsive noise is[7] widespread, and REIN entered the literature []. This is a new area of study and there is still no agreement on how to apply a standard REIN test. Its effect on current generation modems is significant : they interpret all noise as if it were some level of stationary Gaussian noise; they report ‘margin’ in terms of it; and they equalize assuming it (including the bit loading in DMT). So (i) in a REIN environment they are unduly harmed and (ii) different implementations interpret (and suffer) the same levels of noise very differently.

---

[6] it is known that Gaussian noise is the most damaging to theoretical information capacity at any set power level.

[7] Or possibly ‘has become’; its possible that earlier workers did not notice this noise because it was not then common.

In summary, for DSL at present we can categorise noise sources as

- crosstalk
- stationary other
- REIN
- nonstationary other

'Stationary other' includes such slowly varying effects as diurnal variation of RFI; its stationary enough that an equalizer can maintain adaptation to it. 'Nonstationary other' includes those things as yet undiscovered, and virtual noise [see section 0] is a response to it. Note that we have so far assumed crosstalk is stationary; with DSM this assumption may become false, since under DSM systems vary their output power at the same kind of response rate as they perform equalisation.

Regulation & Public Policy

When LLU was mandated in the EU, spectrum management became a matter of public policy. It must be imposed in each country with the authority of the regulator, because it is a limitation imposed on the LLU operators. Regulators typically consult all the interested parties before imposing a particular set of rules, and typically prefer the technical development to be by an industry committee, but the authority lies with the regulator.

So far spectrum management has meant *static* spectrum management, typically discussed in terms of spectral limits (masks) imposed on each transmitter. This is considered technology neutral – it does indeed exclude technologies which cannot meet the limits, but it can admit technologies as yet undeveloped.

However with DSM there may be decisions to be taken on which *services* to allow over others, or even which particular lines may have a particular service. It is not clear how one may get a politically acceptable answer to these decisions. However it is clear that the authority for the process which produces these answers still lies with the regulator.

3 SSM

Static Spectrum Management is control of the noise environment in a cable by limiting the spectral power the signal on each pair into it. The limits are usually prescribed by masks much like those in equipment standards. Access points at different locations can have different masks, but usually all the line ends at one location have the same mask irrespective of who is using them and for what.

The principle objective is that each system has a predictable performance because all of its neighbours are limited. A desirable second feature is that the rules are simple enough that each system operator can understand what limits apply to him.

An example of regulation using SSM is the UK access network frequency plan [5], which has been in operation for some years.

In this section we will take as given the simplest situation, where each shared cable has only one 'head' end (exchange or cabinet but not both), where all the cables in an administration are subject to the same rules, and where all pairs in a cable have the same rules. The following subsections then explore the cases where there are multiple 'head' ends, where the spectral masks may be changed on selected cables, and where different pairs in the same cable have different spectral allocations.

Multiple DSLAM Locations

The problem of line attenuation on DSL performance can be partially addressed by hybrid fibre/copper systems – FTTC and FTTK. However the spectrum management problem space expands considerably when DSLAMs are installed at multiple locations, but with shared spectrum.

In this section we consider full bandwidth ADSL2+ installed at the exchange and VDSL installed at the cabinet. We conclude that static spectrum management can yield valuable improvements in the data rate which can be offered for all users.

This section is substantially the rationale for the current developments in the UK's[8] spectrum management regime.

3.1.1 Background

The use of ADSL2+ from the Exchange or Central Office raises interesting spectral compatibility issues with higher rate systems deployed at Remote Terminals or Cabinet flexibility points that may share the same spectrum (e.g. VDSL or ADSL2+).

The question of to what extent one affects the information capacity of the other depends on the installed base of systems and the location of the remote equipment that may share the same spectrum on adjacent wire pairs within the same binder. This will vary with operator and access network.

BT has carried out an extensive set of capacity simulations to explore the technical options for deploying ADSL2+ at the central office, while avoiding severe degradation of capacity of DSL systems deployed at the cabinet.

Simulations have shown that protection of VDSL2 capacity at the cabinet requires fully shared downstream spectrum from 138kHz to 2.2MHz. This implies full overlap of ADSL, ADSL2 and ADSL2+ downstream spectrum with VDSL2 operating in a downstream band that begins at 138kHz and extends as far as the highest ADSL2+ frequency (2.2MHz).

3.1.2 The Near/Far Shared Spectrum Problem

To date[9] making VDSL at the cabinet "compatible" with ADSL at the exchange has largely consisted in arranging that the cabinet based VDSL simply doesn't use frequencies below about 1.1 MHz. However deployment of ADSL2+ at the central office makes this strategy obsolete. The corresponding strategy of denying the cabinet based VDSL access to frequencies below 2.2 MHz simply has too drastic adverse effect on its capabilities.

If ADSL2+ is permitted at the CO with no downstream power controls (other than the maximum permissible levels according to the ITU-T ADSL2+ standard (G.992.5)), then one way of delivering maximum network capacity from the CO and Cabinet is to use electrical distance related power controls at the Cabinet at which VDSL2 or ADSL2+ may be in use.

---

[8] Strictly, spectrum management in BT's access network. For historical reasons the city of Hull has different rules.

[9] In the UK

Goal

The goal is to define a method that will be acceptable in a regulated environment. Where cabinet deployment is undeveloped, and legacy deployments dominate, the ideal scheme must meet a 'fairness' criterion heavily biased to protection of exchange systems. In the UK it was anticipated that ADSL2+ systems would be deployed from the exchange before significant deployment VDSL from the cabinet. Consequently the criterion of fairness led to a requirement that below 2.2 MHz cabinet based DSL systems are not permitted to exceed a power level equivalent to passing ADSL2+ signals in adjacent pairs arriving from the exchange. However, fairness also permits VDSL to fully exploit frequencies above the maximum than can be exploited by ADSL2+ from the exchange even if this frequency is below 2.2 MHz.

Cabinet close to CO

If the Cabinet location is close to the CO then little or no restrictive power control is required in the region 138kHz to 2.2MHz to prevent the Cabinet systems causing network harm to the CO systems and vice versa.

Cabinet distant from CO

If the Cabinet is a long distance from the CO then no or only minimal power cut-back would be required in the range 552kHz to 2.2MHz. This is because ADSL and ADSL2+ frequencies would not have any usable capacity from the CO and the remote ADSL2+ or VDSL2 would have little or no effect on CO based systems.

Problem region

The capacity problem is the middle region where the Cabinet systems affect CO based systems sharing the same downstream bands. Power control would need to be a function of the CO to RT loop loss, and each Cabinet or RT in theory would need different power control settings. The desire would be to make the crosstalk effect of the RT as no worse than if more ADSL2+ systems were deployed from the CO.

3.1.3 Solving the Near/Far Shared Spectrum Problem

The figure below shows the family of downstream PSD shapes that BT has been experimenting with in simulations. It shows the range of power control required in the region 138kHz to 2.2MHz which has been shown to substantially increase network capacity using Static Spectrum Management.

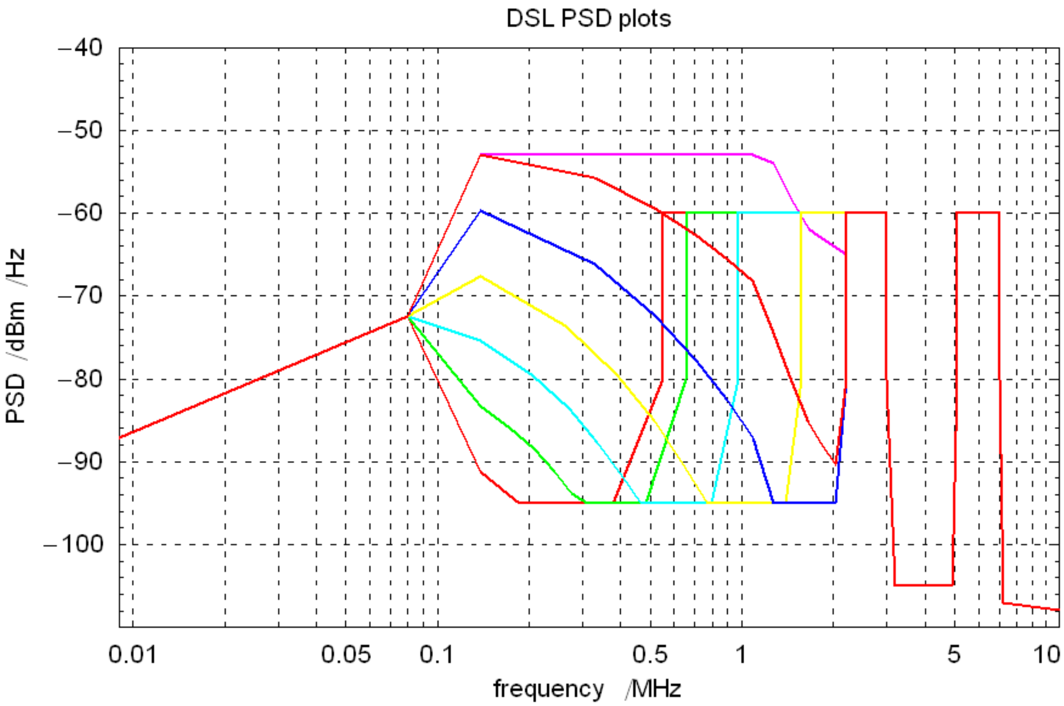

Figure 0.1 Allowed Downstream PSD at 10 dB steps of E-side cable loss to 60 dB

The plots are for illustrative purposes only and the lower limit for shaping in the ADSL2+ band has been assumed to be –95dBm/Hz in accordance with G.992.5.

The particular power profile used would be a function of the loop loss from the CO to Cabinet. In the region above 2.2MHz the PSD aligns with VDSL Plan 997 (the more symmetric plan), rather than Plan 998. BT now feels that Plan 997 offers advantages on longer loops from the cabinet, particularly for upstream.

This downstream power backoff approach is an extension to that captured in BT contribution to the ITU-T SG15/Q4 - "VDSL Requirements" SS-058, section 4.6.2 "Downstream Power Backoff".

3.1.4 An Algorithmic Downstream Power Backoff Approach

An approach to implementing the variable PSD mentioned above is to implement Downstream Power Backoff (DSPBO) as a function of the electrical length L of cables (E-side) from the exchange to the cabinet, but limited to frequencies below a notional 'maximum usable frequency' (MUF).

The function DSPBO(f) would be

- a simplified cable model H(f,L) for f<MUF
- 0 dB for f>MUF

L is a metric of electrical length of the E-side cables to the cabinet, provided by network management.

3.1.5 A simple cable model

A three parameter model has been found to be satisfactory for H(f,L):

$$H(f,L) = \left(a + b\sqrt{f} + c\,f\right)L$$

with
f in Mhz
L in km
H in dB

Using this model is has been found possible to satisfactorily track 0.5mm and 0.4 mm BT cable over the range of frequencies of interest, i.e. 138 kHz to 2.2 MHz, with one set of parameters calibrated for equal loss at 300 kHz.

The graphs below illustrate the degree of fit which has been achieved using this model.

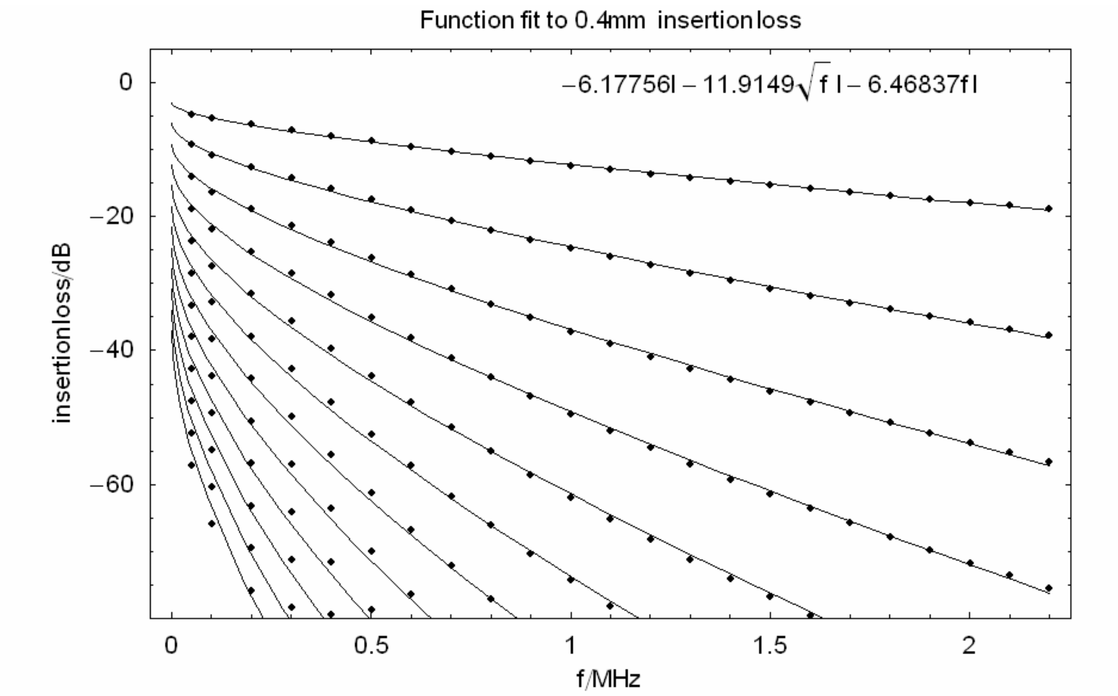

Figure 0.2 Function fit to measured insertion loss on 0.4mm cable

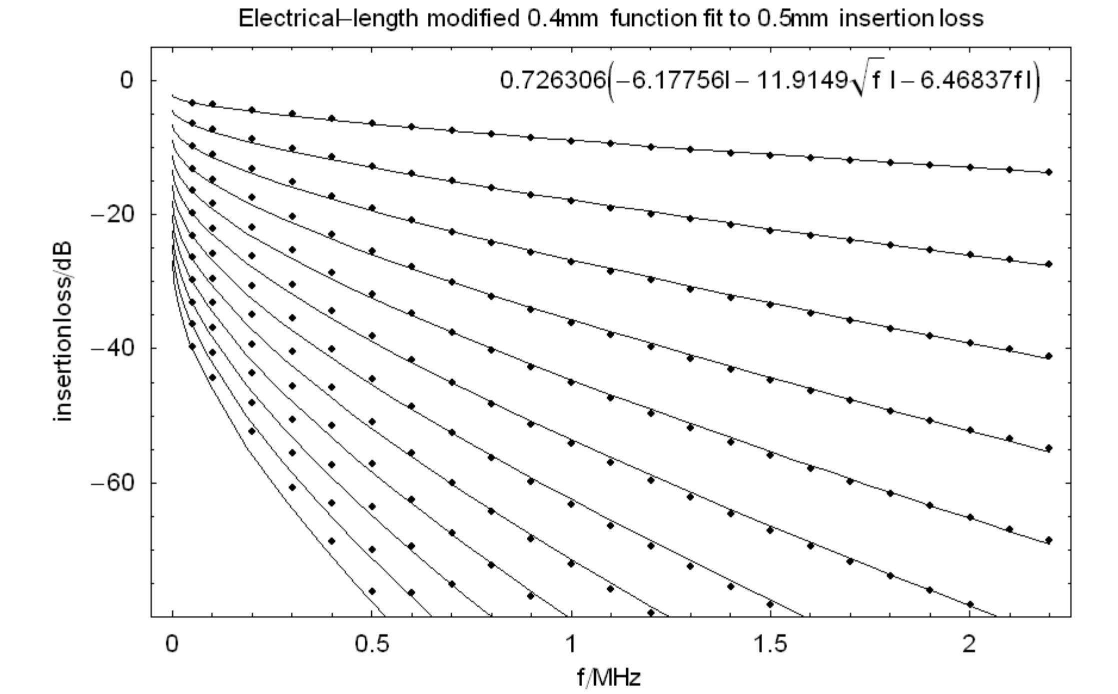

Figure 0.3 Accuracy of 0.4 mm cable model scaled for 0.5 mm cable

NB in the graphs above the 'l' was in each case the length of a measured cable sample. In practical use in a network the value 'L' used is a notional length (perhaps equivalent to 0.4mm cable) and includes the scaling factor. From the above one would expect cables of any mixture of 0.4 and 0.5mm sections to have a meaningful 'L' value, at which H(f,L) tracks the real cable's insertion loss well, and the ratio ((actual cable length) / L) to be somewhere in the range 1 (all 0.4mm) to 0.726 (all 0.5mm).

> Warning: this model is fitted to the frequency band 138 kHz to 2.2 MHz, the ADSL2+ downstream band. The f term in the model is unusual, and was included to get good fit at the lower frequencies. Do not use this model for bands much above 2.2 MHz (VDSL for example) because the fit becomes *very* poor. For example, above 2.2 MHz a better model is $H(f,L)=\left(k+b\left(\sqrt{f}-\sqrt{.3}\right)\right)L$ with k = -14.644 dB and b = -33.254 dB/km/√MHz

3.1.6 The Maximum Usable Frequency - MUF

The maximum usable frequency parameter is a parameter associated with a modulation scheme and a cabinet. In this case, it attempts to estimate an upper frequency bound, above which ADSL2+ modems operating on links from the exchange through this cabinet will not be able transmit any data, even in the absence of crosstalk from cabinet located equipment.

The MUF could be determined from a notional noise floor N, so that the MUF is the lowest frequency for which the transmitted signal power is less than N, except that MUF is never greater than 2.208 MHz for the case of ADSL2+/VDSL interaction.

This parameter is important, in that ADSL2+ equipment located downstream of the cabinet will not be receiving any data on carriers above this frequency. Therefore there is no need to protect the spectrum above this frequency. Indeed at frequencies above the MUF, the VDSL equipment located at the cabinet is allowed to transmit at its nominal power, according to the band plan in use.

3.1.7 Impact of Downstream Power Back-Off on Standards

Currently, ADSL and VDSL equipment does not possess PSD shaping capability sufficient to meet the requirements of such a scheme.

ADSL2/2+ can shape the downstream spectrum to meet differing national requirements, but restrictions in the toolset provided mean that the in-band PSD can be no more than 20 dB below the standard PSD template.

VDSL2's transmit windowing will allow significantly improved D/S PSD shaping, as a result of the more tightly constrained frequency domain 'splash' from the modulation of an individual carrier.

Contributions have been provided to the ITU-T SG15/Q4, ATIS NIPP-NAI and DSL Forum describing the support which would be needed to allow this DSPBO scheme to work.

3.1.8 Simulation Results [funded from outside MUSE]

Four scenarios were simulated. The ADSL2+ capacity from the exchange and the VDSL capacity from the cabinet were tabulated.

"1.1" Non-overlapped spectrum, with transition frequency 1.1 MHz, i.e.
ADSL2+ at exchange, limited to 138 kHz to 1.1 MHz
VDSL at cabinet, operating at -60 dBm/Hz from 1.1 MHz upwards

"2.2" Non-overlapped spectrum, with transition frequency 2.2 MHz, i.e.
ADSL2+ at exchange, operating on its full bandwidth
VDSL at cabinet, operating at -60 dBm/Hz from 2.2 MHz upwards

"O" Overlapped spectrum, no D/S PSD shaping, i.e.
ADSL2+ at exchange, operating on its full bandwidth
VDSL at cabinet, operating at -60 dBm/Hz from 1.1 MHz upwards

"PC" Overlapped spectrum, with D/S PSD shaping at cabinet, i.e.
ADSL2+ at exchange, operating on its full bandwidth
VDSL at cabinet, operating with downstream PSD shaping as described

The charts on the next pages show the ADSL2/ADSL2+ capacity available from the exchange and the VDSL capacity available from the cabinet in the four scenarios.

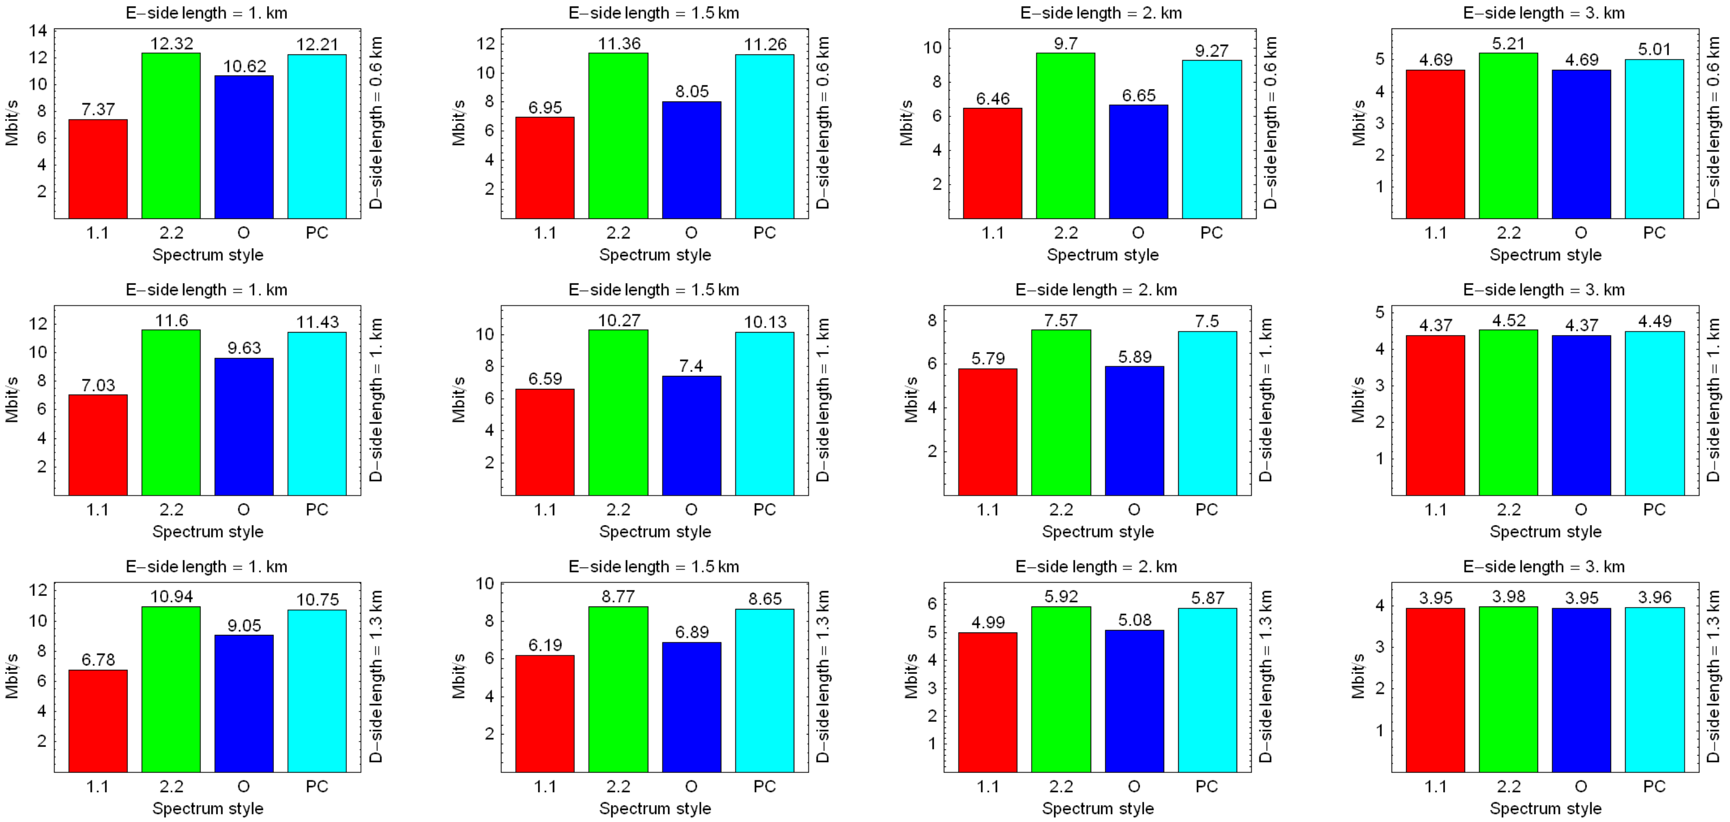

Figure 0.4 ADSL2+ performance with overlapped exchange and remote ADSL2+/VDSL PSDs

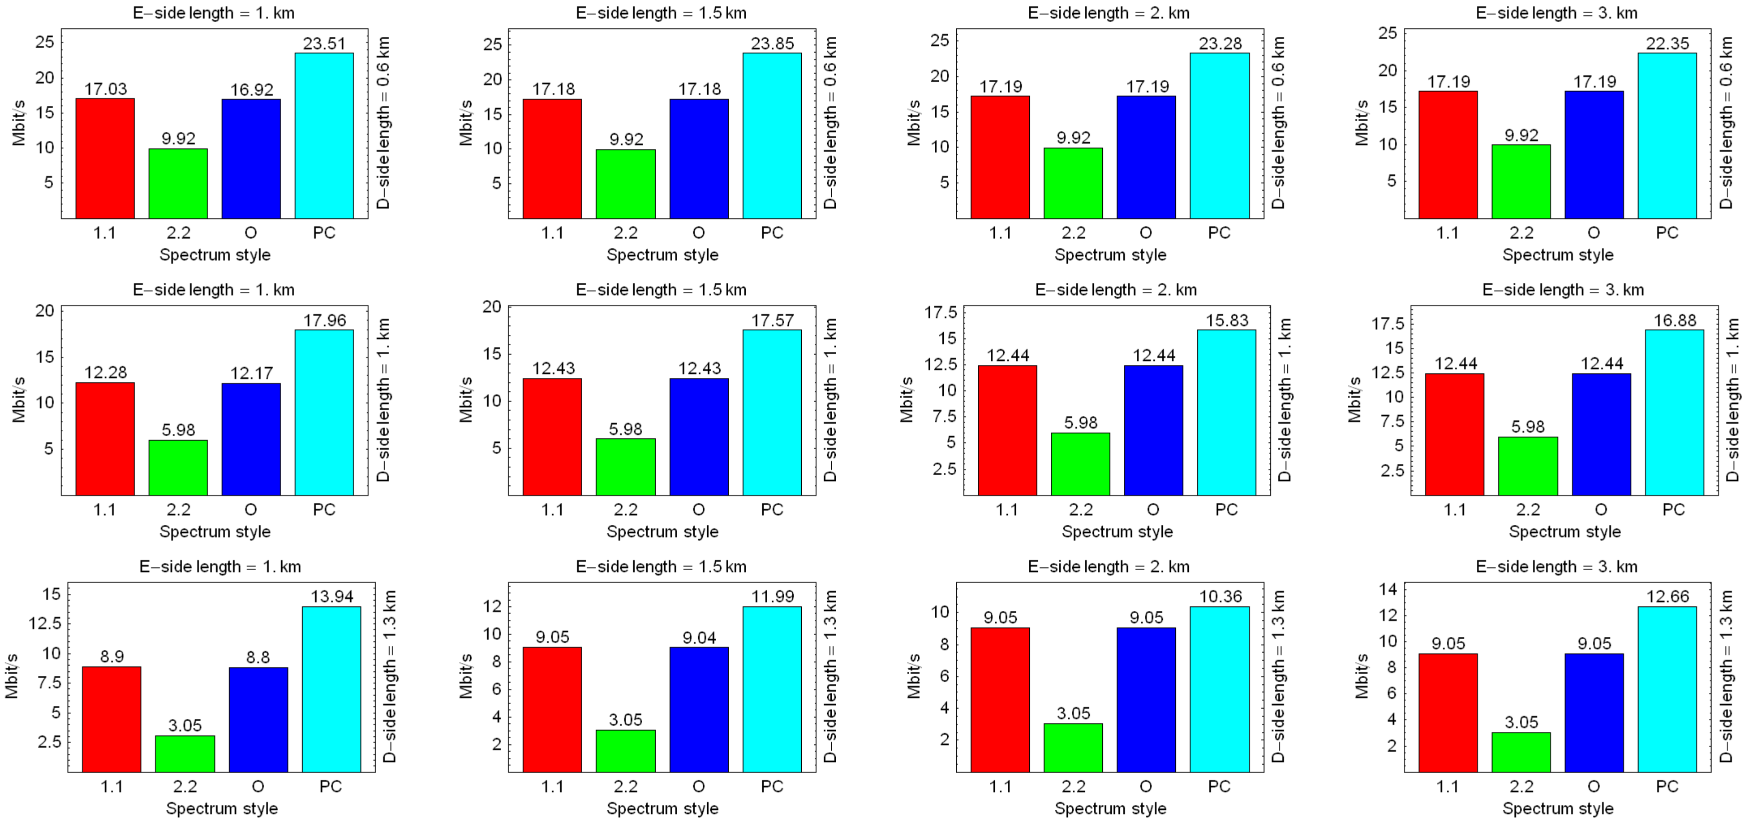

Figure 0.5 VDSL performance with overlapped exchange and remote ADSL2+/VDSL PSDs

The major conclusion to be drawn from these charts is that the “PC” option of shared spectrum with downstream power control is essentially a win-win result.

Service Opportunities Through Engineering SSM

Present ADSL systems typically send more power to line than is necessary for their operation. The DSM problem is how to exploit this excess to best advantage.

This section explores exploitation by static spectrum management, where each modem has a prescribed limit to its transmit spectrum. We intend to study the approach based on static reassignment of bandwidth.

3.1.9 Introduction to ADSL

The ADSL system in the physical layer divides the entire spectrum available in 256 small slots[10]. Each slot is a carrier modulated using N-QAM that can transport information at a variable bit rate (but fixed in the training phase).

In the training phase, when the channel is being analysed and negotiated, the ATU-R and the ATU-C check all the usable slots (the not-reserved ones) to check the capacity of them. This capacity depends on the noise and the attenuation of the line.

The line attenuation depends on the quality of the wire used (material, diameter, insulation, cable construction, etc.) and the distance from ATU-C and ATU-R, and it is higher as the frequency grows.

The noise is introduced by noise of the components in the equipments (typically from the D/A converters[11], not very dependant of the frequency) and external noise. The external noise comes from several sources, but the most important on medium length and short lines is crosstalk, where an affected line is subject to interference from signal energy coupled from systems using companion wires in a binder.

3.1.10 The excess margin and excess capacity problem.

In a noisy line, both maximum length and maximum throughput are reduced as noise increases. If a long line is required, the maximum data rate is limited by the noise present. If a high capacity line is requested, it cannot be very long due to noise.

ADSL modems will try to obtain the maximum raw capacity from the channel, independently of the ATM layer data rate. The modem will fully exploit the bandwidth, and will usually operate at maximum power (sometimes limited power back-off is performed[12]). Therefore most lines exhibit spare capacity or signal-noise ratio margin, while contributing to a high level of crosstalk in the cable.

---

[10] 512 for ADSL2+

[11] typically the A/D converter provides quantisation noise at -140 dBm/Hz,
and thermal noise is insignificant at -175 dBm/Hz

[12] basic ADSL backs off on very short lines to avoid overloading its partner on the line, not for public benefit

The excess capacity/margin is an opportunity. The problem is how to use it to best advantage. SSM and DSM are techniques to manage the process

The main idea of this section is to increase the potential capacity of neighbouring lines by reducing crosstalk. A modem with excess capacity transmits less power so that while its own needs are still met it liberates capacity for others.

3.1.10.1 Simulations

For these simulations, MATLAB will be used, with a simulation software for xDSL systems originally by Telia and then continued by FTW (“Forschungszentrum Telekommunikation Wien” in Austria).

The web page of the xDSL Simulator is http://www.xdsl.ftw.at/xdslsimu/xDSLsimu3.0b1/doc/Simulation_Tool.html

Some modifications have been done to the software to accommodate it to the project requirements.

The exact scenario configuration will depend on what is being tested, but there are some common parts. All the lines will be copper line of 0.4mm in diameter with polyester insulation.

The noise model used is not one of the defined in the ITU specifications, but it is calculated in every simulation with the scenario conditions. This noise is a floor white noise (only applicable if there is only one line in the simulation) and the crosstalk noise produced by the interfering lines.

The varying part in the simulations, as the number of lines or line length, will be specified in every section.

For most of the parts of the section, only downstream (from the network to end users) part will be analysed, for two main reasons:

- The downstream part has a lot of available bandwidth, so the percentage of unused spectrum is considerable.
- The attenuation and noise increases with frequency, so the main advantages of spectrum management will be noticed in downstream.

3.1.10.2 Crosstalk noise impact.

This first simulation is to check that crosstalk noise really is significant. If crosstalk impact were minimal there would be no need to manage it.

The next graph shows the maximum throughput in a 2000m line as the number of ADSL lines increases from just 1 up to 50.

It should be noted that the rate is measured in raw ADSL data rate at physical net ATM layer.

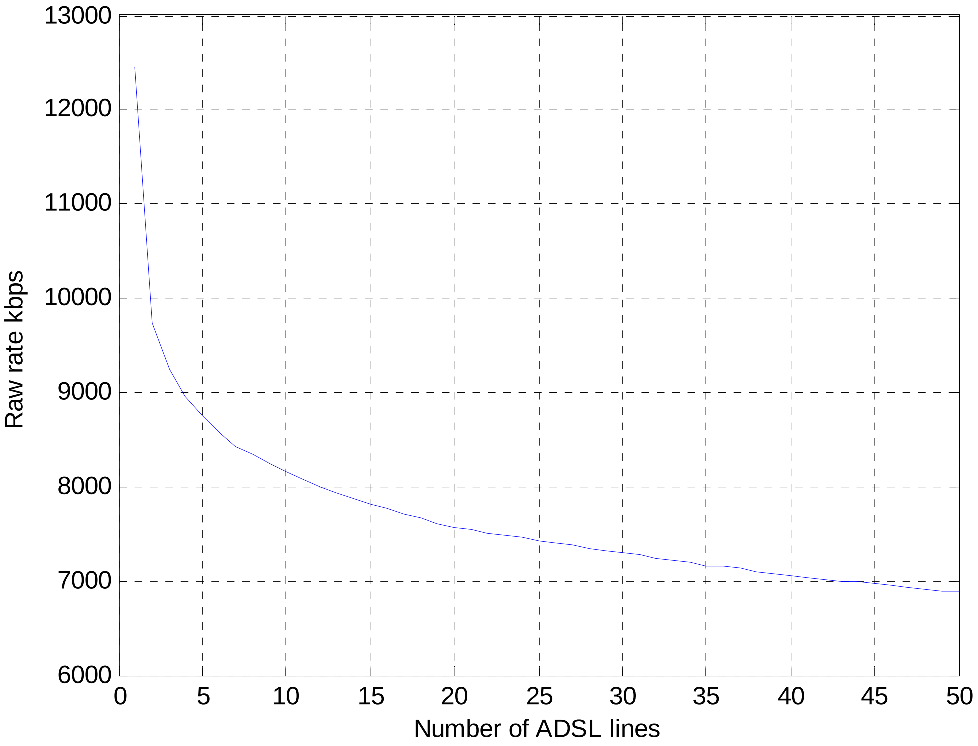

Figure 0.1 Raw rate as no. of lines increases

The impact of the crosstalk noise on the maximum throughput is clear, it being reduced by a factor of two when the number of lines is high.

The loss of maximum throughput is very noticeable when there are a few lines, but less important as the number the lines increase. This can be seen in the decreasing slope of the curve. This means that the important loss is produced when some crosstalk noise is introduced, and less important when the crosstalk noise is further increased.

A first conclusion that can be extracted from this graph is that the crosstalk noise has a very important impact on the signal/noise ratio of the received signal, so SSM/DSM may be interesting.

3.1.10.3 Trying take advantage from SSM

The simulations here will explore the proposition that some lines may be given extra capacity by restricting their neighbours to only part of the bandwidth (but the restricted systems still achieve their own rate needs)

Let us suppose a fixed-length line. For the test, a line of 2 kilometers using a 0.4mm wire will be used.

If the purpose is to increase the maximum rate of all the ADSL lines in the binder, it cannot be done limiting the available spectrum of each line. A noisy slot is better than no slot at all. So the case proposed will be different.

In the binder, there will be a high number of low-rate ADSL lines and a low number of high-rate lines. The low-rate lines will be limited in spectrum so the high-rate lines will have much less noise, increasing the available rate.

The low-rate lines will have a downstream throughput of 512 kbps. As the simulator returns the RAW throughput, 1 Mbps will be used as a conservative value of RAW data needed. This means that the lines will have an wasted data rate that could be removed without any loss of performance.

So, in total, there will be 40 low-rate lines and an increasing number (from 1 to 10) high-rate lines. This increasing number will allow a quick view of how the high-rate lines affects the overall performance.

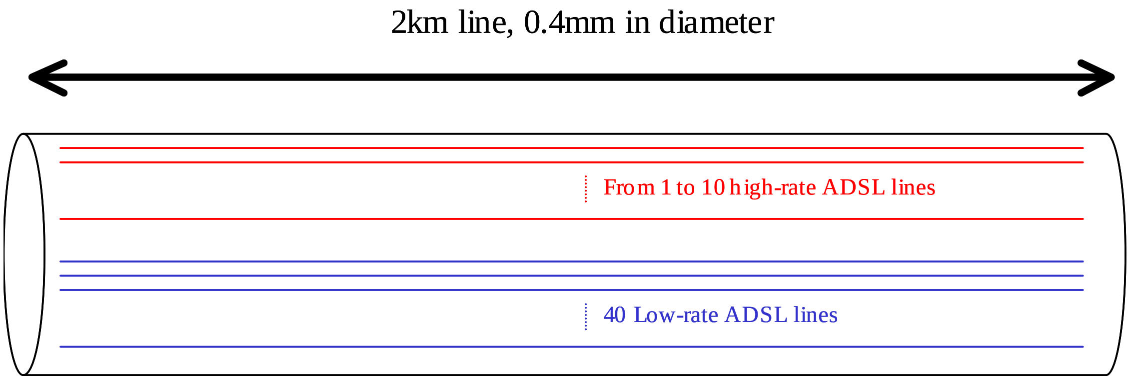

Figure 0.2 Binder configuration

The first step is to check how much bandwidth is necessary for the 40 low-rate lines.
For that, the first simulation contains 50 lines (the worst case with 40 low-rate and 10 high-rate lines) with restricted bandwidth. The restriction will be from 5% of the bandwidth to the 100% of the bandwidth.

The restriction will be introduced in the higher part of the spectrum. So a line with the 50% of the bandwidth restricted can use the first 109 slots in downstream, from the 218 it usually has (from 38 to 255).

The higher part of the spectrum has been chosen because that part is the worst part for crosstalk, as crosstalk increases with frequency. A graph comparing the effects of stripping the higher part or the lower part is showed later in the document.

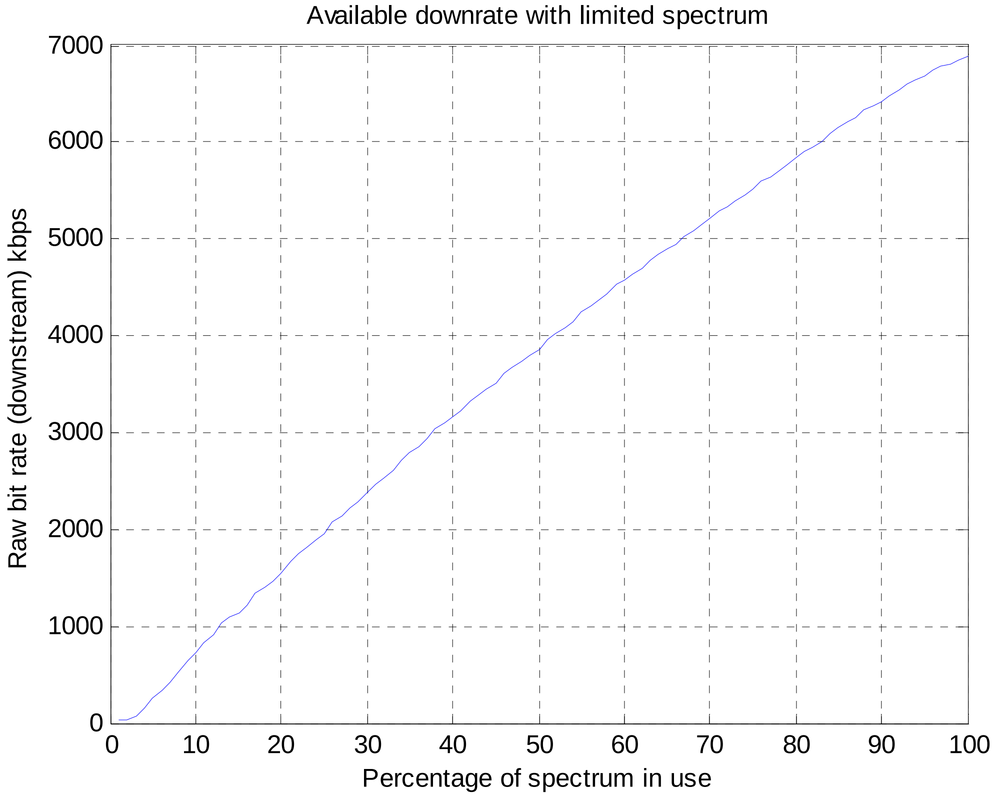

Figure 0.3 Available downrate with limited spectrum

In the graph it can be seen that to obtain a raw rate of 1000kbps, it is needed about 12% of the total spectrum, about 26 slots, so this value will be used for the low-rate ADSL lines.

Once the minimum required bandwidth for the low data-rate lines is known, the next step is to check if the high-speed lines will have a higher rate with the limited speed, so a graph showing the data-rate of the high data-rate lines in both situations (with the low-rate limited and with the low-rate unlimited) would be interesting

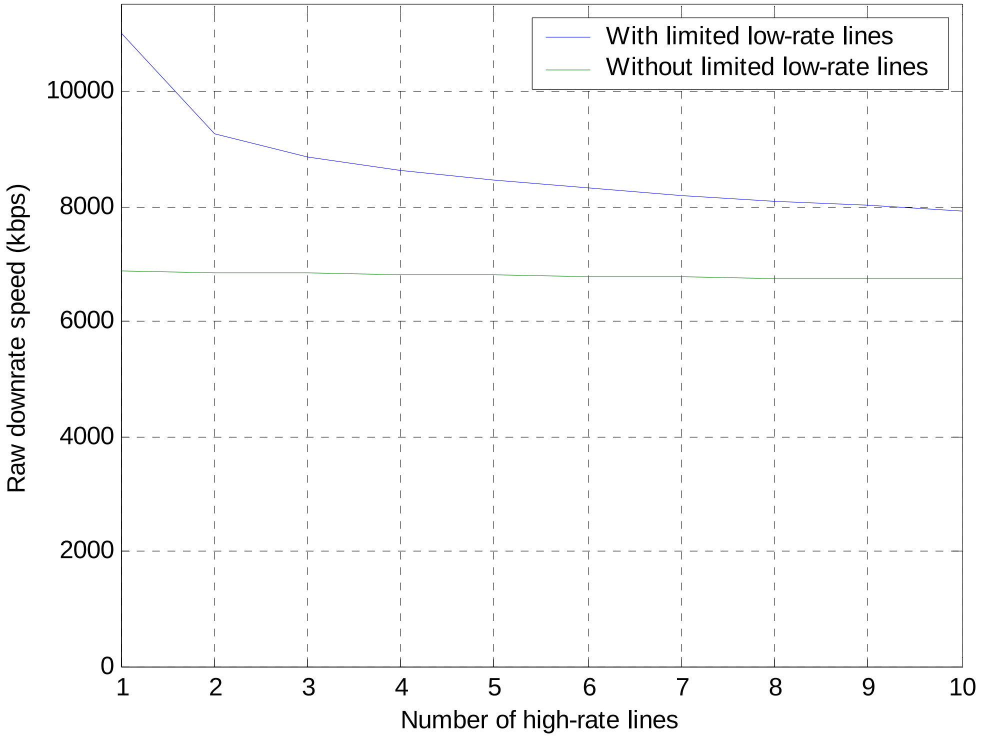

Figure 0.4 Comparison between high-rate lines with and without limits

This graph shows an improvement in the maximum capacity of a line when the others lines are spectrum-limited. This difference is maximum when there is just one high-rate line (the difference is about 4Mbps!) and is progressively reduced as the number of high-rate lines increases.

The previous test has been done using the lower part of the spectrum for the low-rate lines. The next step is to check if using the higher part of the spectrum (thus, the lower part is only for the high-rate lines) can increase the capacity further.

Using the same method, the necessary bandwidth for the low-rate lines can be obtained. In this case, it is necessary about 18% of the spectrum. As higher frequencies have more attenuation and more noise, more spectrum is necessary to obtain the same 1000kbps.

The following graph compares the three situations: with no limitations, limiting in the lower part and limiting on the higher part of the spectrum:

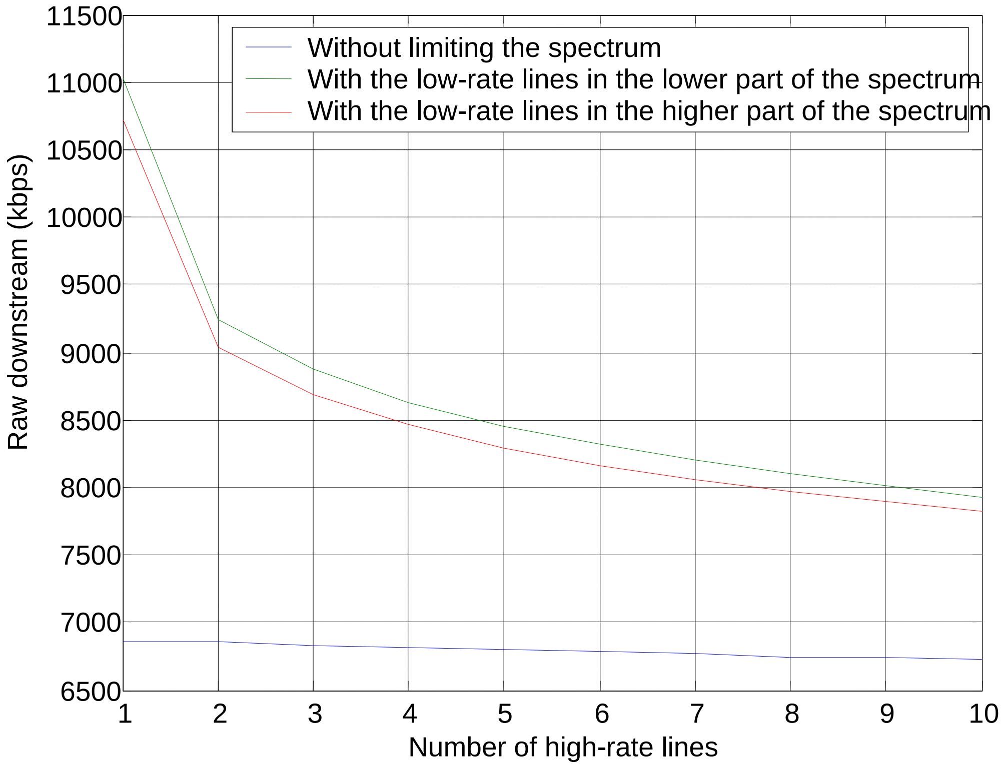

Figure 0.5 Comparison among different limitations on the spectrum

It can be seen that it is better to use the lower part of the spectrum for the low-rate lines, as it produces the maximum throughput.

Another possible situation would be splitting the low-rate lines in two groups. Both will be in the lower part of the spectrum, as it has been seen that it is the better place. One segment of 20 lines followed in the spectrum by another segment of 20 lines.

Using the simulations the values of 12% and 13% of the spectrum for each segment can be obtained.

With these values, the following graph can be generated:

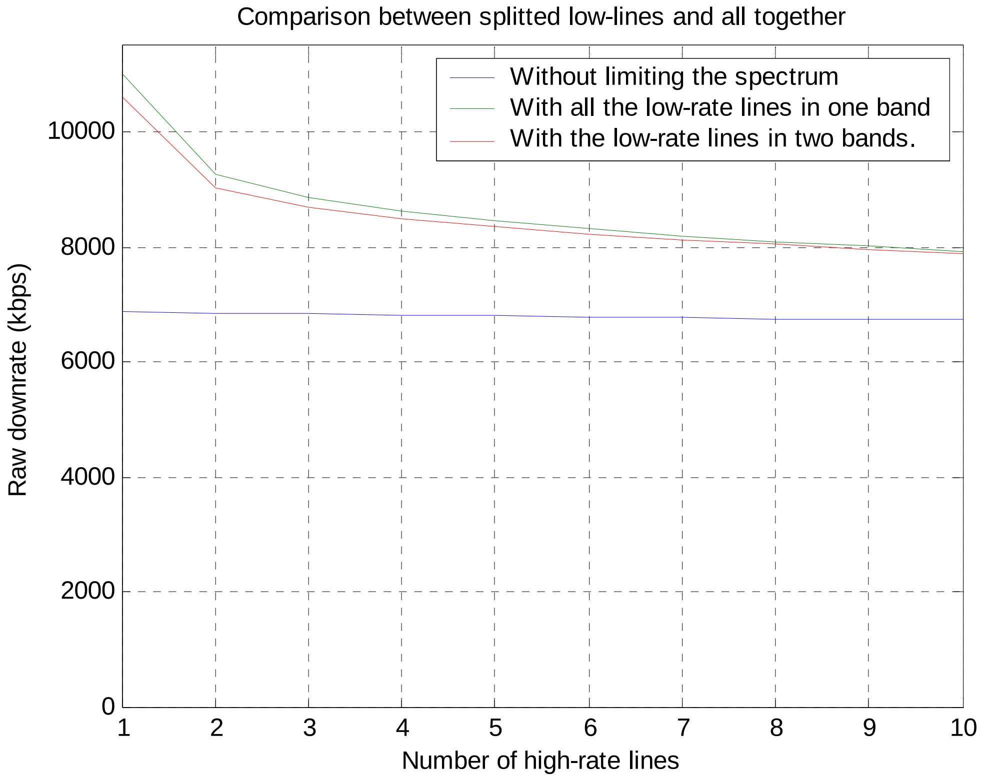

Figure 0.6 Comparison between split low-lines and all together

The results are a little worse for the high-rate lines, so this technique doesn't seem to be very effective. This were a predictable result, as it could be seen in Figure 2.2.1 that the effects of introducing crosstalk were most noticeable when the crosstalk is first introduced than when the crosstalk is further increased.

The main conclusion is that under some circumstances (a few high-rate lines) the use of static spectrum management may improve the maximum capacity of some lines, without degrading the quality of the other lines.

3.1.10.4 Conclusions.

The main conclusion that can be observed is that the noise produced by crosstalk has a big impact in the maximum line length or the maximum throughput. So, a static spectrum management may in some cases improve the quality of some lines, as can be seen in Figure 2.2.4 and Figure 2.3.1 where standard lines and spectrum managed lines are compared.

3.1.11 E1 over ADSL2

In this section we consider the carriage of E1 service over ADSL2 modems and explore what SSM impositions are needed to make it work.

An E1 line over ADSL would require 2244 kbps (for exact calculations see APPENDIX 1 - E1 over ADSL data rate) but symmetrical. The ADSL could provide that speed perfectly in downstream but it is unable to reach 2Mbps in upstream.

The solution proposed here is to use ADSL2 in all-digital mode with an extended upstream channel. The extended upstream allows to use 32 additional sub-carriers (from 32 to 63) and another 6 sub-carriers (from 1 to 6) to reach more than 2 Mbps.

The main problem would be that some sub-carriers are shared in downstream and upstream, as the ADSL2 specification doesn’t change the downstream bands in the extended upstream (ADLU) mode. This will produce a high degree of interference in those sub-carriers making it nearly useless at long distances.

The static spectrum management will allow reducing the available downstream so the overlapping is avoided. The upstream will be increased drastically and the downstream will be little affected.

Another effect that will be studied is the compatibility of this new service (extended upstream and reduced downstream) with the traditional ADSL service

3.1.11.1 ADSL2 with extended upstream and reduced downstream.

The Annex J of the ADSL2 specifications has been used as a reference for the spectral mask of ADSL2. Annex J specifies the extended upstream of an All-digital ADSL2 with spectral compatibility for ISDN.

The first step was to check if the reduction of the downstream was really needed. For that, a simulation with an ADLU-64 standard line and an ADLU-64 with reduced downstream (called from now ADLU-64R) are shown together.

The first scenario represents ten lines in a binder with a varying length from 250 to 4000 meters. In the graph there are two pairs of graphs: the discontinuous line represents the data rate reached by 10 ADLU-64 standard lines. The continuous line shows the data rate of the reduced ADLU-64R lines.

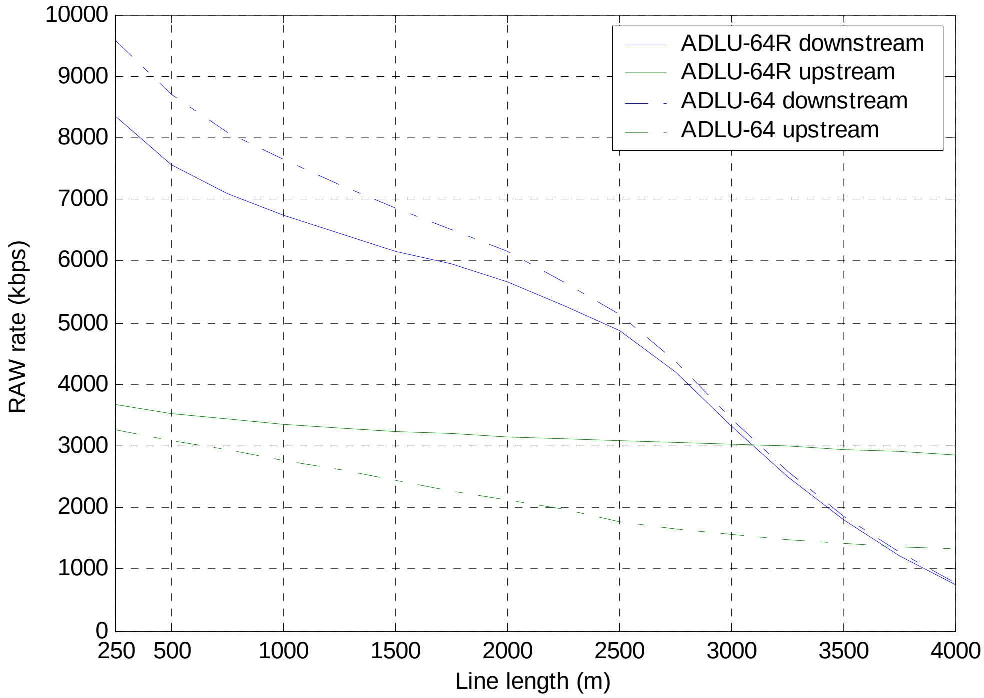

Figure 0.7 Rates for ten ADLU-64 and ten ADLU-64R

The graph shows that to reach 3Mbps in upstream, it is necessary to reduce the downstream spectrum. The discontinuous (standard ADSL2) line quickly reduces the available data rate due to spectrum overlapping with the downstream.

The continuous (reduced-downstream ADSL2) line reaches 3Mbps at long distances of more than 3km, while the downstream is mostly affected at short distances.

So in this first simulation is clear that an overlapping spectrum is not a good choice if an E1 is to be offered over ADSL2.

3.1.11.2 Effects of crosstalk with ADLU-64R

Now another scenario is presented. It has been seen that it is possible to offer E1 over ADSL2 with static spectrum management where there are ten lines in the binder.

The next graph shows the data rate resulting when more than 10 lines are in use. Up to 50 lines are simulated.

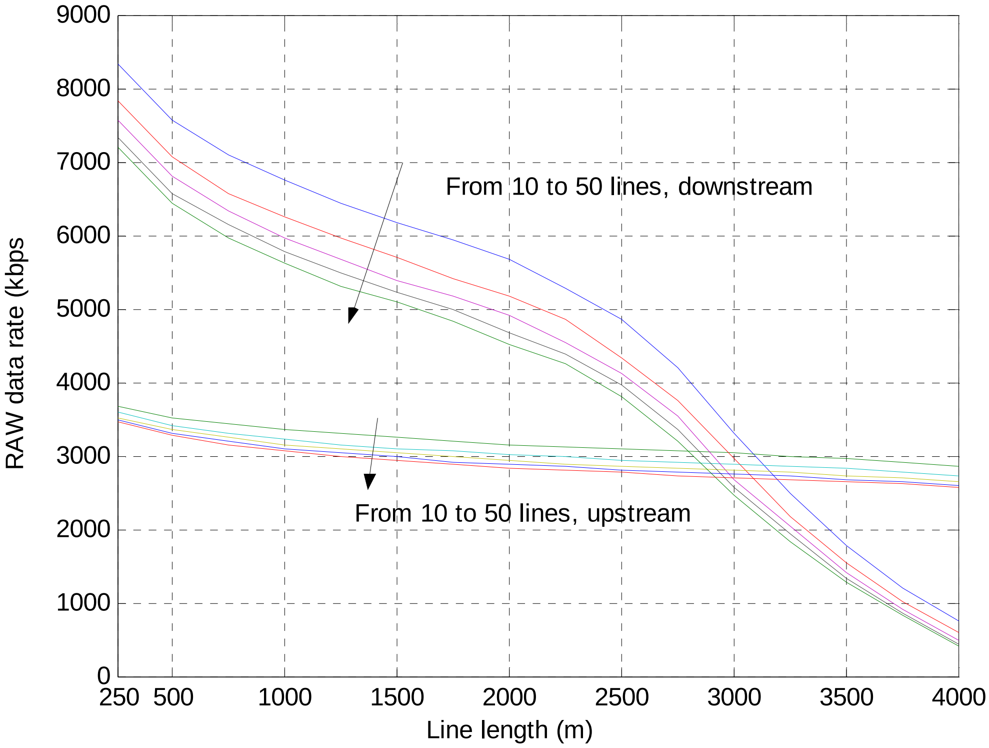

Figure 0.8 ADLU-64R lines with reduced downstream

Downstream uses the upper part of the spectrum, the most affected by crosstalk and attenuation, so the downstream data rate decreases faster than the upstream data rate.

The zone where downstream and upstream lines cross is show in detail in Figure 0.9 and the values where downstream rate and upstream rate are equal are the following:

Number of lines	Line length (m)	Data rate (kbps)
10	3090	3008
20	3020	2880
30	2960	2816
40	2940	2752
50	2920	2688

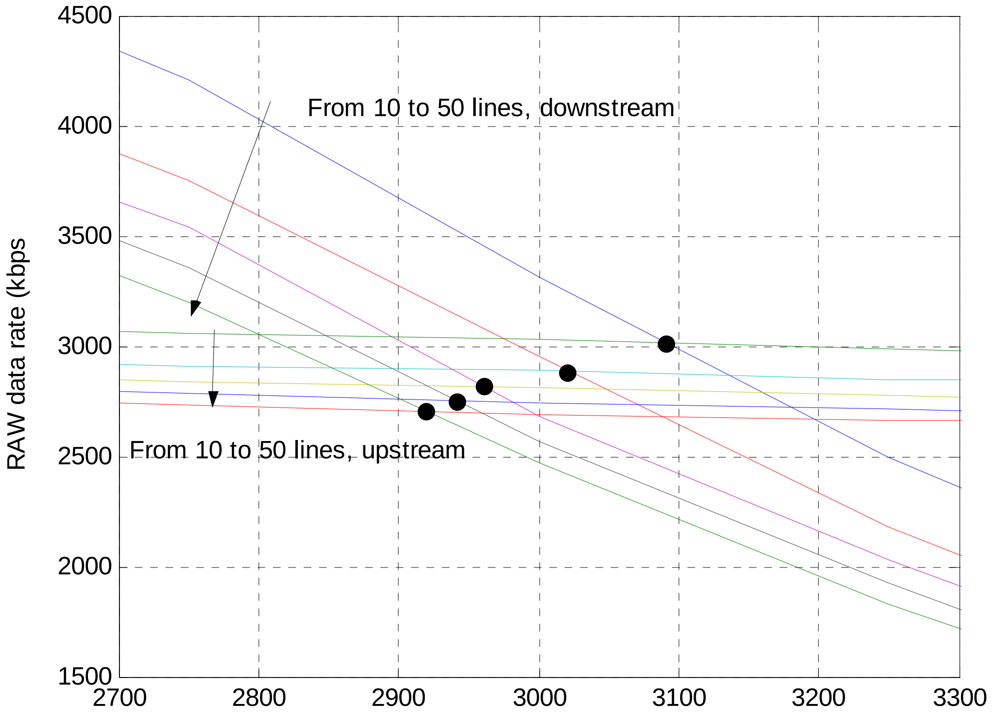

Figure 0.9 Ranges for symmetrical max rates

As the number of lines increases, the upstream data rate is lower, and only with ten lines it is higher than 3Mbps at 3km.

3.1.11.3 Compatibility with standard ADSL lines

An E1 line over ADSL2 will be only required by a limited number of people (or enterprises) so it won't be the most solicited service. It will share the binder with some traditional ADSL lines. It is important to check if the crosstalk from the standard ADSL lines will interfere with the upstream of the ADLU-64R lines.

A worst-case scenario is proposed. 10 ADLU-64R lines with 40 ADSL lines are together in the same binder and the available rate for each class is shown in the graph.

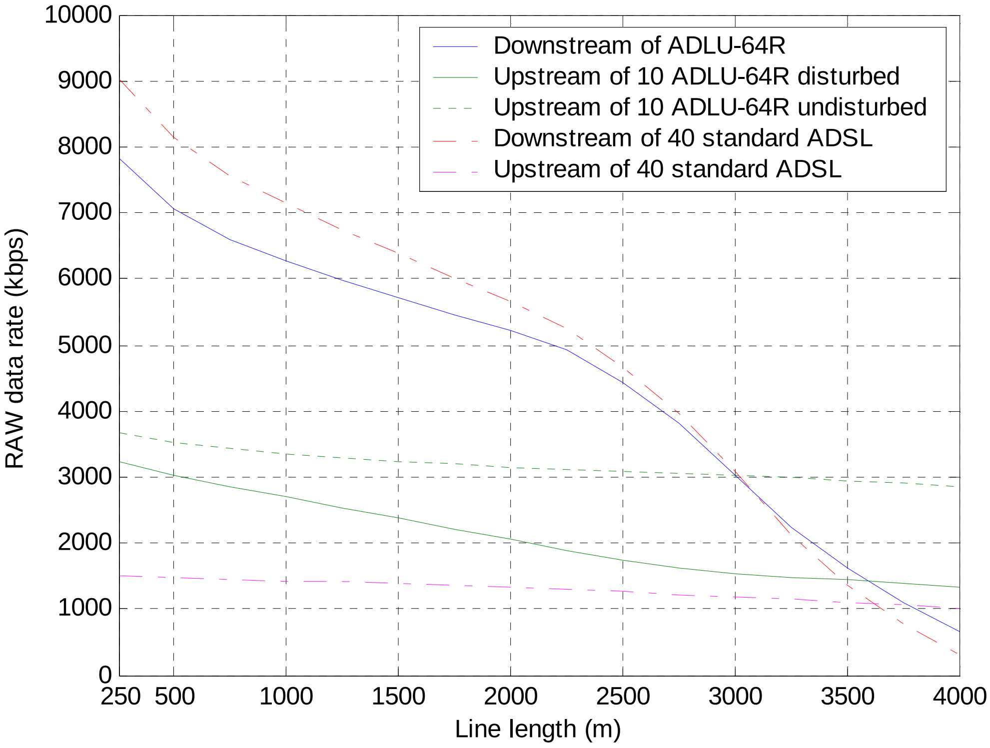

Figure 0.10 Spectral compatibility between ADSL and ADLU-64

The continuous line shows the data rate in downstream and upstream of one of the ten ADLU-64R lines. The “dot-line” graphs correspond to the upstream and downstream of one of the 40 ADSL lines.

There is an additional graph in dotted line. It shows the upstream of the previous graph with 10 lines, that is, the upstream of 10 ADLU-64R lines *alone,* if there were not ADSL lines. It is shown here for comparison.

The graph clearly shows that the crosstalk introduced by the ADSL lines deteriorates the upstream rate of the ADLU-64R and makes it unusable to send an E1 over it.

Reached this point, it looks like an E1 over ADSL2 would be incompatible with ADSL if both were in the same binder. However, as seen in the previous section, it is possible to restrict also the downstream spectrum to avoid crosstalk.

A possible solution would be to reduce the ADSL spectrum in downstream so it begins at sub-carrier 64, avoiding overlapping with ADLU-64R upstream. This solution, by the way, is very unlikely to be acceptable in practise, as it would mean all the ADSL infrastructure is changed just to accommodate some ADSL2 spectrum-reduced lines.

A simulation done with this scenario, where ADSL lines have a spectrum reduction is shown here:

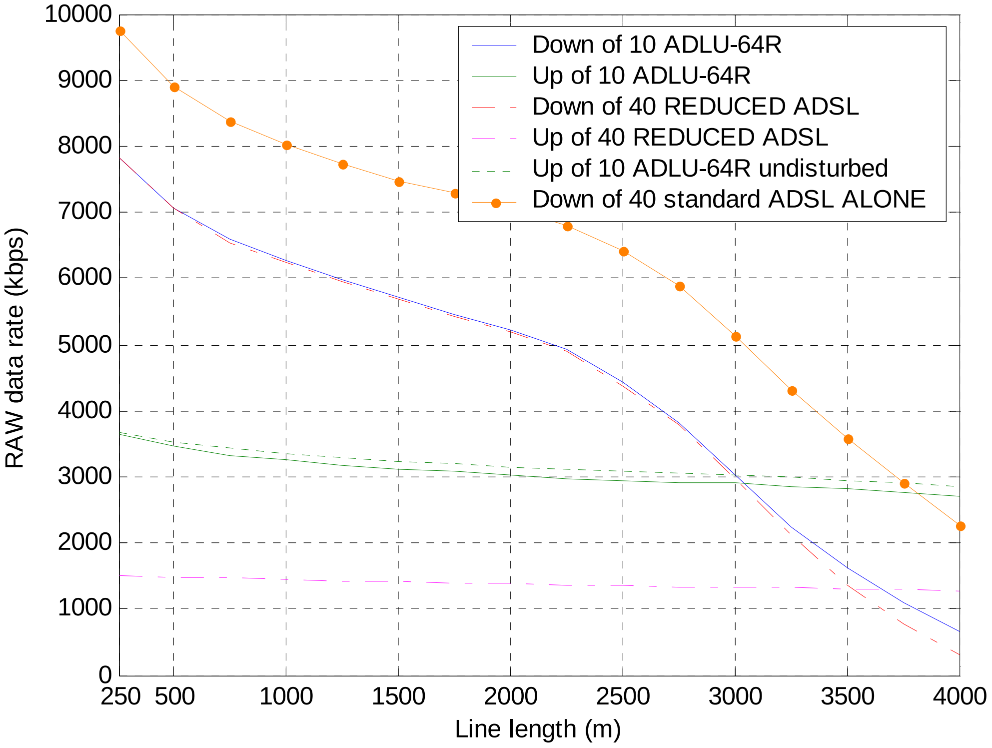

Figure 0.11 Spectral compatibility between reduced ADSL and ADLU-64

The graph shows the downstream rate and upstream rate of both ADLU-64R lines (the continuous ones) and the ADSL lines (the “dot-line” ones).

Is it also shown in the dotted line the upstream of 10 undisturbed (alone) ADLU-64R lines like in the previous graph for comparison. In this case, the upstream is nearly unaffected by the ADSL lines, as there is no overlapping spectrum now.

The problem in the reduction of the downstream spectrum of ADSL lines could be a loss of downstream data rate. To check this, another line is drawn. The continuous line with dotted marks represents the downstream of 40 ADSL lines without any spectrum reduction and without sharing the binder with ADLU-64R lines. Clearly, the bandwidth is reduced by about 2Mbps, and the difference is constant in all the line lengths. Note that the lower frequencies carry more bits/Hz and travel further, so this proposition is also likely to be impractical; the lower spectrum is CRITICAL for long-reach ADSL.

This data rate reduction could also be a problem if high-speed ADSL (like high speed Internet or Video over ADSL) and E1 over ADSL2 are both offered in the same binder.

3.1.11.4 Conclusions.

As it can be seen in the graphs, it is possible to offer E1 over ADSL2 at long distances of about 3km only if the downstream of the ADLU-64 line is statically reduced to avoid overlapping with the upstream (Figure 0.11). If ADSL lines are not reduced, the ADSL2 is severely affected (Figure 0.7).

Also it is possible to offer both ADSL and E1 over ADSL2 in the same binder, but only if the downstream spectrum of the ADSL lines are reduced to avoid crosstalk with the upstream of ADLU-64R lines.

This reduction in the downstream data rate of ADSL lines could be incompatible with high-speed services offered over ADSL and be available only at short distances. It is also incompatible with the crucial long reach ADSL.

This reduction of the downstream data rate of ADSL lines could be compatible with services like High Speed Internet access (4 Mbps at 2.6 km are reached) but may be incompatible with high speed services like TVoADSL at 6 Mbps in ATM layer, like Imagenio TVoADSL access in Spain.

In any case, this problem is also present if SHDSL is used to offer E1 or other symmetrical service over the local loop, as SHDSL crosstalk is more aggressive with ADSL signals.

3.1.12 E1 over ADSL2+

The main difference between ADSL and ADSL2+ is that ADSL2+ has much more downstream spectrum to use. As well as using slots from 38 to 255, ADSL2+ can use up to slot 511, providing in theory more than the double of downstream data rate at short distances.

The xDSL signal is more degraded at high frequencies than at lower frequencies, so this extra spectrum of ADLS2+ suffers more from distance and crosstalk. The result is that the increased data rate is only useful at short distances.

The next graph (Figure 0.12) shows a comparison of data rate between 10 ADSL lines and 10 ADSL2+ lines. The simulations are different, so the two kinds of lines don’t share the binder.

The ADSL2+ lines follow the specifications in ADSL2+ ITU-T G.992.5 Appendix I.

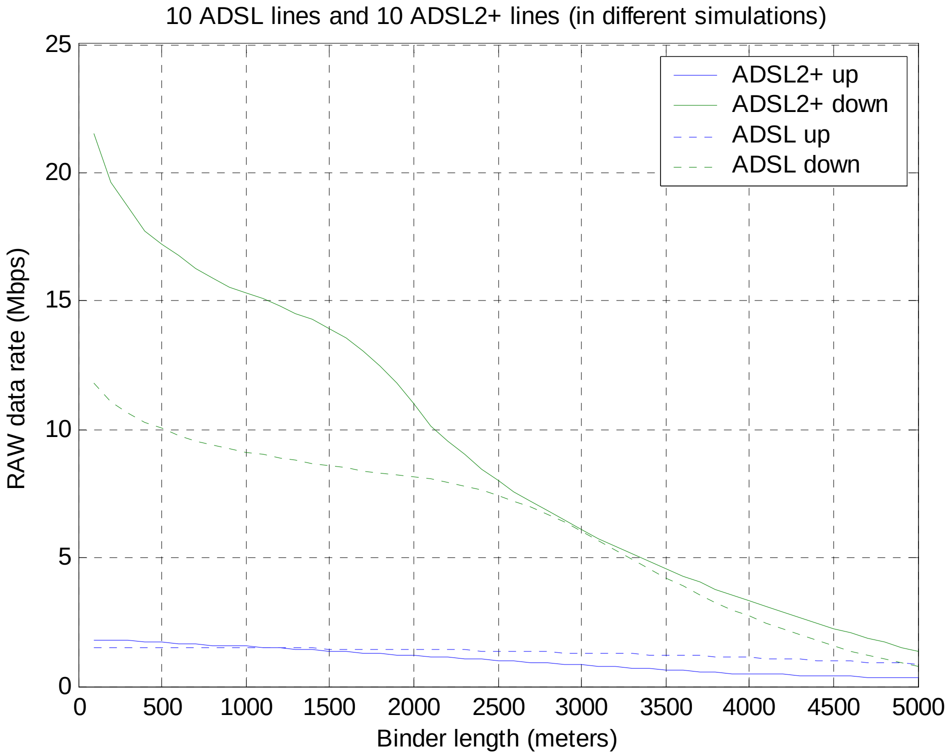

Figure 0.12 Performance of ADSL compared to ADSL2+

It is clear the big improvement in downstream speed, about double at very short distances. Over 2500m both data rates are very similar.

The upstream is slightly affected due to different spectral masks defined both in ADSL and ADSL2+ in upstream and downstream.

The next step, as with the ADSL2 lines, is to extend the upstream band to slot 63 and to limit the bandwidth used by downstream, so no crosstalk is produced in each line with itself. The objective is to check if it is possible to extend the range to use an E1 over ADSL2+, compared to ADSL2 lines with ADLU-64R.

3.1.13 E1 over ADSL2+ with restricted dowstream

A new kind of ADSL has been defined, ADSL2+R (R is from restricted downstream). It has the same upstream as ADLU-64R and the downstream is the same as ADSL2+ Appendix I but with the spectrum limited.

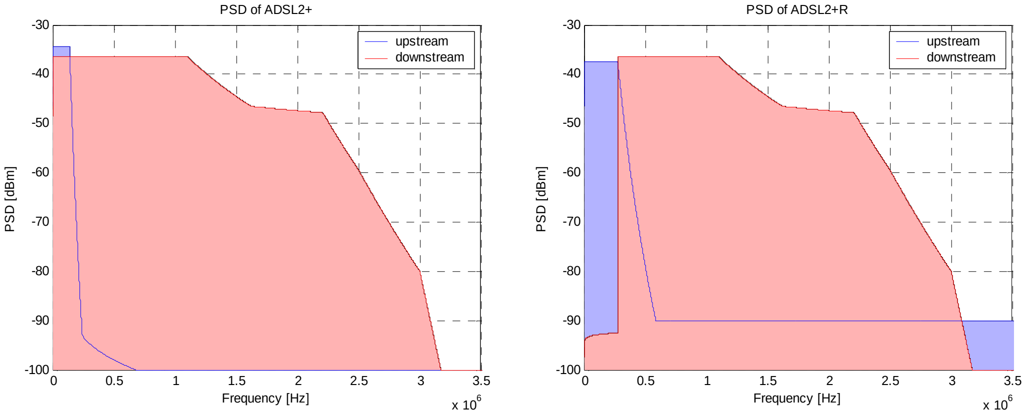

Figure 0.13 ADSL2+ and ADSL2+R Downstream PSD masks

The next graph compares both ADSL2+R (the modified with restricted downstream bandwidth) and the previous ADLU-64R lines. Both are simulated in different simulations (not sharing the binder) with ten lines per binder.

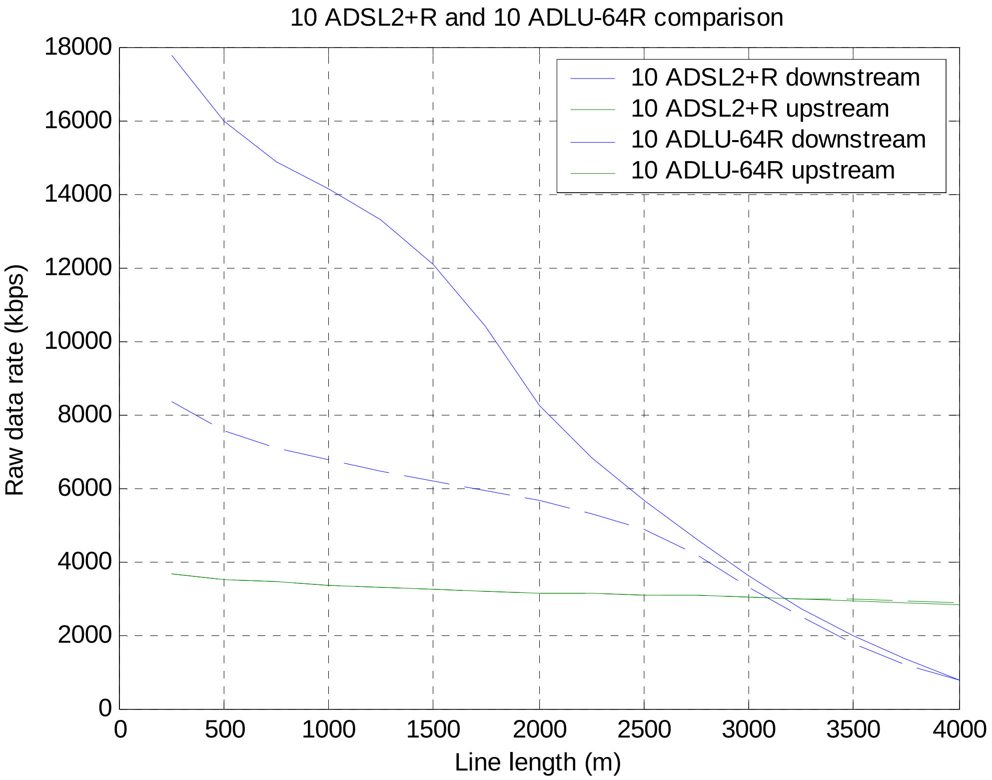

Figure 0.14 Performance of ADSL2+ compared to ADLU64R

The difference in upstream is minimal, but downstream is, as expected, almost the double for short distances.

At the significant point of about 3000 meters where downstream and upstream are equal, ADSL2+R provides a slight improvement, reaching about 50 more meters in the simulation.

A conclusion that can be obtained is that ADSL2+ with static spectrum management cannot provide much benefit in this situation.

3.1.14 Comparision of different ADSL2+ Annexes

Using the knowledge about the xDSL simulator, it is possible to create some graphs to compare the different speeds and distances that can be achieved.

The comparison has been made in identical conditions: different annexes of ADSL2+ G.992.5 are tested in different line lengths and with different number of lines in the binder. The number of lines has been chosen to representative values: 1, 10, 25, 45, 50, 80 and 100 lines, as used in the binders of different countries.

The noise present in the environment is only noise produced by the crosstalk of the lines themselves.

The annexes considered are the following:

- Annex A: ADSL2+ over POTS.
- Annex B: ADSL2+ over ISDN.
- Annex I: All-digital mode with spectral compatibility with POTS. Only the non-overlapped version has been considered.
- Annex M: ADSL2+ over POTS with extended upstream. Only EU-64 mode (the maximum upstream) has been considered.

3.1.14.1 G.992.5 Annex A

This Annex defines an ADSL2+ mode compatible with POTS. In this mode, the local loop carries a POTS line and the ADSL2+ service. To allow this, slots from 6 to 31 are used in upstream and slots 32-511 for downstream, providing a high downstream data rate and a limited upstream data rate.

Next graph shows the rate reached at different line lengths and line concentrations. Note the big impact of the presence of more than one line in the binder.

All the unnumbered lines in the lower part of the graph are the upstream part, much less affected by noise, crosstalk and attenuation, so they are almost together.

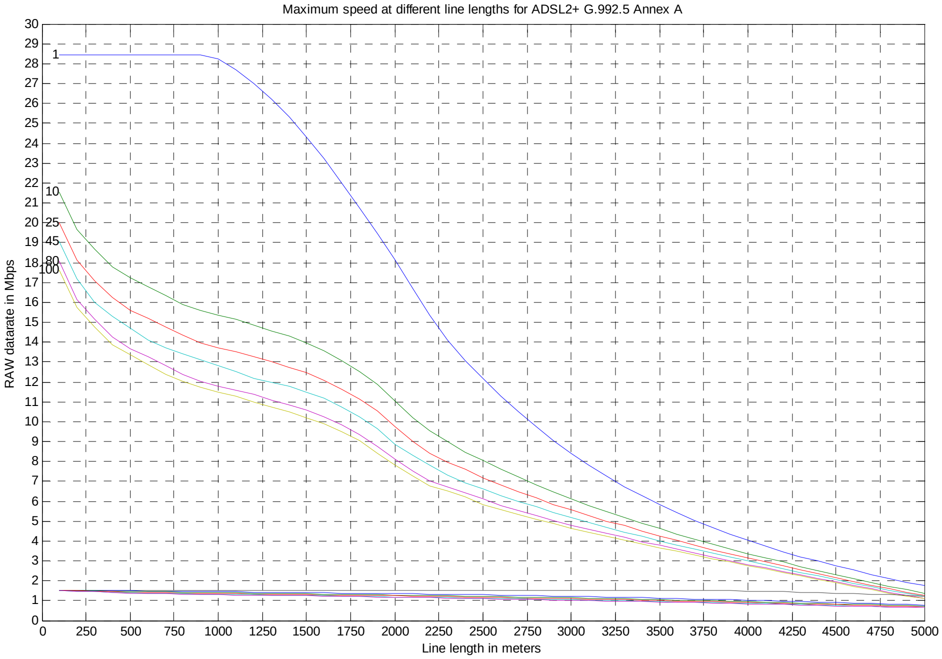

Figure 0.15 Performance of ADSL2+ (G992.5 annex A)

At a line length of 2500m and 3000m, two typical values that cover most of the potential customers, the raw data rate available is as follows:

Number of lines	Downstream kbps 2500m	Downstream kbps 3000m	Upstream kbps 2500m	Upstream kbps 3000m
1	12164	8392	1500	1500
10	5036	6132	1296	1216
25	7152	5568	1216	1144
45	6608	5172	1172	1096
50	6528	5108	1160	1088
80	6116	4808	1120	1044
100	5824	4640	1100	1024

3.1.14.2 G.992.5 Annex B

Annex B defines an ADSL2+ with a reserved spectrum for an ISDN channel on the same local loop. To reserve enough spectrum for ISDN, slots 29 to 63 are used in upstream and 64 to 511 in downstream. This reduces slightly the downstream data rate, compared to Annex A, where downstream is from 32 to 511.

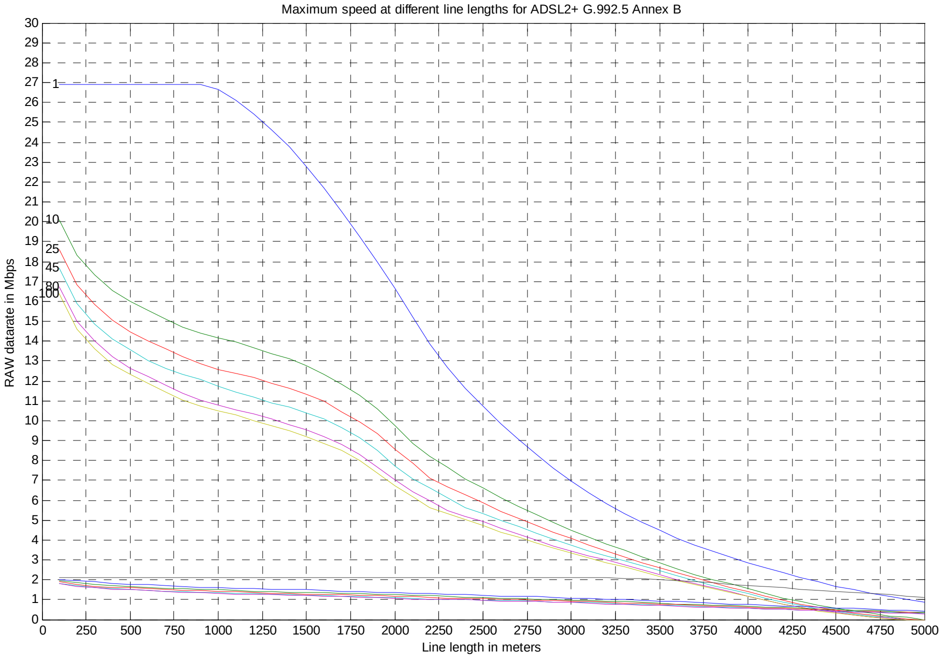

Figure 0.16 Performance of ADSL2+ (G992.5 annex B)

The data rate values at the considered lengths are as follows:

Number of lines	Downstream kbps 2500m	Downstream kbps 3000m	Upstream kbps 2500m	Upstream kbps 3000m
1	10724	6976	2100	2100
10	6596	4492	1200	1084
25	5856	4084	1108	984
45	5336	3756	1052	916
50	5260	3712	1040	908
80	4920	3468	988	864
100	4712	3328	968	844

3.1.14.3 G.992.5 Annex I

Annex I provides an all-digital mode for ADSL2+ where no POTS or ISDN service can be used in the same local loop. It is, anyway, compatible with other POTS services deployed in the same binder in different local loops.

There are two possible modes for this Annex: overlapped and non-overlapped spectrum.

The overlapped mode uses all the slots for downstream, and slots 1-31 for upstream. This overlapped spectrum will reduce the upstream for long distances and will produce more crosstalk in the binder to the upstream band in systems using different annexes.

The non-overlapped spectrum reduces the slots available for downstream to 32-511, as in Annex A, reducing the downstream, but the upstream is enhanced at long distances.

The resulting data rate graph is as follows:

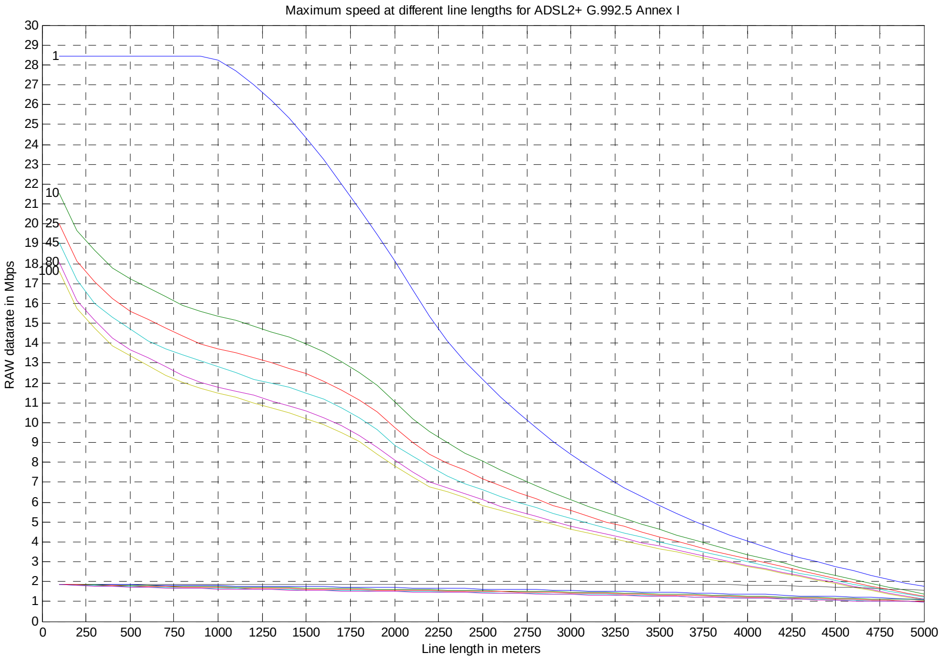

Figure 0.17 Performance of ADSL2+ (G992.5 annex I)

The data rate values at the considered lengths are as follows:

Number of lines	Downstream kbps 2500m	Downstream kbps 3000m	Upstream kbps 2500m	Upstream kbps 3000m
1	12164	8392	1860	1860
10	8036	6132	1624	1552
25	7152	5568	1548	1476
45	6608	5172	1500	1428
50	6528	5108	1484	1412
80	6116	4808	1448	1372
100	5824	4640	1432	1352

3.1.14.4 G.992.5 Annex M

Annex M provides a mode similar to Annex A. It permits use a POTS line in the local loop, but it provides an extended upstream spectrum. So not only can slots 6-31 be used, but the upper limit can be extended to slots 35, 39, 43, 47, 51, 55, 59 or 63, in different modes called EU-36, EU-40, EU44, etc. up to EU-64. Only EU-64 is considered here.

The downstream part remains unchanged, so there is a spectrum overlap between the extended upstream and the downstream, which will degrade the improved upstream data rate at long distances.

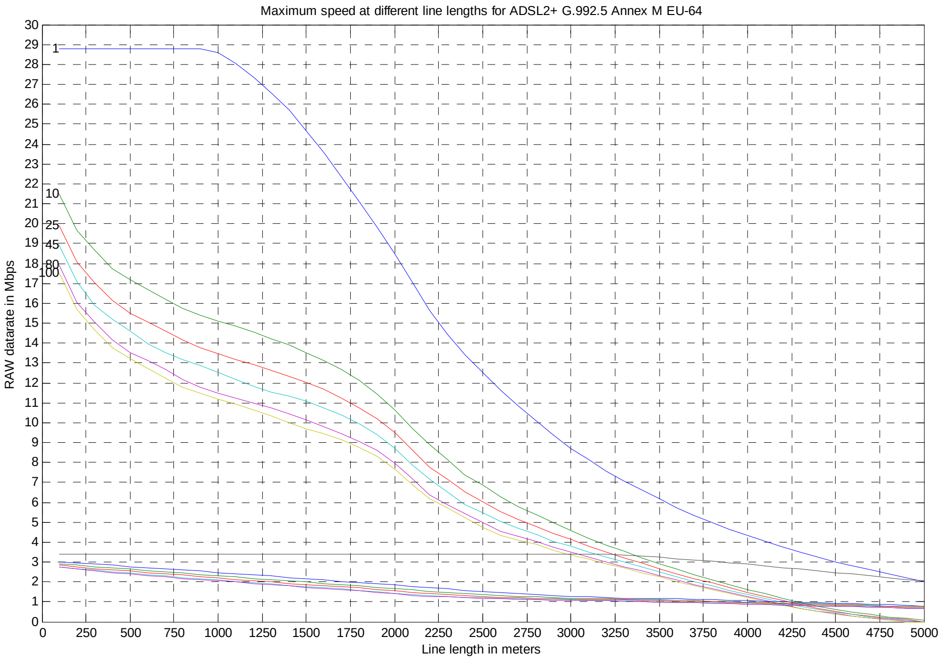

Figure 0.18 Performance of ADSL2+ (G992.5 annex M (EU-64))

The data rate values at the considered lengths are as follows

Number of lines	Downstream kbps 2500m	Downstream kbps 3000m	Upstream kbps 2500m	Upstream kbps 3000m
1	12496	8720	3420	3420
10	6844	4592	1500	1280
25	6004	4140	1352	1188
45	5484	3780	1276	1132
50	5388	3740	1264	1124
80	5000	3492	1204	1096
100	4748	3368	1168	1080

3.1.14.5 Comparison among different annexes.

For comparison purposes, two graphs will be presented. The first shows all the downstream rates for the four previously shown annexes, for a total of 25 lines. The second graph shows the same but with the upstream instead of the downstream.

This will allow a quick comparison among the different annexes explained, for a total of 25 lines. The results for a different number of lines would be different, but only quantitatively, and not qualitatively.

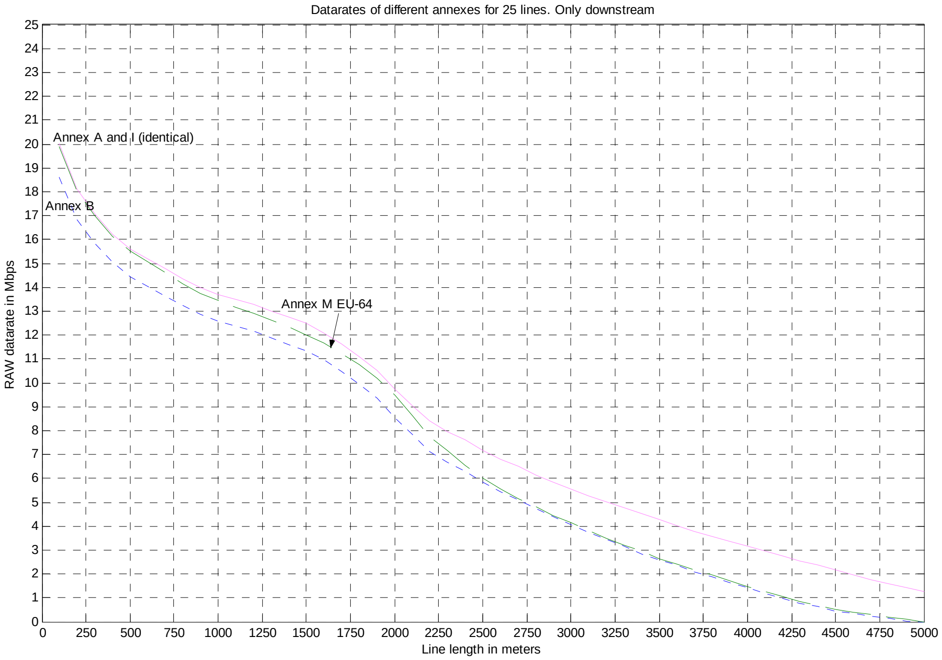

Figure 0.19 ADSL2+ annexes compared downstream

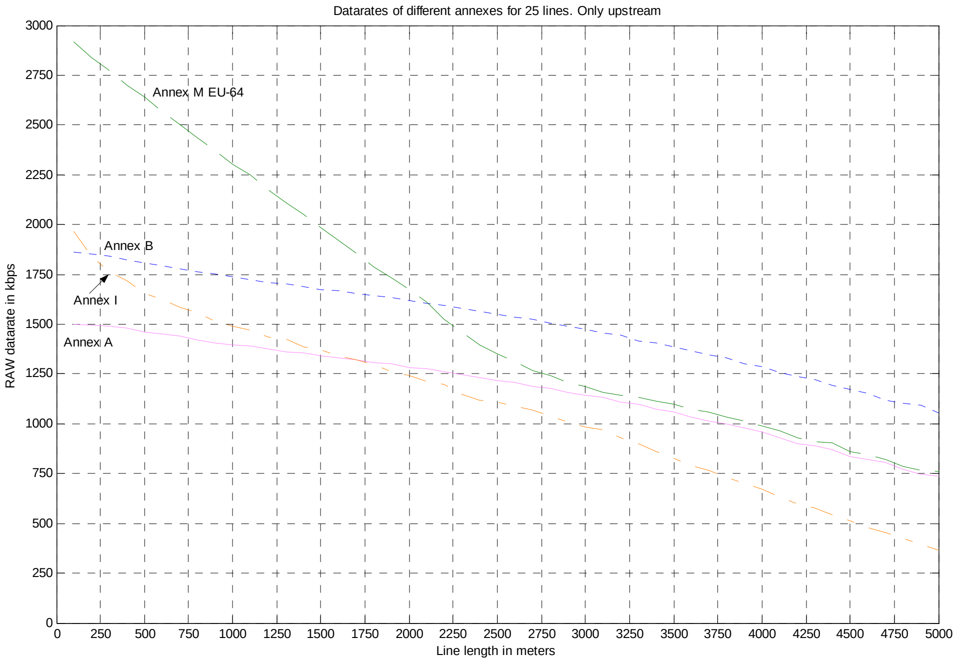

Figure 0.20 ADSL2+ annexes compared upstream

Intra cable cellular bandplan

Let us consider an attempt to reduce crosstalk by assigning each pair a frequency band which is not used by any of its immediate neighbours – a cellular approach similar to that used by mobile phone operators. The strongest crosstalk coupling is between immediate neighbours, and more distant neighbours are typically at least 15 dB better isolated.

This would only be possible in a network where pair adjacency is known and is maintained, so we may exclude the USA (where cables have unstructured bundles) and the UK (where ‘random jointing’ means pair adjacency is not maintained between cable sections). It might be viable in countries which use quad construction.

The cellular operators use a tessellation of seven cells. For this study lets suppose ideal hexagonal packing in the cable and use a tessellation of three cells. So each pair uses a third of the spectrum, but experiences 15dB less crosstalk and so gets 5 more bits per Hz in that spectrum.

Break even for this approach is when ordinary spectrum management gives 2½ bits/Hz. Above this figure the strategy of all users using all the bandwidth is better.

A channel whose SNR gives only 2½ bits/Hz is a poor quality channel by most standards : the original ADSL could not use a carrier whose quality was much below 2 bits/Hz. So this cellular approach might be of interest for poor quality capacity – say on longer lines. It will not be of interest for the higher data rate environments DSM seems to address.

4 DLCM ANALYSIS

In this paper 'Dynamic Line Code Management' concerns the engineering of performance of a line independently, using methods which do not change signal spectra and so do not cause changes to the crosstalk environment[13] of other lines in the same cable.

Currently the main method of interest to the literature is the adjustment of error correction configuration in response to the recent noise experiences of a line and the declared intentions of the user of that line.

The scope of DLCM includes *any* internal adaptation by a system in response to its channel, including bit loading in ADSL systems. Formally it would include all adaptive equalization.

FEC configuration : RS & interleaving against REIN

The vast majority of installed ADSL lines are configured to use the FAST path and do not use interleaved Reed Solomon coding. It is not clear exactly why this happened, but it may have been due to the perception that games users would be very sensitive to delay, and that video transport was not imperative.

- In DSM a need is identified for seamless power control – adaptation without disturbing the data streams. For DLCM it seems likely that the main adaptation will be changing the Reed-Solomon interleaving dimensions, necessarily changing the latency delay of data. For constant-bit-rate services such a change will cause discontinuities in the data. So DLCM can only be specified as *near* seamless.
- DLCM is largely concerned with adapting against impulsive noise. Impulsive noise in domestic systems appears to be mostly due to mains-borne transients which then couple into the house telephone wiring. It may be of interest that if triple play is achieved then one should expect electrical isolation between the line and the house telephone wiring, and a consequent gross reduction in the incidence of impulsive noise.
- DLCM may of course still be useful in mitigating the noise picked up from mains distribution cables running alongside telephony ducts outside the boundary of the customer premises, even in a triple play scenario.

4.1.1 Field experience

In a recent survey of long lines in the BT network, experiments have shown that stability of lines can be substantially improved by enabling the interleaver (and RS FEC overhead) on lines when it is believed that impulse noise is causing the line to retrain frequently. There was however no increase in the bit rate achieved in rate adaptive mode after enabling the interleaver, just a reduction in the number of times the modem lost sync.

---

[13] this comment is only true assuming SSM. DLCM engineers an internal state of the system it serves, and this internal stste could become externally visible under DSM because the system might make different back-off decisions.

4.1.2 Laboratory experience

Experiments in the lab echo these findings. In an experiment using simulated self-crosstalk on cables of 30 dB and 40 dB insertion loss at 300 kHz, the bit rate at which the DSL modem would remain in sync was determined.

Then simulated repetitive electrical impulse noise (REIN) was introduced, with a repetition rate of 100 Hz, duration 250 us and PSD -100 dBm/Hz.

In non-interleaved mode, the bit rate at which the link could remain in sync dropped substantially when the simulated REIN was introduced. In reality TCP/IP throughput will almost certainly drop further, as the errors introduced by the REIN are not all corrected when operating in FAST mode.

Allowing use of the interleaver with 4 ms interleaver delay resulted in a modest increase in the bit rate at which the link could remain in sync.

Increasing the interleaver depth to 8 ms saw a further substantial increase in the bit-rate at which the modem could remain in sync. No further increase was seen when setting the interleaver depth to 16 ms, however.

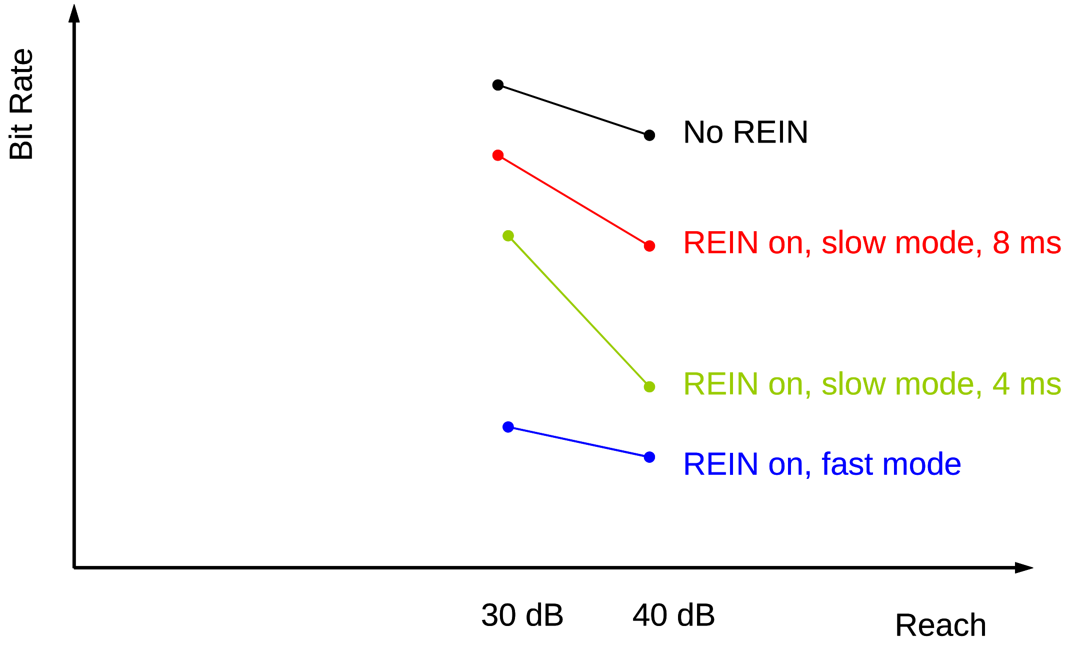

Figure 0.1 Sensitivity of data rate to interleaver depth

This chart illustrates that the process of determining the minimum interleaver depth and RS codec settings at which the end user can be assured of error-free, stable operation, is the foundation of a DLCM system

Virtual Noise

VN is a proposition in play for developing standards. It proposes that noise ‘other’ (not stationary like crosstalk, not regular like REIN) be combated by reserving extra margin.

This is an alternative use for the excess margin/capacity, the same resource DSM seeks to exploit, and so may be seen as a competitor to DSM.

Dynamic Response Under SSM

A system which transmits at a constant spectrum, say using all the spectrum available under an imposed SSM mask, will experience changes in its channel as changes in available capacity. For instance ADSL systems turn tones on and off as channel quality varies. So if services are engineered to survive changing capacity then SSM systems can exploit them too.

Under the current static spectrum management rules many operators are offering fast internet services that can flex up to the maximum data rate that the line can support (with a set noise margin), or the maximum the modem can sustain. This approach tends to minimize noise margin, so such systems stand to gain more from the FEC configuration aspects of DLCM.

Conclusion

The primary benefit of DLCM is improved service stability. However, link capacity and service stability are generally tightly coupled. DLCM does not provide systematic capacity gain, but in an environment where extrinsic electrical impulse noise is a significant factor it may in practice provide greater capacity benefits for the individual subscriber than DSM, by enabling service where otherwise no service is possible.

Bearing in mind the fact that DSM tends to minimize signal power across the community of modems, and the hence noise margin of each system individually, if DSM is deployed we should expect systems to become more susceptible to noise; and DLCM may be expected to become more desirable still.

Unfortunately the benefits of DLCM are even harder to predict with confidence than for DSM, since so much depends on highly variable electrical noise conditions in the network. No two customers’ lines are the same, nor can the conditions on a particular line be expected to stay the same for long. However we note there is a rising trend in electrical interference from domestic appliances, and worrying signs that this trend is continuing – even accelerating – as switched mode voltage converters become more common. So despite the uncertainties it is expected that DLCM will become increasingly important.

5 DSM : DYNAMIC SPECTRUM MANAGEMENT

In this paper ‘Dynamic Spectrum Management’ concerns the engineering of capacity of a group of lines by altering their transmit spectra co-operatively.

To Quantify Advantages

This section explores the behaviour of multiple systems in an autonomous DSM environment. Although based on quantitative simulations, the object here is to get a feel for qualitative behaviour.

5.1.1 The Simulated Configuration

50 DSL systems are simulated, sharing a single cable section 1 km long. A novel crosstalk model is used [see Appendix 2 - Per-Line FEXT Model], which seeks to be realistic rather than pessimistic. Given the transmit spectrum of each system the capacity of each system may be calculated.

Crosstalk is modelled as FEXT, so the simulation represents ADSL-like channels in a FDD regime. For simplicity only channels in one direction are analysed here; we suppose that domestic users will continue to be very asymmetric in their downstream/upstream needs, so the direction represented is supposed to be downstream.

We also assume modems are subject to transmitter quantization noise at -130 dBm/Hz. This defines our noise floor.

To represent a static spectrum management environment suited to ADSL2+, each system is limited to a PSD mask of -40 dBm/Hz within a 2 MHz band, and no power outside it. DSM power back-off is simulated as frequency independent attenuation.

To simulate autonomous DSM each system is assigned a target capacity (possibly different for each system), and each system is started at the PSD mask. Iterations calculate capacity, and modify each system's transmit power[14] in pursuit of its target. Iteration stops when a stable set of powers is reached. We find:

- Stability is reached (we never found a case which failed to converge[15])
- If any system did not achieve target it is at maximum power
- If all systems achieve target then convergence is *slow*

The slow convergence is believed to be a real effect : the systems fairly rapidly find an arrangement of their powers where each has a slight excess capacity, and these are roughly equal; then they reduce their powers in very small steps, giving up a fraction of their apparent excess, and the steps are the pretty much the same size.. While the noise floor is not significant one such step has no effect on SNRs. There can be a lot of such steps before the noise floor becomes significant and equilibrium is approached.

---

[14] subject to the mask upper limit of -40 dBm/Hz and a lower limit of -140 dBm/Hz

[15] from this work we cannot exclude the possibility of chaotic oscillation in a real system where couplings vary with frequency

5.1.2 Full Power (i.e. SSM)

As an introductory simulation, consider the capacity for each system if every line is operated at full power:

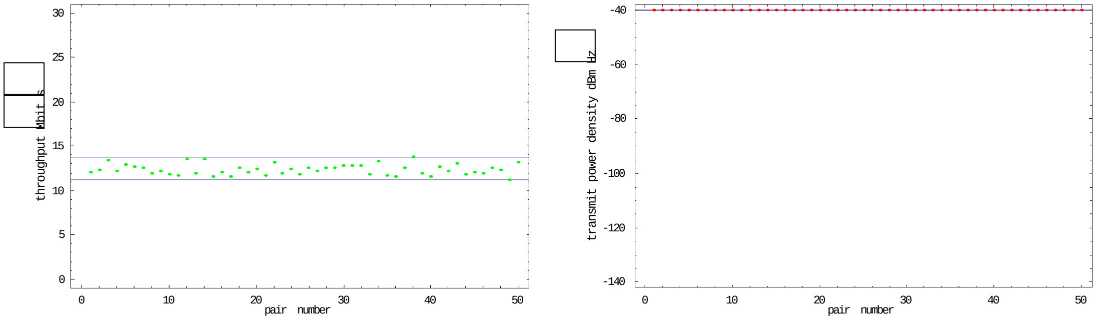

Figure 0.1 Downstream data rate scatter with all systems at full power

the graphs show the power applied to each line (right) and the data rate obtained on each line (left). The cursors show the max and min values.

The lines attained data rates in the range 11.2 to 13.8 Mbps. This represents the performance if all systems in the configuration seek their maximum data rates subject to the PSD mask.

The close bunching of results, given the wide range of the individual crosstalk couplings, is because each system experiences the sum of crosstalk from all the couplings into it, so the individual variations are largely averaged out.

(For consistency the simulations presented here use one instance of a cable : the crosstalk model [in Appendix 2 - Per-Line FEXT Model] is random and so can generate an infinite number of different instances, but we use the same one thoroughout.)

5.1.3 Equal Capacity

Suppose DSM were used to obtain the same capacity for all users of this cable, and the highest rate was wanted. In this simulation systems started at maximum power, and those with larger capacities backed off to benefit their peers until all attained the same capacity.

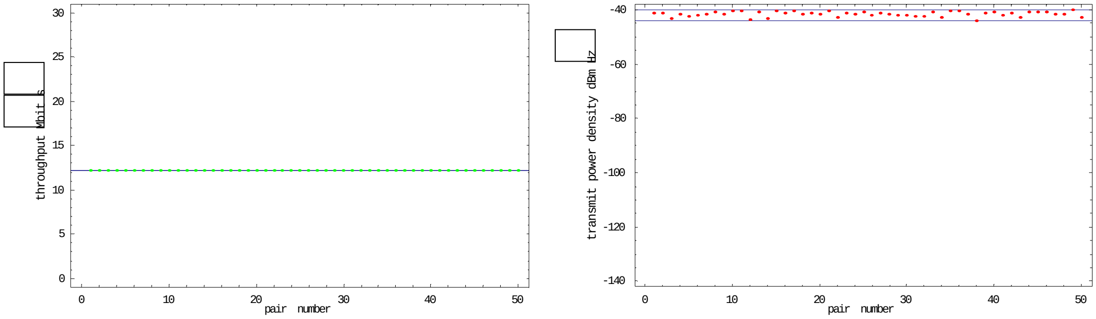

Figure 0.2 Equal rate for all, operating at saturation

The capacity attained here was 12.16 Mbps. This represents the potential for DSM on the simulated configuration when a central authority dictates rate to obtain the highest possible data rate for all equally. The range of transmit powers is -43.9 to -40.0 dBm/Hz. Note that at saturation all systems achieve their targets and one or more systems is at the power limit.

To provide an operating margin, lets suppose this ensemble is set targets slightly below this, at 12 Mbps

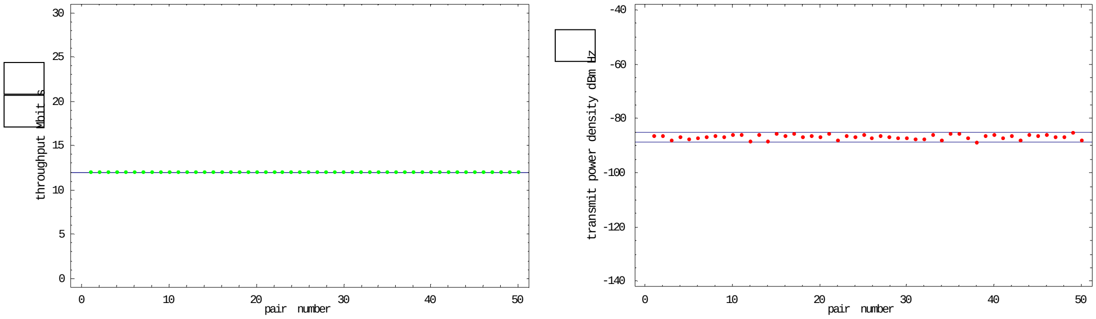

Figure 0.3 Equal rate for all, operating near saturation

The range of transmit powers is -88.8 to -85.2 dBm/Hz, a fall of 45 dB for a cost of 1% of capacity. This suggests terrible sensitivity at saturation; and that one could be operating close to saturation and be unaware of it from the transmit power levels.

5.1.4 Cost of margin

In practise robust operation may require an operating margin, say to allow service to survive during the ensemble's reaction to small environmental changes. Let us consider the cost of a 2 dB increase in margin:

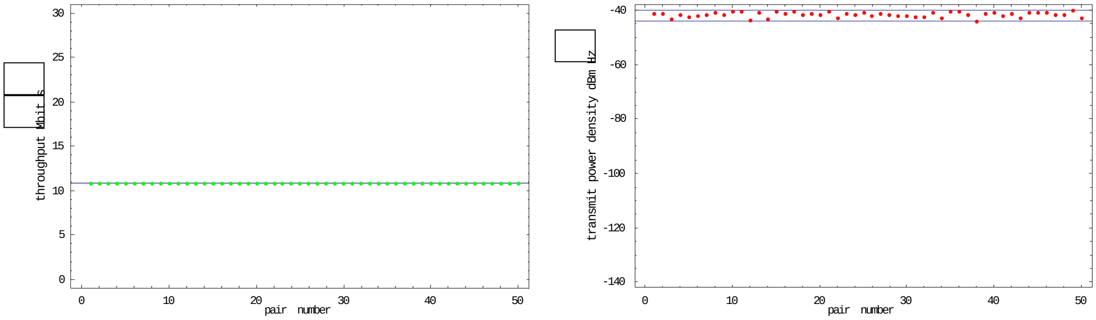

Figure 0.4 An extra 2dB margin, operating at saturation

The capacity attained here was 10.85 Mbps. The range of transmit powers is much as before, -43.9 to -40.0 dBm/Hz. The data rate reduction is consistent with the simulation's bandwidth, of 2 MHz; the 2dB margin is under these circumstances exactly a reduction of 2dB in each line's SNR, costing ⅔ bit/Hz, or 1.3 Mbps, which is what we see here.

So the data rate reduction due to even a modest practical margin is significant.

5.1.5 Conservative rate

The traditional conservative telco will offer near the worst case rate available[16], and will also want extra margin as their comfort factor (although at present some telcos are going away from this for their domestic broadband offerings.)

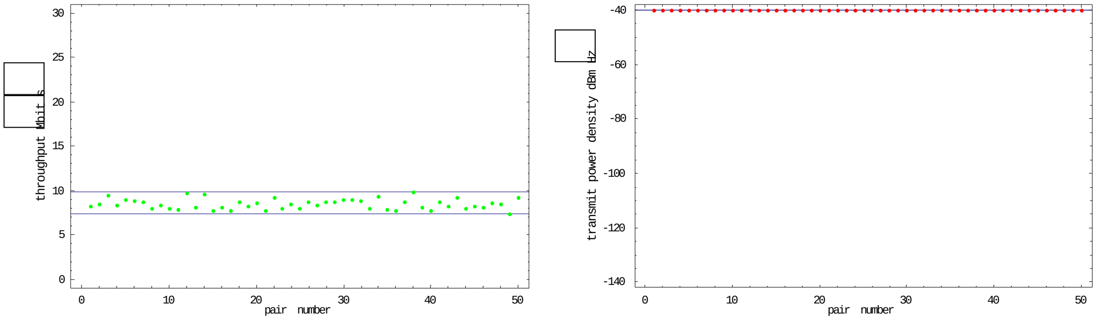

Figure 0.5 A conservative telco would only offer 7 Mbps

Simulating with an extra 6 dB obtained a range of similar width to the full power case and about 4 Mbps lower (here from 7.4 to 9.8 Mbps). This suggests the simulated configuration could be offered a data rate of 7 Mbps by a traditional conservative telco.

5.1.6 Bandwidth on demand

To enable bandwidth on demand, the 'normal' state of the ensemble must have spare capacity. Let us suppose this background state is set by each line having a target of 8 Mbps; then of course all systems can meet their targets:

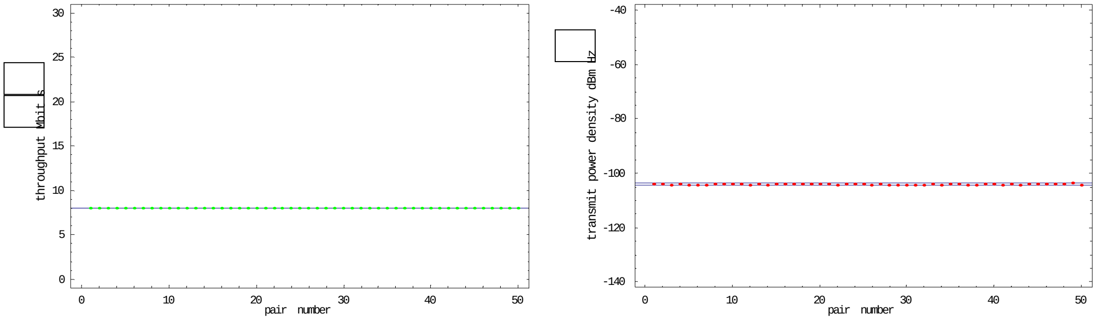

Figure 0.6 Downstream power scatter when use is less than cable capacity

The transmit PSDs are in the range -104.6 to -103.7 dBm/Hz, which is about 16 dB above the noise floor. Of course this level would track the noise floor. The transmit PSDs here for 8 Mbps are lower than those of the 12 Mbps case in section 5.1.3 above, but the decrease is much less dramatic than the decrease there from saturation to 12 Mbps.

---

[16] here taking the worst performance of 50 lines to get the 99%ile worst case performance.

Case	Common data rate	PSD range
saturation	12.16 Mbps	-42.0 to -40.0 dBm/Hz
12 M	12.00 Mbps	-88.8 to -85.2 dBm/Hz
8 M	8.00 Mbps	-104.6 to-103.7 dBm/Hz

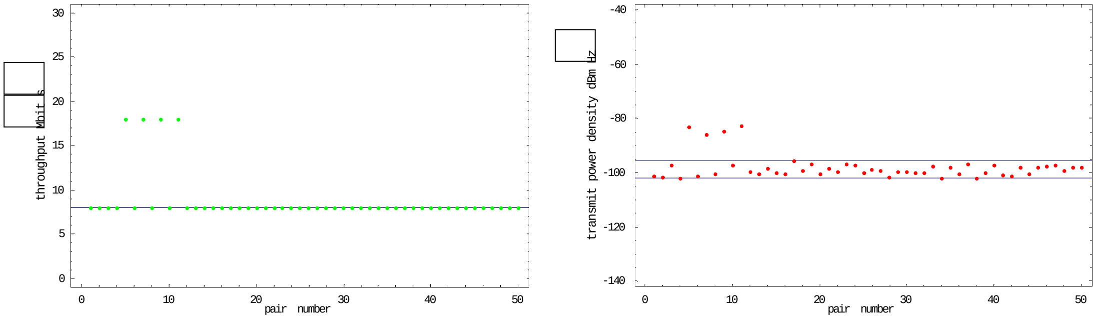

Figure 0.7 A few systems can attain higher targets

A few systems can now attain higher rates. Consider, say, 18 Mbps. Empirically we find this cable can support about four such higher rate systems, although it depends which lines. The graph above shows four arbitrarily chosen lines attaining the higher target. The cursors on the graphs now show only the bounds of the 8 Mbps systems, it is evident that the power required by the ‘normal’ systems have increased a few dB to (-102.0…-95.5 dBm/Hz). Indeed a fifth line can also run at this rate

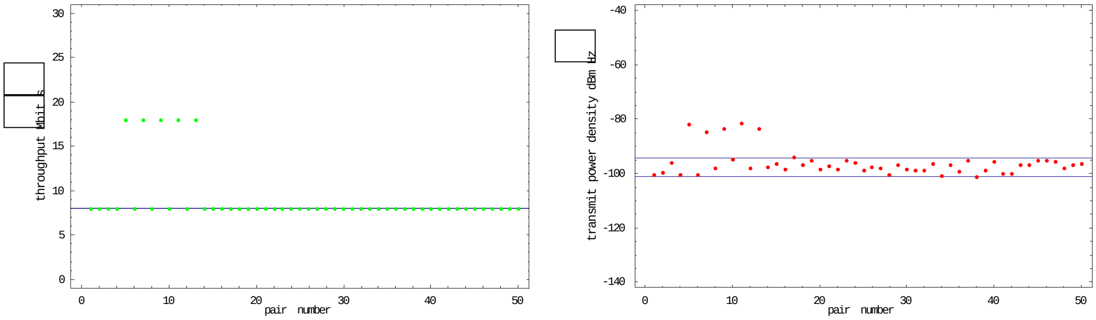

Figure 0.8 Five systems all achieve the extra rate

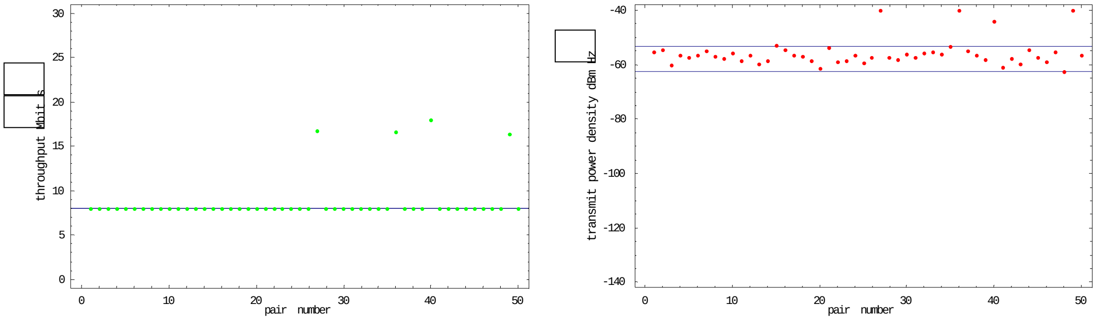

Figure 0.9 Four systems do not all achieve the extra rate

However a different choice of lines gives a case where even four lines don't get the higher rate. So the rates available are arbitrarily variable between lines (perhaps to be expected, given random coupling). We note the characteristic maximum power on those lines failing to attain their target, and the consequent gross increase in the power needed by those lines which do attain their target.

5.1.7 Non-cooperative sources

For a real network there will be noise sources which are not party to the DSM cooperation deal. These include legacy systems predating DSM, faulty modems which do DSM wrongly, external sources such as RFI and REIN, and LLU free spirits who don't want to do DSM.

To represent such a non-cooperative source, consider the case where all attempt a data rate of 8 Mbps but one modem does not back off:

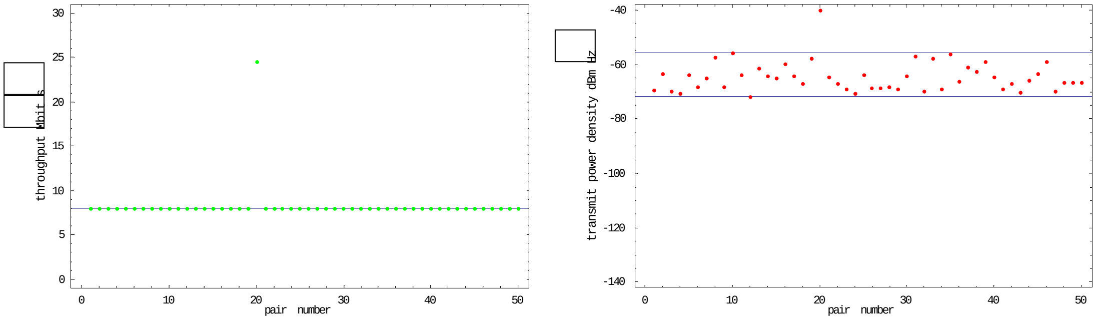

Figure 0.10 One non-cooperative system

Unsurprisingly that one system has a higher rate[17]. We note that, as before, the power needed by the other lines has risen grossly. We also note the range of those powers is wider than previously (-72.0 to -55.8 dBm/Hz), as the variation in couplings is now *not* averaged out; every line suffers mostly from that one system via the single coupling into its own line.

---

[17] *potential* rate. It could alternatively be following the philosophy of virtual noise [section 0] and running with a large margin.

5.1.8 *Sensitivity to small changes*

In practise there will be changes to the ensemble of systems from time to time, as individual systems are turned on and off. Given the sensitivity sometimes evident in a DSM ensemble, let us study the effect of one new system turning on.

Firstly, to appreciate the change in equilibrium (i.e. after the ensemble has settled down) consider the case of all systems equal, near saturation (the 12 Mbps case as in figure 0.3), but with one system turned off, and then turn it on. Beforehand one might expect the noise environment for each of the already-on systems to increase by about a fiftieth (0.09 dB), there being 50 lines in the simulation.

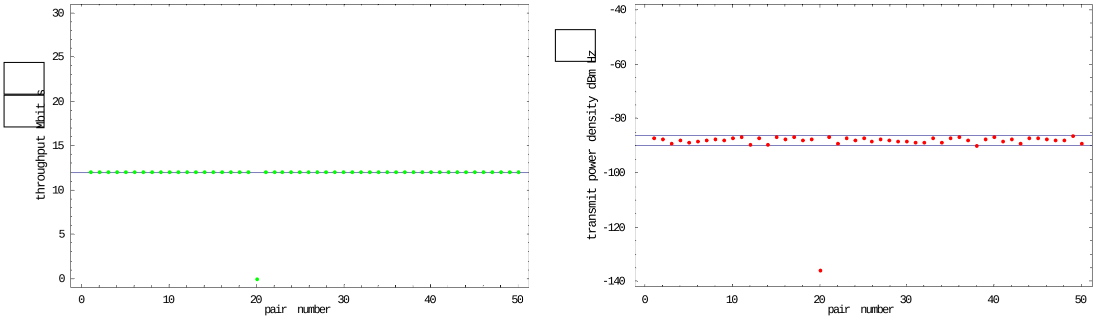

Figure 0.11 Just before new system joins ensemble

In the simulation here the range of transmit powers is -90.0 to -86.6 dBm/Hz, about 1.2 dB below the all-on case. So turning on the extra line increases the equilibrium PSD levels by about 1.2 dB. (For the ensemble operating at 8 Mbps the change is about 0.01 dB)

While this increase is more than might be expected, its still small enough to be engineered round by allowing each line to have a couple of dB margin – although that will itself cost as a reduction in the available data rates (see section 5.1.4).

Potentially more serious is the transient while the new system trains : present ADSL systems perform the training phase at near full power, so there would be a gross impact on all the other systems (15 to 30 dB: compare with section 5.1.7). For DSM use modems would have to train starting at very low power.

5.1.9 *Realistic networks*

In practise we are interested in more complex networks, the dendritic networks of a typical access network, with DSLAMS at the exchange and possibly elsewhere too.

The simplest form of DSM presumes that in a community of modems sharing the plant, each will use the minimum power necessary to meet a single service data rate, uniformly applied across the ensemble. If the uniform service rate is set well below the ultimate capacity of the network, then it might be expected that the ensemble power levels would all fall to low levels, eventually limited by the extrinsic noise level, or a minimum power limit set by the network operator.

A more complicated arrangement has two kinds of service sharing the plant, one hosted from the exchange and one hosted from a remote cabinet. Here the network operator has a two dimensional optimisation to control. The simplest approach is again to implement a uniform service. This may be appropriate where the cabinet is provisioned primarily with a view to extending the reach of an exchange-based service to reach customers that are otherwise too far from the exchange[18]. However current interest in the UK is for cabinets provisioned to provide higher data rate services, the locations closer to the customers benefiting from lower loop loss. In this case the cabinet based and exchange based service data rates will be different, and a challenge is how to determine these rates. This selection could be driven by commercial needs, but would normally be the subject of unbundling regulations. This then implies the need to define these rates based on some notion of fairness in the allocation of a right to pollute.

However the commercial drivers to provide multiple data rates, best effort service types, as well as premium connections, means that optimisation to simultaneously meet multiple rate targets is even more complex – even in the case where only one network end signal injection point is considered. This section describes some initial simulations that explore the behaviour of a modem ensemble, where lines normally operate at a rate below the best that could be provided equally, but where some are permitted to operate at significantly higher rates than could be guaranteed across the ensemble. For now it is assumed that some kind of network control would be necessary to prevent the cable spectrum resource being exhausted in a 'cocktail party effect'.

5.1.10 Initial Impression

The main attraction of DSM is the belief that there is a mass of untapped capacity waiting to be exploited. Such a belief might be encouraged by figure 0.3 above, where the systems are known close to the cable capacity and there is a rather large gap between the power levels the systems need and the power level limit. However the simulations above suggest this gap does not translate to capacity. If the simulations are realistic then DSM is motivated by a mirage.

5.1.10.1 In detail

The backoff regime gives the collection of live systems great sensitivity as the cable approaches saturation. Running a network in this regime means that the set of systems will experience great changes in power settings from time to time as saturation is entered and left. It remains to be seen whether such a transition will disrupt services on lines which are fully serviceable before and after the transition.

If not saturated, the equilibrium tracks the noise floor[19]. However near saturation the ensemble exhibits positive feedback, and the response to changes is disproportionate. Credible external noise changes are movements in the noise floor (such as diurnal RFI changes), and systems turning on and off.

---

[18] This may become necessary for universal service, which is MUSE’s objective

[19] in a simulation environment we can of course define noise floor, apply it uniformly, and know what it is a priori. In the real world we would presumably measure the noise floor by observing the equilibrium point.

Presumably this sensitivity would still be present if the systems reserved themselves some fixed margin – the equilibrium points would move but the new equilibria would be equally sensitive. However the services carried would be buffered against small changes, and then the new equilibrium can be sought without service loss.

The use of DSM to increase the rate promised to all users in a community seems to offer only a small increase. In the simulated examples, without DSM the lowest common denominator was about 11 M bps [section 5.1.2], while with it the greatest common rate was about 12 M bps [section 5.1.3].

5.1.10.2 Comparing with initial field experience elsewhere

Administrations elsewhere have reported good things about 'DSM', which for them has meant manually adjusting transmit powers on some shorter lines on their own DSLAMs such that those lines still attain their margins while their neighbours become viable. The beneficial effects are most pronounced under inhomogeneous conditions, for example when exchange and cabinet solutions co-exist in the same plant.

The analysis above has assumed conditions which are believed typical for European administrations[20], in particular that there is a well engineered SSM environment and that there is a significant presence of systems not operated by the incumbent. Under these circumstances the capacity gains are quite small, even with our rather idealised assumptions.

However we note that complicated techniques such as those discussed in section 0 will be needed to keep SSM working well if a network must mix deployment from exchange and cabinet in the same cables for systems which use the same frequencies.

Engineering DSM

This section discusses the engineering developments which will be needed to make DSM viable. This discussion is necessarily speculative in places since we have taken a view where the literature is still discussing strategic questions (such as what constitutes DSM, what the excess capacity is to be used for, and what to do when capacity is outstripped by demand).

---

[20] Or perhaps aspirational for their regulators

5.1.11 DSM Objectives

The use for the higher capacity DSM allows is open to discussion. The main contenders appear to be:

- Providing higher equal capacity for all players [2]
- Providing higher capacity on demand [1]

5.1.11.1 DSM Algorithms

All the DSM algorithms seem to assume that a modem has a target data rate and license to change its power settings in pursuit of that rate. Related to this will be a protocol whereby each system may know what maximum data rate it may attempt. This could be by command from some central authority, by arbitration with the central authority (who might even charge more for a greater license to pollute); or it could be that a separate spectral bound is provided as part of ordinary spectral management (the current UK position).

Most authors do not address what should happen when the systems' targets are not all reachable. This author[21] believes that, for protection of the community, there should be a separately specified max power limit for each system (i.e. ordinary SSM, for the usual reasons).

The optimal method for a civic-minded system to relinquish unneeded signal power is still a matter for research. The options include:

  - Classical water filling
  - White backoff
  - Reddening backoff
  - Blueing backoff

5.1.11.2 Classical water filling (IWF)

In the classical literature 'water filling' is the method to design a signal with minimum transmit power, on a fixed channel with a fixed noise spectrum. Cioffi claims it is also viable and stable for DSM.

It is discussed in the position paper [1] under the title 'iterative water filling'.

5.1.11.3 Optimal Spectrum Balancing algorithms

The position paper [1] discusses algorithms, which have the freedom to change the FDD channels' spectral allocations. They are essentially constrained maximisation algorithms, for example the normalised-rate iterative algorithm (NRIA), developed at FTW [3],[4].

These are not pursued in this paper: it is suspected they will not be viable in Europe because of the apparent need for ordinary static spectrum management to support LLU, and the apparent need for DSM to be subject to its local static spectrum management regime to protect LLU systems.

---

[21] Rob Kirkby

5.1.11.4 White backoff

In which power is adjusted by equal attenuation at all frequencies. The simplest back-off strategy, used in the simulations of section 4.

5.1.11.5 Reddening backoff

In which power is withdrawn from highest frequencies first. This is beneficial to neighbours in that crosstalk coupling increases with frequency, so withdrawing from the high frequencies disproportionately reduces coupled power. In practice it is only of benefit to systems on short lines, as the longer lines cannot use the higher frequencies anyway.

As noted earlier, ADSL already does this to a certain extent : when a carrier becomes unusable it is turned off.

5.1.11.6 Blueing backoff

In which power is withdrawn from the lowest frequencies first, in order to benefit longer lines which can only use the low frequencies.

Note that longest lines are often limited by their internal noise, rather than crosstalk, because they only use low frequencies and crosstalk is not the dominant noise source at those frequencies. So the blueing approach seems unpromising.

5.1.11.7 Comparison

An initial comparison suggests vacating spectrum completely is only really valuable to a neighbour who is then the only remaining user of that spectrum. If there are two or more remaining users then white backoff appears better.

5.1.12 Stability

We have already noted that a DSM community reacts sensitively to small changes in usage near saturation . While this need not of itself cause oscillatory behaviour, reaction time lags may produce oscillations (or even chaos). In other branches of engineering this is controlled by damping of responses.

It is of course highly desirable that power level swings in the community do not harm systems which are serviceable before and after the swing. Present ADSL systems do behave pathologically in these circumstances; each system adapts by retraining, and the act of retraining is itself harmful to the neighbours. The stressed ensemble ‘flaps’.

5.1.13 Modem changes

A DSM capable modem will need these features:

- Seamless adaptation (of bitrate and of transmit power)
- Identify traffic of different statuses
- Participation in the process which assigns status to each data stream
- Discard data if the channel cannot carry it.
- Rearranging data between the form at its interfaces and the form on line.

Clearly such new capabilities will mean new standards (or possibly a substantial revision to the existing ADSL standard). New tests will be needed, although from experience it is not possible to predict these until the engineering is done. The new capabilities are a significant increase in the complexity of an ADSL system's behaviour, but perhaps not of the systems which succeeded them. In particular the backoff features are similar to the backoff regime for VDSL, so testing should be similar. Development time for the enhanced ADSL standard should be comparable to the development time for VDSL.

Depending on the algorithms used to determine the transmit power in a DSM community, the definition of margin may need to be tightened. In G.992.3, the definition of SNR margin is "The signal-to-noise ratio margin is the maximum increase (in dB) of the received noise power, such that the ATU can still meet all the target BERs over all the frame bearers."

There is no guidance as to how impulse noise should be handled in the calculation of margin. Indeed, wide variation in reported margin has been seen across modems from different vendors when operating in the presence of impulse noise. This could lead to widely varying transmit powers, if DSM control is attempting to equalise "margin" as reported by the modems.

5.1.14 Service changes

Perhaps the most demanding issues for DSM are the level 2+ design issues for a video service which can seamlessly survive increases and decreases in available capacity. They are out of scope for this document but they must be solved for DSM to have any significant worth.

It is interesting to note that triple play also requires this seamless survival if it is to work over contended networks. So if seamless survival should prove unfeasible then triple play will rule out DSM, and will also need the backhaul networks to be reengineered to provide guaranteed capacity.

DSM might provide service benefits to an individual customer and at a particular point in time. These benefits may be of significant value to the customer, and could fit into a service definition based on an "up to X Mbps" with a lower guaranteed rate. Nevertheless, minimizing transmit power and thereby receiver margin means that the majority of subscribers may be more at risk of disturbance from extrinsic noise than is the case without DSM.

Threats to DSM Viability

5.1.15 The capacity benefit may be too small

From the simulations above one gets the impression that the apparent excess in capacity, exploiting which is the point of DSM, may not really be there. In our simulations if a set of systems is operated close to but below the capacity of the cable then they jointly back off a large amount, and so there is a tempting gap between what they each are transmitting and what their limits are. However an attempt to exploit this in any way simply causes all systems to increase power – the gap closes for virtually no increase in actual capacity.

5.1.16 The engineering may prove infeasible

The simulations suggest that an ensemble of systems becomes extremely sensitive near the saturation of the actual capacity of the cable, and approaching this capacity causes gross increases in transmit power in all the systems. When a system of the ensemble tries to access capacity beyond saturation this power swing will actually happen, and these conditions will place great demands on the engineering of the back-off (here, back-on) protocol to remain stable. We note that the power-up response of a system must be rapid in order to avoid a break in service, much faster than a conventional equalizer could track.

5.1.17 Stability Margins

In fixed rate mode, DSM fits the classical defined service model. Even using the politeness power backoff capability of VDSL margins for the most part considerably exceed the minimum target of 6 dB noise margin. Some operator even then find that service stability is enhanced if modems operate at full power all the time – with politeness power backoff disabled.

As new higher fixed rate services are introduced into the network, the operating margins are degraded, and higher frequency of service instability can be expected. In its fixed rate mode DSM has the disadvantages of reducing flexibility in data rate, and reduces network wide operating noise margins. This should be expected to lead to greater service instability.

Triple play quality of service requirements imply that only the fast internet element may be flexed to exploit variable transmission capacity, since the voice, and in particular the video elements that would have both stringent service availability requirements. A critical issue for triple play services is that re adjustment of transmitter spectrum must not disrupt the service, so it may not be feasible to perform such power optimization without co-ordination with service control and network management functions – unless the changes to power profile happen very slowly. This is not likely to be possible when also changing the data rate.

5.1.18 Operators' Concerns

5.1.18.1 How may the operator derive a benefit from DSM?

Guaranteed rate. The traditional telco business model for broadband has a service being guaranteed. Such a model offers no benefits to the telco if capacity does exceed the guarantee unpredictably. There are however considerable disbenefits if capacity falls below the guarantee, even occasionally. The telco builds in margins for its own comfort to avoid this. A benefit could only be obtained if a higher guarantee can be made.

At present broadband is typically used only for net browsing, and service is contended, so there is no guarantee on the rate of access. However if triple play is to be carried then the delay-intolerant services will want assurance their minimum capacity is served, which makes rate guarantees attractive for the future.

Average rate. If revenue can be made from level of traffic – “pay per bit” – then obviously the average capacity is significant; DSM is attractive principally because it allows revenue from the actual capacity of the line.

Headline rate. Some marketing strategies just cite a headline maximum data rate (56k anyone?) and then don't guarantee it on any particular installation. In this business model the headline rate will obviously be when all other users are DSM compliant and idle, producing a large headline number (which of course doesn't mean anything in the real world…).

5.1.18.2 How does DSM help with universality of service?

It does not. DSM is aimed at higher bit rate, not greater service reach. Indeed, by increasing the expectation of high bit rates, DSM may actually make universality harder to achieve, not easier.

5.1.18.3 Will users make their modems non conformant?

This is an issue but is not special to DSM. It is general to spectrum management. In the past it has been seen as users manually changing the gain settings of equipment, and in downloading unofficial firmware.

Separately, it seems likely that LLU operators will not be coerced to use DSM anyway, so a non DSM conformant modem cannot even be considered cheating. If this is the case 'conformant' means whatever the local (ordinary) spectrum management rules dictate.

5.1.18.4 How does one deal with complaints of unfairness?

The typical complaint is "I can only get 2 Mbps but my neighbour gets 8". It already happens in the traditional telco approach, where there are deployment limits for each service and at the range limit a line will be within the limit while a neighbour is not.

With DSM the equivalent is services working for some neighbours but not others; we should expect such inequalities to be more commonplace, since DSM exploits the minute details of each line's coupling. However the telcos' present complaints answering approach should be equally effective for this complaint.

There is of course scope for entirely new problems too, such as a service only sometimes working on a line. The operator will have to review and possibly extend their complaints answering approach.

5.1.19 Non-cooperative sources

We noted above that there will be components in the noise environment which do not co-operate with DSM, including older modems, faulty equipment, electrical noise, and LLU operators who do not want to participate. The simulations suggest, that even in a widespread DSM culture, if any of these becomes significant then the whole ensemble increases power to near the limits allowed. Any benefits of DSM would cease to be available because there is no excess margin/capacity to exploit.

5.1.20 Public acceptance of allocation process

Many of the scenarios assume multiple grades of service. When demand is less than the capacity available this is not a problem. However from time to time the users will jointly want more than the cables will carry, and some must be disappointed. In this case how does one get universal acceptance of the rationing process?

The classical solution is to price them differently, and ration by pricing (a congestion charge perhaps?), but this seems to go against LLU, and also to go against the current trend of bandwidth costs being negligible. Indeed if cost is to be an effective load control then it *must* be appreciable – by the load which is to be deterred.

The method to be used will be a matter of public policy, and it may be that the only solution simple enough to be agreed by a regulator is a single grade of service for all.

6 THE GAINS

DSM is primarily motivated by seeking more capacity, and DLCM is primarily motivated by seeking service stability, let us summarise the gains from the different sources.

DSM

If DSM is used to engineer the improvement of a single common service (i.e. improve the lot of the worst line) the simulated cases suggest typical improvement of 1 Mbps in a bandwidth of 2 MHz : an improvement of +½ bps/Hz.

If DSM is used to engineer higher rates on demand, the simulated cases suggest modest improvement for a few lines only, and at the cost of enforced hard limits on the remainder. In the simulated example the extreme cases of full capacity of the cable were:
SSM : all systems achieving 11-14 Mbps
and
DSM : all but one operated much below maximum (8 Mbps) and the chosen one at 24 Mbps.
This extreme case only gives a high rate twice the non DSM case, and the costs include artificially restricting the aspirations of all the other lines. So the gain seems rather modest compared to the effort needed to achieve it.

There will be service losses due to DSM : the operation at lowest practical power, with modest margins, means systems will be more vulnerable to external noise increases. This is true whatever the target rate of any particular system.

Capacity increases are adventitious, the actual network capacity is highly dependent on the topology of the cable plant – dendritic with a tree and branch structure in contrast to point to point topology with short tails from a single remote flexibility point. It is also impacted by the prevalence of crosstalk enhancing cable faults, and variable extrinsic noise. The capacity increases also depend upon the impact of legacy services, which large, and difficult to predict.

So it is hard for a network operator to determine what is the actual capacity that can be delivered, so the initial determination of target data rates, and how many users may be allowed extra data rate, has to be a rough estimate.

However, a presumption of DSM is that in return for some customers having their rate limited for the general good, there will be assured service at the common rate. This however may not be the case if deliberately or by accident the initial service rate is set close to or above the maximum sustainable target rate. So despite the objective of certainty, services must be designed to be resilient to rate changes[22].

---

[22] although the converse isn't automatic : if services are not viable with varying capacity then DSM is not viable,

DLCM

DLCM is not expected to give any theoretical capacity gain, but in practice does improve service rates by increasing immunity to extrinsic noise and releasing SNR margin that can be traded for increased data carrying capacity

Triple Play

If triple play proves feasible then the engineering which achieves it will provide some gains..

6.1.1 Embedded Voice

When voice is embedded into DSL the modem will use the whole bandwidth of the line, including the voice band. There is a capacity gain from the extra bandwidth, and a service gain because house wiring is not connected to the line.

Current ADSL systems typically use tone 6 and above, so do not use the bottom 25 kHz of the line – the gap between classical voiceband at 4 kHz and this is to engineer the splitter filter. The voice band can reasonably be expected to be very high quality, as noise coupling and attenuation decrease with reducing frequency. We might reasonably expect 14 bits/Hz in this band[23], that's 350 kbps. This may be available in both directions, given careful engineering of the echo canceller. (There have however been comments that the increasing group delay of the line at low frequencies may adversely affect the design requirements for the TDEQ and echo canceller, so a more realistic expectation may be 200 kb/s U/S and 0 D/S)

While 350 kbps is small compared to the 4 Mbps plus we have been talking about elsewhere, it is very significant when we consider universal service. The voice band degrades more slowly than the higher frequencies as the line gets longer.

Separately, isolating the line means any electrical pick-up by the house wiring does not interfere with modem service, neither for the premises concerned nor for other customers. This is expected to give great reduction in the incidence of REIN.

6.1.2 Service resilience to capacity changes

Triple play will require that services are resilient to capacity changes, whether DSM is used or not, because the backhaul network does not promise a minimum capacity. Hence there is no point in promising an actual line rate; so the operators don't need to build in margins. That means we may use the capacity actually present, not the fraction the telco can be confident of.

---

true, but the backhaul network must also be reengineered to guarantee capacity.

[23] a conservative figure : 56k modems achieve this (56 kbps in 4 kHz is 14 bps/Hz) and they have to also contend with a digitizer at one end which is not of their choosing. Present ADSL systems have a maximum of 15 bps/Hz, believed set by their choice of A/D converter.

In the simulation above the difference is from 7 Mbps[24] to 11-14 Mbps, which appears more of a gain than DSM[25] per se offers.

Common SSM is also potentially 'dynamic' in this way : for SSM a constant transmit power in a changing world gives a changing capacity, so if services are engineered to survive changing capacity then SSM systems can exploit them too.

Summary

Effects on:	**capacity**	**stability**
DSM	+½ bps/Hz, common service modest increases for a chosen few, heterogeneous service	Worsened : increased susceptibility to noise
DLCM	no change	Increased greatly
Embedded voice	+350 kbps both directions	Improved : decreased noise ingress
Service resilience	+2 bps/Hz	Improved

7 CONCLUSIONS

The main conclusions of this report are as follows:

- There is scope for further refinement of SSM techniques to increase service rates
- DLMC can provide substantial benefits - higher service reliability and service rates
- DSM does not offer significant benefits compared to optimised SSM with DLCM

There is still considerable scope for further innovation in static spectrum management methods. This is likely to remain central to regulation for the foreseeable future.

We see for example, that it is possible to exploit ADSL2 and ADSL2plus to provide a symmetric 2 Mb/s E1 service. As expected this requires a performance trade off that reduces asymmetric ADSL performance. It was also found that ADSL2plus offers little advantage over ADSL2 in this application.

---

[24] Even this conservative rate for a given line is only fully available if the marketing process changes to make a line-specific rate promise; conventionally only a few rates are offered ,leading to further rounding down.

[25] Recall the 11-14 Mbps rates were attained when every system was transmitting at its (ordinary) spectrum management limit, so DSM was not in use in this case.

We also find that in the case of mixed use of cabinet and exchange deployment, shared spectrum use can be managed statically to provide a better overall service mix than simple spectrum separation, though there are still service tradeoffs involved. Nevertheless, there remains the certainty that most lines can provide more capacity than the figure used as the basis for planning. So there is much flexibility for operators to trade excess margin for higher data rate, or retain it for improved service reliability.

Another factor is that commercial and marketing pressure for high headline data-rates drives exploitation of this capacity, using rate adaptive operation, but potentially at the expense of network stability.

Classical automatic DSM on the other hand requires strict adherence to a policy data rate, with the likelihood, but not the certainty, that the policy rate can be set higher than the worst case used in static spectrum management. We have considered more complex rules for DSM based on demand driven exploitation of higher data rate by a proportion of lines.

A shared cable operating in a DSM regime seems to behave as if there is a total capacity limit[26]. When the users in aggregate operate within this limit, their transmit power is reduced substantially to the point that background extrinsic noise limits line capacity. If the target data rates are set close to the aggregate limit then power increases sharply as the limit is approached. Given that it is difficult to predict the actual DSM maximum capacity it is significant that the typical transmit power is a poor indicator that substantial extra capacity is available.

A further observation is that a single loop forced to operate near maximum power by local conditions is likely to drive the ensemble towards typical transmit power that approaches the upper limit. This means that a small proportion of long lines, necessarily using full power to meet their service objectives, will tend to drive up the typical transmit power across the ensemble.

If the DSM target rate is well below the that at which the mechanism collapses, then it can provide a means of re-allocating spare capacity on a demand basis to a minority of customers, albeit with the possibility that on a particular line the request for extra capacity may fail. Simulations suggest that a ratio of maximum rate to guaranteed ensemble rate of about 2:1 is feasible, if the number of higher rate lines is strictly limited.

Embedding voice gives a capacity increase, which, while small compared to the capacities of interest when discussing DSM, may be significant when considering universal service.

Some services will need reengineering in higher layers to be able to exploit capacity as it appears and survive its loss as it disappears. Currently only data transfer is resilient. Reservation of capacity is a network QoS issue. Linking dynamic loop properties with QoS remains a significant architecture challenge. (It should also be noted that this is linked with the issue of whether DSL should support two latency paths with different PHY QoS properties (latency and error protection)).

---

[26] For an ensemble of N systems this is almost certainly a roughly convex N dimensional blob of complicated shape

The overall conclusion is that static methods will be the mainstay of spectrum management for some time to come. Dynamic methods need to operate within static PSD limitations, but within them they have considerable scope for optimizing service value. In both static and dynamic approaches the key issue for operators and regulators is how to balance service value against public policy for use of the physical infrastructure. Indeed, how to express fuzzy public policy in hard spectrum management rules remains an underlying challenge.

APPENDIX 1 - E1 OVER ADSL DATA RATE

These calculations are based in the standard AF-VTOA-0078 of the ATM Forum about ATM AAL1 Circuit Emulation Service (CES).

A structured E1 circuit consists in 8000 frames per second of 31 bytes each one. The 31 bytes correspond to the 31 timeslots without the byte needed for synchronizing the frames.

It is sent using 8 ATM cells with AAL1, where the payload is 47 bytes except in one of them where one byte is reserved to be used as a pointer.

So the total rate is:(31 x 8000 x 8) /375 = 5290.667 cells per second.

So considering that the every ATM cell is 53 bytes long, the bit data rate is::

5290 cells per second x 53 bytes per cell = 2243242.6 bps

But in ADSL2 and ADSL2+, the bit data rate is proposed to be multiple of 4000 bits/s and it is likely to be used, so the bit data rate obtained must be rounded to the next data possible:

Given $\frac{2243242.6}{4000} > 560.81$ the Allowed data rate = 561 x 4000 = 2244 kbps

So the required bit data rate at ATM layer is 2244 kbps.

APPENDIX 2 - PER-LINE FEXT MODEL

DSM seeks to exploit the individual characteristics of real lines, so simulations of it need a realistic model of the per-line characteristics. The usual Werner models are acknowledged to be pessimistic, as they aim to represent the 99%ile worst case couplings[27]. So a model for a population's crosstalk was sought.

Empirical FEXT Model

For the simulation work presented here only FEXT was of interest, and it was found to be convenient to model with the term for line attenuation removed. Call the remaining term X.

$$\text{FEXT} = X \,.\, |H|^2 \qquad X = \text{FEXT} / |H|^2$$

A 50 pair cable sample 1 km long was measured and its FEXT / $|H|^2$ statistics inspected[28]. The individual $X_{i,j}(f)$ appear to be slowly varying random functions, all different (save that measuring between the same two pair ends gets the same function in either direction). The expected f^2 trend was evident.

The values at 1 MHz appeared to be uniformly distributed between -65.5 to -40.5 dB.

So for this work the FEXT was modelled as $X_{i,j}$ is a random number uniformly distributed between -65.5 to -40.5 dB, with independent values except $X_{i,j} = X_{j,i}$.

Symmetry

In real measurements, FEXT between the end of a pair at one end of a cable section and the end of a different pair at the other end of the cable section is symmetrical : insertion loss measures the same either way round.

The coupling between the other ends of the two pairs is also symmetrical, but *different* to the first ends. However it does have strong correlation : if the pair couples strongly in one measurement then it couples strongly in the other.

So in this FEXT model we take the view $X_{i,j} = X_{j,i}$

---

[27] and simulations using them suggest little benefit for DSM

[28] the insertion losses for the lines (the $|H|^2$) were very similar, as would be expected for pairs in the same single cable section, and we do not distinguish amongst them here

Validation

The classical Werner models feature a term $N^{0.6}$, where N is the number of active pairs in the binder group. This was obtained empirically but usually puzzles students[29]. In order to explore this, the statistics $Sum_{N,j} = \sum_{i=1}^{N} X_{i,j}$ were simulated. For a given N the 99%ile strongest noise sum may then be represented by the largest value over the set of 50 lines[30]. A typical run[31] obtained these results:

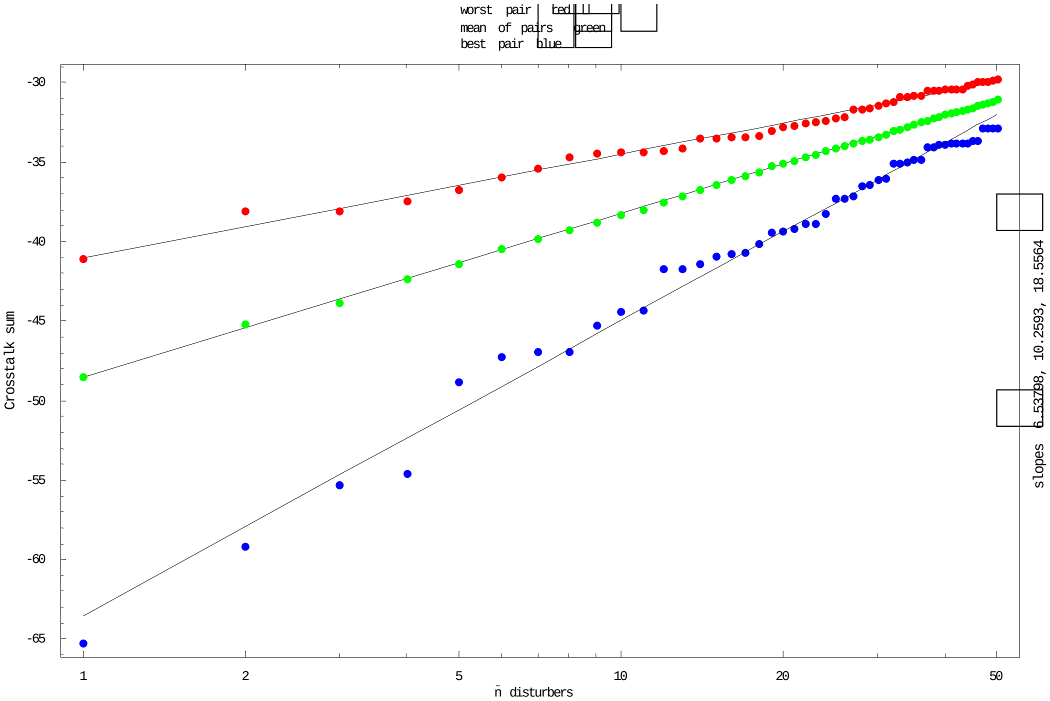

Appendix 2 Figure 1 Crosstalk power summation – worst average and best case lines

Since we are interested in power law models, the figure is plotted on log/log axes. The top set of points is the simulation of the worst affected line, the bottom set is the best line and the middle set is the power average for the lines. The straight lines are best fits. The slopes of those lines are 6.5, 10.3, and 18.6 dB/frequency decade – power laws of $N^{0.65}$, $N^{1.03}$, and $N^{1.86}$ respectively.

---

[29] because it appears to violate the conservation of energy.
(From conservation considerations we would expect a model for *mean* FEXT to have a term N^1)

[30] in a cumulative distribution the 'largest of 50' is the step from 98%ile to 100%ile; it is usual to take the value to represent the middle of the step at 99%ile

[31] the runs being different in a simulation using randomness

Theory suggests the middle set's power law should have an exponent of exactly 1. Repeated runs of the simulation gave a variety of values scattered around 1, so the difference here suggests sampling error.

The top set of points corresponds to the Werner 99%ile worst case, and its line slope is comfortingly close to Werner's exponent of 0.6. Repeated runs gave values scattered from 0.6 to 0.75.

The bottom set of points is for interest, since we have no specific interest in a model for 1% worst (=99%ile best) case.

We conclude this random model credibly emulates the $N^{0.6}$ behaviour of Werner's model.

Shortcomings

In the model described here each coupling is independent of frequency, so simulations using it cannot explore spectral variations. Real couplings seem to be random functions of frequency; in which we note that the stronger couplings seem to vary slower than the weaker couplings.

Use

In simulations of capacity SNR is the main interest. The model here excludes line attenuation, and by also omitting it from the transmission path SNR may still be calculated correctly.

$$SNR_i = {S_i} \Big/ {\sum_{j \neq i} S_j X_{i,j}}$$

The simulation at present uses the frequency independence to simplify calculation of capacity;

capacity [Mbps] = 2 log2(1 + (SNR – 13 dB))

in a 2 MHz channel (and it doesn't matter which 2 MHz band)
and supposing 13 dB gap plus margin