Iterative Spectrum Balancing for Digital Subscriber Lines

Raphael Cendrillon, Marc Moonen
Department of Electrical Engineering
Katholieke Universiteit Leuven, Belgium
cendrillon@ieee.org

***Abstract*— Dynamic spectrum management (DSM) is an important technique for mitigating crosstalk in DSL. One of the first DSM algorithms proposed, Iterative waterfilling (IW), has a low complexity and demonstrates the spectacular performance gains that are possible. Unfortunately IW tends to be highly suboptimal in mixed CO/RT deployments and upstream VDSL. Another DSM algorithm, Optimal spectrum balancing (OSB), uses a weighted rate-sum to find the theoretically optimal transmit spectra. Unfortunately its complexity scales exponentially with the number of lines in the binder N. Typical binders contain 25-100 lines, for which OSB is intractable. This paper presents a new iterative algorithm for spectrum management in DSL. The algorithm optimizes the weighted rate-sum in an iterative fashion, which leads to a quadratic, rather than exponential, complexity in N. The algorithm is tractable for large N and can be used to optimize entire binders. Simulations show that the algorithm performs very close to the theoretical optimum achieved by OSB.**

I. Introduction

Crosstalk is a major issue in modern DSL systems such as ADSL and VDSL. Typically 10-20 dB larger than the background noise, crosstalk is *the* dominant source of performance degradation.

Crosstalk cancellation is one possible solution and can remove crosstalk completely with minimal noise enhancement[1][2][3]. Unfortunately, in many scenarios crosstalk cancellation is inapplicable as a result of unbundling, mixed *central office* (CO) / *remote-terminal* (RT) deployment, or complexity constraints. In this case crosstalk must be mitigated through spectrum management.

Dynamic spectrum management (DSM), a new paradigm, designs the spectra of each modem to match the specific topology of the network[4]. These spectra are adapted based on the direct and crosstalk channels seen by the different modems. With DSM each modem attempts to achieve its desired datarate whilst causing as little disturbance as possible to the other modems in the network. [1]

Iterative waterfilling (IW) was one of the first DSM algorithms proposed and demonstrates the spectacular performance gains that are possible[6]. IW has a complexity that scales linearly with the number of lines in the binder N, an important quality since a full binder typically contains 25-100 lines. Unfortunately, since IW is based on a greedy algorithm, it converges to the selfish-optimum. This tends to be highly suboptimal in near-far scenarios such as mixed CO/RT deployments and upstream VDSL.

To address this, the *optimal spectrum balancing* (OSB) algorithm was proposed[7]. This algorithm is provably optimal and achieves the best possible balance between the rates of the different modems in the network, allowing operation at any point on the rate region boundary. OSB is based on a weighted rate-sum, which forces each modem to account for the damage done to other modems in the network when deciding on its own transmit spectra. This allows the selfish-optimum to be avoided and leads to significantly improved performance[7].

Unfortunately OSB is a centralized algorithm, requiring all PSDs to be calculated jointly at a centralized *spectrum management center* (SMC). Furthermore, it has a complexity that scales exponentially with N, which makes it computationally intractable for use with more than 5-6 lines.

This paper presents a new iterative algorithm for spectrum management in DSL. Like OSB the algorithm is based on a weighted rate-sum, which makes it possible to avoid the selfish-optimum. However unlike OSB the optimization of the weighted rate-sum is implemented in an iterative fashion, which leads to a quadratic, rather than exponential, complexity in N. The resulting algorithm is computationally tractable for large N and, as will be shown, leads to near-optimal performance.

The price to pay for this improved performance is the loss of some autonomy. Each modem must have knowledge of the noise PSDs and crosstalk channels of all modems in the binder, which was not necessary in IW. This increases the overhead required for communication with the SMC. However, since the twisted-pair channel is slowly time-varying, the additional overhead is minimal.

This work was carried out in the frame of IUAP P5/22, *Dynamical Systems and Control: Computation, Identification and Modelling* and P5/11, *Mobile multimedia communication systems and networks*; the Concerted Research Action GOA-MEFISTO-666, *Mathematical Engineering for Information and Communication Systems Technology*; IWT BANITS Project, *Solutions for xDSL Interoperability, Deployment and New Technologies*; FWO Project G.0196.02, *Design of efficient communication techniques for wireless time-dispersive multi-user MIMO systems* and was partially sponsored by Alcatel-Bell.

[1] Recent standardization activities [5] have broadened the scope of DSM to include not only spectra adaptation, but also signal processing algorithms such as crosstalk cancellation and vectoring. This paper uses the term DSM in the narrowly defined sense.

II. System Model

This paper only considers DSM as applied to DMT modulated modems. Whilst some form of DSM can also be applied to single carrier modems it often leads to inferior performance since dynamic shaping of the transmit spectra is not possible. As such it is assumed that any non-DMT systems form part of the background noise. Assuming that *discrete multi-tone* (DMT) modulation is employed, transmission can be modelled independently on each tone

$$\mathbf{y}_k = \mathbf{H}_k \mathbf{x}_k + \mathbf{z}_k. \tag{1}$$

The vector $\mathbf{x}_k \triangleq \left[x_k^1, \cdots, x_k^N\right]$ contains transmitted signals on tone k. There are N lines in the binder and x_k^n is the signal transmitted onto line n at tone k. $\mathbf{y}_k$ and $\mathbf{z}_k$ have similar structures. $\mathbf{y}_k$ is the vector of received signals on tone k. $\mathbf{z}_k$ is the vector of additive noise on tone k and contains thermal noise, alien crosstalk, single-carrier modems, RFI etc. Recall that $1 \leq k \leq K$ where K is the number of tones within the system. We denote the noise PSD on line n as $\sigma_k^n \triangleq \mathcal{E}\left\{|z_k^n|^2\right\}$. $\mathbf{H}_k$ is the $N \times N$ channel transfer matrix on tone k. $h_k^{n,m} \triangleq [\mathbf{H}_k]_{n,m}$ is the channel from TX m to RX n on tone k. The diagonal elements of $\mathbf{H}_k$ contain the direct-channels whilst the off-diagonal elements contain the crosstalk channels. We denote the transmit PSD $s_k^n \triangleq \mathcal{E}\left\{|x_k^n|^2\right\}$. For convenience we denote the vector containing the PSD of user n on all tones as $\mathbf{s}_n \triangleq [s_1^n, \ldots, s_K^n]$. We denote the tone spacing as Δ_f and DMT symbol rate as f_s.

In (1) it is implicitly assumed that each tone operates independently and free from inter-carrier interference (ICI). If the modems in the network are not synchronized this will not strictly be the case. Sidelobes from the FFT operation in the receiver will give rise to ICI between users. The use of transmit and receive windowing can help to suppress the sidelobes and minimize their impact on performance[8]. Transmit and receive windowing is mandatory in VDSL standards[9] and appears as an option in ADSL standards[10]. In the remainder of this paper it is assumed that the modems are either synchronization or windowing is employed such that sidelobe effects are negligible.

It is assumed that each modem treats interference from other modems as noise. When the number of interfering modems is large the interference is well approximated by a Gaussian distribution. Under this assumption the achievable bitloading of user n on tone k is

$$b_k^n \triangleq \log_2\left(1 + \frac{1}{\Gamma}\frac{|h_k^{n,n}|^2 s_k^n}{\sum_{m\neq n}|h_k^{n,m}|^2 s_k^m + \sigma_k^n}\right), \tag{2}$$

where Γ denotes the SNR-gap to capacity, which is a function of the desired BER, coding gain and noise margin[11]. The data-rate on line n is thus

$$R_n = f_s \sum_k b_k^n.$$

In practice the relationship between the received *signal-to-interference-plus-noise ratio* (SINR) and the bitrate may be more complex as it is dependent on the coding scheme employed within the modem. In this paper (2) will be used for simplicity however the algorithms presented can be applied to any arbitrary function that relates the bitloading to the SINR on each tone.

Each modem is typically subject to a *total power constraint*

$$\Delta_f \sum_k s_k^n \leq \overline{P}_n, \ \forall n,$$

where $\overline{P}_n$ denotes the total power that modem n can transmit. This arises from limitations on each modem's analog front-end. For convenience this is reformulated as

$$\sum_k s_k^n \leq P_n, \ \forall n, \tag{3}$$

where $P_n \triangleq \overline{P}_n/\Delta_f$

III. The Spectrum Management Problem

The spectrum management problem is defined as

$$\begin{aligned} \max_{\mathbf{s}_1,\ldots,\mathbf{s}_N} R_1 \quad \text{s.t.} \quad & R_n \geq R_n^{\text{target}}, \forall n > 1, \\ \text{s.t.} \quad & \sum_k s_k^n \leq P_n, \forall n, \end{aligned} \tag{4}$$

where R_n^{target} denotes the target data-rate for the nth modem. The total power and target data-rate constraints cause the optimization (4) to be coupled across frequency. Furthermore, since the data-rate constraints form a non-convex set, solving (4) directly results in an exponential complexity in the number of tones K. Since $K = 256$ in ADSL, and $K = 4096$ in VDSL, this leads to a computationally intractable problem.

Following the approach of [7], the original optimization (4) is replaced with the Dual Problem

$$\max_{\mathbf{s}_1,\ldots,\mathbf{s}_N} J(\mathbf{s}_1, \ldots, \mathbf{s}_N), \tag{5}$$

where

$$J(\mathbf{s}_1, \ldots, \mathbf{s}_N) \triangleq \sum_n w_n R_n - \sum_n \sum_k \lambda_n s_k^n.$$

The weight for the first user w_1 is set to unity. This will maximize the rate of the first user subject to the target rate constraints on the other users. In fact, the choice of w_1 is arbitrary and any constant, positive value will achieve the same effect. The Lagrangian multipliers λ_n and w_n are chosen such that the KKT conditions are satisfied

$$\lambda_n \left(P_n - \sum_k s_k^n\right) = 0, \forall n, \tag{6}$$

$$w_n \left(R_n - \sum_k b_k^n\right) = 0, \forall n. \tag{7}$$

Provided that these conditions hold, the dual problem (5) is equivalent to the original optimization (4) and also yields the optimal transmit spectra. However, unlike (4), the dual

Algorithm 1 Optimal Spectrum Balancing

repeat
  for each k: $\left(s_k^1, \ldots, s_k^N\right) = \arg\max_{s_k^1,\ldots,s_k^N} J_k$;
    (solve by N-D exhaustive search)
  for each n: $w_n = [w_n + \epsilon\,(R_n^{\text{target}} - \sum_k b_k^n)]^+$;
  for each n: $\lambda_n = [\lambda_n + \epsilon\,(\sum_k s_k^n - P_n)]^+$;
until convergence

problem can be decomposed into a set of sub-problems that are decoupled across frequency. The sub-problem on tone k is

$$\max_{s_k^1,\ldots,s_k^N} J_k(s_k^1, \ldots, s_k^N), \tag{8}$$

where

$$J_k(s_k^1, \ldots, s_k^N) \triangleq \sum_n w_n f_s b_k^n - \sum_n \lambda_n s_k^n. \tag{9}$$

Note that $\sum_k J_k = J$, so maximizing the sub-problems is equivalent to maximizing the dual problem.

IV. Optimal Spectrum Balancing

Since J_k is non-convex, it must be solved through an exhaustive search of $(s_k^1, \ldots, s_k^N)$. Define the granularity in the transmit PSD as Δ_s. This results from the limited accuracy of each modem's AFE. In current standards Δ_s is set to 0.5 dBm/Hz[12]. The value of s_k^n can now be limited to the set $\{0, \Delta_s, \ldots, P_n\}$. A such, the exhaustive search of $(s_k^1, \ldots, s_k^N)$ has a complexity $\mathcal{O}\left(\left(P_n/\Delta_s\right)^N\right)$. So the KN-dimensional, non-convex optimization (4), can be solved through a set of K decoupled non-convex optimizations (8), each of dimension-N. This allows the spectrum management problem to be solved with $\mathcal{O}\left(K \exp(N)\right)$ complexity, instead of $\mathcal{O}\left(\exp(KN)\right)$. So for small N the spectrum management problem becomes tractable. This is the basis behind the OSB algorithm, which is listed as Alg. 1, where the function $[x]^+ \triangleq \max(0, x)$[7][13].

Despite this complexity reduction, due to its exponential complexity in N, for large N OSB is still intractable. This prevents the direct implementation of OSB since binders typically contain 25-100 lines.

V. Iterative Spectrum Balancing

OSB is intractable for large N. To address this problem, we now present an iterative algorithm that is tractable for large N. Like OSB this algorithm is based on a weighted rate-sum, which allows the selfish-optimum to be avoided. However, the weighted rate-sum optimization is implemented in an iterative fashion as is now described.

In OSB the transmit PSDs are searched jointly (8), which leads to an exponential complexity in N. This is why OSB is intractable for large N. An alternative approach is to search the PSDs of each user in an iterative fashion. The PSD of each user is updated one at a time. When updating the PSD of user n, the PSDs of all other users are fixed at their present values. The optimization is then

$$\max_{s_k^n} J_k(s_k^1, \ldots, s_k^N). \tag{10}$$

Algorithm 2 Iterative Spectrum Balancing

repeat
  for $n = 1 \ldots N$
    repeat
      for each k: fix s_k^m, $\forall m \neq n$, then
        $s_k^n = \arg\max_{s_k^n} J_k$;
        (solve by 1-D exhaustive search)
      $w_n = [w_n + \epsilon\,(R_n^{\text{target}} - \sum_k b_k^n)]^+$;
      $\lambda_n = [\lambda_n + \epsilon\,(\sum_k s_k^n - P_n)]^+$;
    until convergence
  end
until convergence

The algorithm iterates through the users, optimizing the PSD of each user in turn. The complete *iterative spectrum balancing* (ISB) algorithm is listed as Alg. 2.

The algorithm consists of an outer loop and an inner loop. In the inner loop the PSD of user n is optimized. In a similar fashion to OSB, the update of each user's PSD is based on a weighted rate-sum (9), which allows the selfish-optimum of IW to be avoided. However, unlike OSB, the optimization is only done on the PSD of a single user. So the N-dimensional exhaustive search is replaced by a 1-dimensional exhaustive search. This leads to a complexity which is quadratic, rather than exponential, in N.

The inner loop also updates the Lagrangian multipliers λ_n and w_n. The update rule for w_n, based on sub-gradient descent, is

$$w_n = \left[w_n + \epsilon\left(R_n^{\text{target}} - \sum_k b_k^n\right)\right]^+.$$

Constraints are added to ensure w_n remains positive. One can interpret w_n as the priority given to user n in the optimization. If the data-rate of user n is below its target, then w_n is increased to allocate more priority to user n. The process is repeated until user n achieves its target rate, or $w_n = 0$. This defines the KKT condition (7). Effectively user n chooses the least possible priority w_n required to achieve his target rate, thereby minimizing the disturbance caused to the other modems in the network.

Similarly the update rule for λ_n is

$$\lambda_n = \left[\lambda_n + \epsilon\left(\sum_k s_k^n - P_n\right)\right]^+.$$

Constraints are added to ensure λ_n remains positive. One can interpret λ_n as the price for power. If user n is below its total power budget, then the price for power is decreased and user n will be allocated more power. The process is repeated until user n reaches its power budget, or $\lambda_n = 0$. This defines the KKT condition (6).

Algorithm 3 Iterative Waterfilling

repeat
  for $n = 1 \dots N$
    repeat
      $w_n = 1;\ w_m = 0,\ \forall m \neq n;$
      for each k: fix $s_k^m,\ \forall m \neq n$, then
        $s_k^n = \arg\max_{s_k^n} J_k$
        (J_k convex: solve in closed form)
      if $\sum_k s_k^n > P_n$, then $\lambda_n = [\lambda_n + \epsilon\left(\sum_k s_k^n - P_n\right)]^+$,
        else $\lambda_n = [\lambda_n + \epsilon\left(\sum_k b_k^n - R_n^{\text{target}}\right)]^+$;
    until convergence
  end
until convergence

The outer loop of Alg. 2 repeats the inner loop for each n, optimizing the PSD of each user in turn. The outer loop terminates when the PSDs of the users converge.

The total complexity of ISB is $\mathcal{O}\left(KN^2\left(P_n/\Delta_s\right)\right)$. For comparison the complexity of OSB is $\mathcal{O}\left(KN\left(P_n/\Delta_s\right)^N\right)$, and the complexity of IW is $\mathcal{O}\left(KN\right)$.

VI. Iterative Waterfilling

The IW algorithm is listed as Alg. 3. The essential difference between IW and ISB is that ISB makes use of a weighted rate-sum. In IW each user greedily tries to maximize their own data-rate. To ensure a fair-allocation of rates, the outer loop of the IW algorithm decreases the power of each user, through the waterfilling level, to ensure that they do no exceed their target rate. By setting the target rates on the near-end modems sufficiently low, a degree of protection can be afforded to the far-end modems.

The short-fall of IW is that it tries to protect the weaker users indirectly through the power constraint Lagrangians λ_n. Consider the IW PSD of user n

$$s_k^n = \left[\frac{1}{\lambda_n} - \frac{\sum_{m\neq n}\left|h_k^{n,m}\right|^2 s_k^m + \sigma_k^n}{\left|h_k^{n,n}\right|^2}\right]^+. \qquad (11)$$

Consider the case when user n is the only near-end user in the binder. Since user n is near-end, it will experience negligible crosstalk from the far-end users in the binder. So the IW PSD (11) is well approximated by the single-user waterfilling PSD

$$s_k^n \simeq \left[\frac{1}{\lambda_n} - \frac{\sigma_k^n}{\left|h_k^{n,n}\right|^2}\right]^+. \qquad (12)$$

From (12) it can be seen that decreasing the waterfilling level from λ_n^{-1} to $\widetilde{\lambda}_n^{-1}$ causes the PSD level to decrease by $\lambda_n^{-1} - \widetilde{\lambda}_n^{-1}$ on all tones. So with IW the degree of *power back-off* (PBO) is always constant with frequency. This is a significant limitation since crosstalk coupling varies dramatically with frequency. For optimal performance the degree of PBO should adapt to match the severity of the crosstalk coupling on each specific tone.

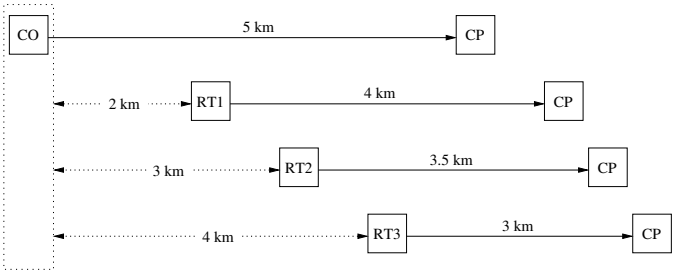


Fig. 1. 4 User Scenario

IW is not capable of implementing such frequency variable PBO because the penalty for loading power λ_n is not frequency selective. This is because IW incorrectly tries to use the power constraint Lagrangian λ_n to play the role of the target rate constraint Lagrangians $w_1, \dots, w_N$. In contrast to λ_n, the use of w_n in the weighted rate-sum optimization allows the PBO to vary with frequency; it explicitly takes into account the disturbance caused to other modems on the network when optimizing the PSD of each user. As is shown in the next section, the result is significantly improved performance for ISB over IW.

VII. Performance

This section evaluates the performance of ISB in downstream ADSL. For all scenarios the line diameter is 0.5 mm (24-AWG). The target symbol error probability is 10^{-7} or less. The coding gain and noise margin are set to 3 dB and 6 dB respectively. The PSD granularity $\Delta_s = 0.5$ dBm/Hz, the tone spacing $\Delta_f = 4.3125$ kHz and the DMT symbol rate $f_s = 4$ kHz. A maximum transmit power of 20.4 dBm applies to each modem[10]. Background noise includes crosstalk from 16 ISDN, 4 HDSL, and 10 non-DSM capable ADSL disturbers which transmit at a spectral mask of -60 dBm/Hz[10].

Comparison is made with the optimal, but highly complex, OSB algorithm, and the lower complexity algorithm IW. Flat PBO is also included for comparison, and consists of each user transmitting at the minimal possible PSD required to support their target rate. Flat PBO gives an idea of the rates that can be achieved with existing ADSL transceivers, and is subject to a spectral mask of -60 dBm/Hz. Spectral masks are not applied to IW, ISB or OSB.

A. 4 User Scenario

The first scenario consists of a mixed CO/RT deployment. A 4 user scenario has been selected to make a comparison with the OSB algorithm possible, since for $N > 4$ OSB becomes extremely complex. As depicted in Fig. 1 the scenario consists of one 5 km CO distributed line, and 3 RT distributed lines: RT1, RT2 and RT3. The RTs are located at 2 km, 3 km and 4 km from the CO respectively. The corresponding line lengths are 4 km, 3.5 km and 3 km.

The target rates on RT1 and RT2 have both been set to 2 Mbps. For a variety of different target rates on RT3, the CO attempted to maximize its own data-rate either by transmitting

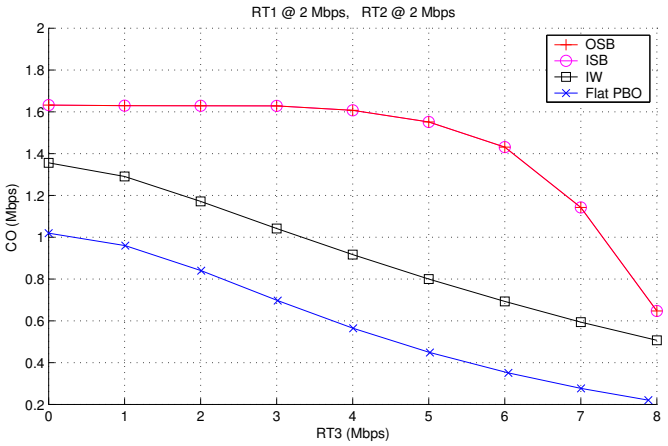


Fig. 2. Rate Region - 4 User Scenario

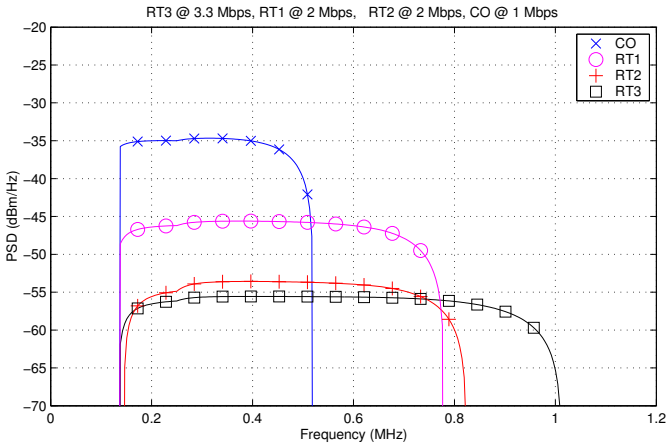


Fig. 3. Iterative Waterfilling (IW) PSDs

Method	RT3 (Mbps)	CO (Mbps)
Flat PBO	0.3	1.0
IW	3.3	1.0
ISB	7.3	1.0
OSB	7.3	1.0

TABLE I
RATE COMPARISON - 4 USER SCENARIO (RT1-2 @ 2 MBPS)

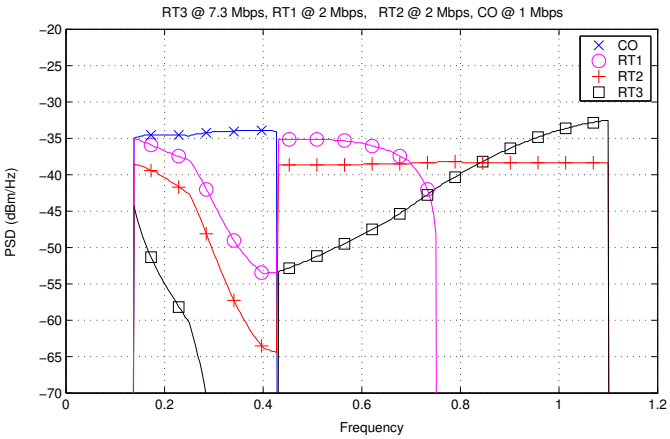


Fig. 4. Iterative Spectrum Balancing (ISB) PSDs

at full power in IW, or by setting its corresponding weight w_{co} to unity in ISB and OSB. This produced the rate regions shown in Fig. 2. Each rate region shows that rate combinations that are achievable with a given algorithm.

The rate regions in Fig. 2 show the substantial gains that ISB achieves over IW. For example, consider the case when a minimum service of 1 Mbps must be provided to the CO line. Fig. 2 shows that with IW the maximum achievable rate on RT3 is then 3.3 Mbps. Compare this with ISB where the rate on RT3 can be increased to 7.3 Mbps whilst still maintaining 1 Mbps on the CO line. So the achievable rate on RT3 can be doubled through the use of ISB. These results are summarized in Tab. I.

The corresponding PSDs are shown in Fig. 3 for IW and Fig. 4 for ISB. The PSDs from OSB are not shown since they are nearly identical to those from ISB. Note that with IW the PBO on the RTs is flat with frequency, as discussed in Sec. VI. Contrast this with ISB where the PBO varies dramatically with frequency. Crosstalk coupling is minimal at low frequencies so with ISB the RTs transmit at full power on the lower tones. As frequency increases the RTs reduce their power to protect the CO. The level of PBO increases with the nearness of an RT's transmitter to the receiver of the CO line. At 430 kHz the CO line becomes inactive due to poor channel-SNR. Above this frequency the CO line no longer needs to be protected and the PSDs of the RTs increase abruptly. RT3 still does some PBO to protect RT1. At 750 kHz RT1 becomes inactive due to poor channel-SNR on its line. As a result the PSD on RT3 increases again.

It should be clear that optimal performance requires PBO that varies with frequency. ISB adapts the transmit spectra to match the crosstalk coupling strength and the type of active users on each particular tone. This leads to a large performance gain over IW, which can only implement frequency flat PBO.

Note that, as the ISB and OSB rate region coincide in Fig. 2, ISB gives close to optimal performance in this scenario. After simulating ISB in a broad range of scenarios, it appears to be near-optimal in general. A detailed study of why ISB yields near-optimal performance is an important area for future work. We postulate that this is due to the hierarchal structure of crosstalk, by which we mean that far-end users do not cause substantial crosstalk to near-end users. For example, in this scenario the CO causes significant interference to no-one, and RTn only causes significant interference to the CO and RTm, $\forall m < n$. This appears to enable an iterative, user-by-user line-search to converge to the globally optimal solution. Simulations also show ISB to be near-optimal in VDSL.

Method	RT1-9 (Mbps)	CO (Mbps)
IW	0.6	1.5
ISB	2.0	1.5

TABLE II
RATE COMPARISON - 10 USER SCENARIO

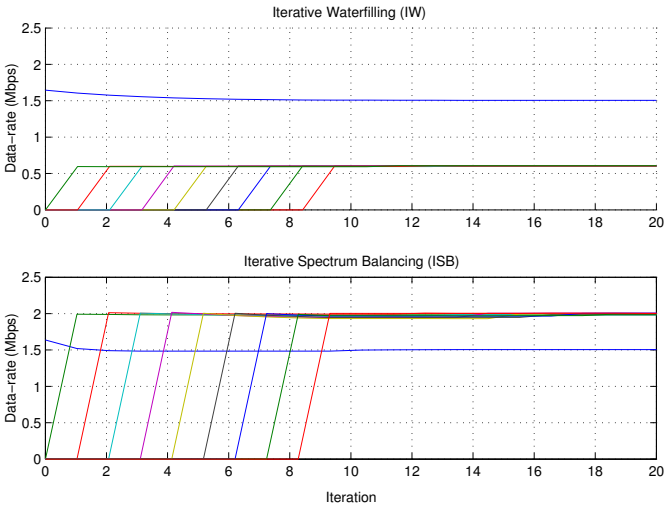


Fig. 5. Rate Convergence - 10 User Scenario

B. 10 User Scenario

A 10 user scenario has been simulated to evaluate the convergence of ISB and performance in large networks. This scenario consists of a 5 km CO distributed line and 9 RT distributed lines. The RTs are located at 2 km, 2.25 km, . . ., 4 km from the CO consecutively. The corresponding line lengths are 4.5 km, 4.1875 km, . . ., 2 km.

The target rates on the RTs were equally set, and chosen such that a minimum service of 1.5 Mbps could be achieved on the CO. With IW the RTs could achieve 0.6 Mbps. ISB increased this to 2 Mbps, whilst still ensuring a 1.5 Mbps service on the CO line. So again the achievable rate on the RTs can be doubled through the use of ISB. These results are summarized in Tab. II.

Fig. 5 shows the data-rate convergence of IW and ISB. As can be seen, the data-rates of both algorithms converge within $2N$ iterations. After simulating ISB in a broad range of scenarios, such rapid convergence appears to be the norm. Convergence behavior was also seen to be independent of the iteration order amongst users.

VIII. CONCLUSIONS

Dynamic spectrum management (DSM) is an important technique for mitigating crosstalk in DSL. One existing algorithm, iterative waterfilling (IW), is simple but converges to the selfish-optimum, which leads to poor performance in near-far scenarios. Another algorithm, optimal spectrum balancing (OSB), gives optimal performance but is computationally intractable for networks with many users.

This paper presented a new iterative algorithm for spectrum management in DSL. Like OSB the algorithm makes use of a weighted rate-sum to avoid the selfish-optimum. However, unlike OSB the optimization of the weighted rate-sum is implemented in an iterative fashion, which leads to a tractable complexity even with a large number of users.

Simulations show that this algorithm leads to near-optimal performance in a large number of scenarios. In mixed CO/RT distributions, the proposed algorithm often achieves double the data-rate of IW.

Unlike IW, the proposed algorithm is non-autonomous, requiring knowledge of the crosstalk channels in the network. An important area for future work is the development of a fully autonomous algorithm with near-optimal performance.

REFERENCES

[1] R. Cendrillon, M. Moonen, T. Bostoen, and G. Ginis, "The Linear Zero-Forcing Crosstalk Canceller is Near-optimal in DSL Channels," accepted for publication in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)* 2004.

[2] R. Cendrillon, G. Ginis, and M. Moonen, "Improved Linear Crosstalk Precompensation for Downstream VDSL," in *Proc. IEEE Int. Conf. on Acoust., Speech and Sig. Processing (ICASSP)*, vol. 4, Montreal, May 2004, pp. 1053–1056.

[3] G. Ginis and J. Cioffi, "Vectored Transmission for Digital Subscriber Line Systems," *IEEE J. Select. Areas Commun.*, vol. 20, no. 5, pp. 1085–1104, June 2002.

[4] J. Cioffi, *DSL Advances*. Prentice Hall, 2002, ch. 11 - Dynamic Spectrum Management.

[5] *Dynamic Spectrum Management*, ANSI Draft Std. T1E1.4/2003-018, Rev. 15, 2004.

[6] W. Yu, G. Ginis, and J. Cioffi, "Distributed Multiuser Power Control for Digital Subscriber Lines," *IEEE J. Select. Areas Commun.*, vol. 20, no. 5, pp. 1105–1115, June 2002.

[7] R. Cendrillon, W. Yu, M. Moonen, J. Verlinden, and T. Bostoen, "Optimal Multi-user Spectrum Management for Digital Subscriber Lines," in *Proc. IEEE Int. Conf. on Commun. (ICC)*, Paris, June 2004, pp. 1–5.

[8] F. Sjoberg, M. Isaksson, P. Borjesson, *et al.*, "Zipper: A Duplex Method for VDSL Based on DMT," *IEEE Trans. Commun.*, vol. 47, no. 8, pp. 1245–1252, Aug. 1999.

[9] *Very-high bit-rate Digital Subscriber Lines (VDSL) Metallic Interface*, ANSI Std. T1.424, 2004.

[10] *Asymmetrical digital subscriber line transceivers 2 (ADSL2)*, ITU Std. G.992.2, 2002.

[11] T. Starr, J. Cioffi, and P. Silverman, *Understanding Digital Subscriber Line Technology*. Prentice Hall, 1999.

[12] *Physical layer management for Digital Subscriber Line (DSL) transcievers*, ITU Std. G.997.1, 2003.

[13] W. Yu, R. Lui, and R. Cendrillon, "Dual Optimization Methods for Multiuser OFDM Systems," accepted for publication in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)* 2004.