Designing Tone Reservation PAR Reduction

Niklas Andgart,[1] Per Ödling,[1] Albin Johansson,[2] and Per Ola Börjesson[1]

[1] *Signal Processing Group, Department of Information Technology, Lund University, Box 118, SE-221 00 Lund, Sweden*
[2] *Ericsson AB, 126 25 Stockholm, Sweden*

Tone reservation peak-to-average (PAR) ratio reduction is an established area when it comes to bringing down signal peaks in multicarrier (DMT or OFDM) systems. When designing such a system, some questions often arise about PAR reduction. Is it worth the effort? How much can it give? How much does it give depending on the parameter choices? With this paper, we attempt to answer these questions without resolving to extensive simulations for every system and every parameter choice. From a specification of the allowed spectrum, for instance prescribed by a standard, including a PSD-mask and a number of tones, we analytically predict achievable PAR levels, and thus implicitly suggest parameter choices. We use the ADSL2 and ADSL2+ systems as design examples.

1. INTRODUCTION

With *discrete multitone modulation* (DMT) as the dominating modulation scheme in *digital subscriber line* (DSL) systems, there is a problem with high signal amplitudes. This is caused by several independent sequences adding up to a signal that approximately will adhere to a Gaussian distribution and is commonly referred to as a high *peak-to-average ratio* (PAR). Several methods have been presented to alleviate this problem [1–7].

We focus on the tone reservation method, which has been presented in [1, 2], with further improvements in [8–12]. Following the constraints set up by the standards, the achievable performance is limited, and can be determined by mathematical analysis in combination with some sound engineering assumptions. Construction of a system where the designer is unaware of the limitations will likely lead to a severe violation of the *power spectral density* (PSD) mask, or to a worse performance than what could be expected. This is illustrated in Figure 1, where the target PAR level is of significant importance. Aiming at a too low PAR level will lead to a violation of the PSD, or to a much worse result if the PSD is somehow enforced. In this paper we explain this relationship and develop means to predict what can be done when applying tone reservation PAR reduction to a practical DMT system. The aim is to produce results that are valid without having to run extensive simulations for each individual case. Hence, we will look at a number of bounds and engineering approximations that will tell us what can be done in the complete system.

When a system designer is contemplating whether it is worthwhile to include PAR reduction in a system or not, it normally requires a lot of work to develop a simulation chain in order to evaluate the potential gain. With the results presented here, the "worthwhile or not" question can easily be answered in an afternoon by a skilled engineer. A simulation chain would then only be developed when needed, that is, for the precise determination of PAR reduction parameter values. A practicing engineer, who only at a later stage would like to enjoy the theory, could for now skip reading Sections 2 and 3, and move directly to Section 4.

Earlier work has discussed the existence and effect of a PSD bound for tone reservation [13] and algorithms suitable for implementing it [12]. This paper aims at explaining in which situations this PSD constraint is an issue. It also intends to show what levels on PSD constraints and expected PAR reduction performance that can be used in a system design.

The paper starts out with discussing practical standardised systems and what would be proper engineering assumptions. In Section 2, the system and its requirements are defined and a set of theoretical results necessary for the analysis is given. Thereafter, Section 3 analyses what impact this has on reduction performance. The results from this section are summarised in Section 4, where practical instructions for using them are given in a "how-to" style. Section 5 applies the results to an ADSL2 system and also extends the analysis to include ADSL2+ systems.

2. A SYSTEM IN PRACTICE

The Gaussian distribution of the transmit signal [14] implies a possibility of very high peak amplitudes. This may lead to that the signal is clipped or, if the amplitude span of the line driver is increased, high power dissipation. This is especially the case with the commonly used class AB line drivers, where the power dissipation in the line driver is directly dependent on the supply voltage. Notably, the line driver, or power amplifier, is responsible for the major part[1] of the total power consumption in many communication systems, for example, DSL systems [15].

2.1. PAR reduction in standardized systems

As a first example system, we will look at an ADSL2 system using an FFT size of 512 in the downstream direction. Since the lowest part of the frequency band, covering the first 32 tones, is used for analogue telephony and upstream transmission, only tones 33–255 are available for downstream data transmission. These parameters are the same as for the well-established ADSL1 system [16], with the ADSL2 standard [17] more closely defining the requirements on subcarriers that could be used for PAR reduction. We quote the most relevant text in Appendix A. The ADSL2 standard is also a base for the ADSL2+ standard [18], which follows similar specifications. The main difference is the double downstream bandwidth (tones 33–511), and that ADSL2+ systems thereby operate with an FFT size of 1024 in the downstream direction.

In the ADSL2 standard, the PSD mask on the reduction tones is set to −10 dB relative to the PSD for the data tones [17], see also Appendix A. The question now is how the PSD should be properly measured. Since the PSD commonly is averaged over time, the instantaneous PSD may be allowed to vary between symbols, with certain symbols exceeding the average PSD constraint. We will consider two extreme cases of averaging time, as well as a reasonable intermediate point.

One extreme would be to average over a long time. Thereby, we could sometimes allow high instantaneous PSD levels, if we most of the time use little or no power. This interpretation of the definition restricts only the average level, without imposing any restrictions on individual symbols. This can generate occasional large reduction signals, with the average still being below the limit. There are reasons to why this may be undesirable, for instance, large amplitudes on the reduction tones could generate intermodulation interference to neighbouring data-carrying subcarriers by exciting nonlinearities in the line driver. Other users or other systems in the same cable bundle can also be affected through crosstalk. Although this interpretation of the PSD limitation is standard compliant, it is neither neighbour friendly, nor necessary according to the spirit of the standard.

The other extreme is when the PSD of each particular symbol has to conform with the −10 dB limitation. This
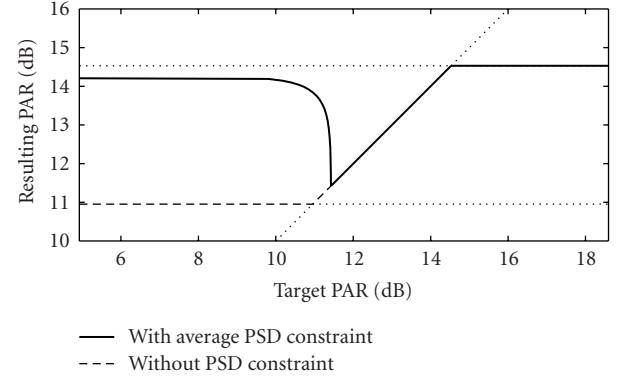

FIGURE 1: Relationship between target PAR and the resulting PAR level for reduction with 6 tones, with and without constraints on average reduction PSD. The narrow optimum and its lack of tolerance for design errors illustrates the nonintuitive behaviour of PAR reduction under PSD constraints. It also emphasizes the need of design methods such as those given in this paper.

results in a very strict peak PSD constraint and a quite low PAR reduction performance.

A compromise between these two cases, with a limit on average PSD as well as a looser limit on maximum instantaneous PSD, would give a standard-compliant system that is not overly inefficient but still friendly to neighbouring users. The value of the peak PSD limit can be obtained from the acceptable amount of disturbance that can be put on other users.

2.2. Tone reservation

The goal with all PAR reduction methods is to generate a transmit signal that has a low amplitude swing. Different approaches to PAR reduction exist, but only a few are viable alternatives to include in standardised DSL systems. The schemes possible to use are those that are transparent to the receiver side, meaning that the receiver does not have to know about the existence of PAR reduction, nor which method is being used.

One of these viable methods is the tone reservation method [1, 2], which adds a reduction signal $c[n]$ to the data signal $x[n]$, see Figure 2. The goal is to make the resulting signal $\bar{x}[n] = x[n] + c[n]$ have a lower amplitude span than before. The PAR is defined as

$$\text{PAR}\{\bar{\mathbf{x}}\} = \frac{\max_n |x[n] + c[n]|^2}{\sigma^2}, \tag{1}$$

where the peak power is compared to the average power *before* the PAR reduction is applied, $\sigma^2 = E[|x[n]|^2]$.

The reduction signal $c[n]$ is constructed of a set of reserved subcarriers, which are not used for data transmission. These may be tones that cannot transmit data reliably, or tones that are explicitly reserved for PAR reduction. Naturally, the reduction performance will increase as the number of reserved tones is increased. At the same time, excluding too many tones from the set of data-carrying tones will

---

[1] Anything between 50%–80% could be considered normal, where the numbers have been increasing over time as the digital system parts have become more and more efficient.
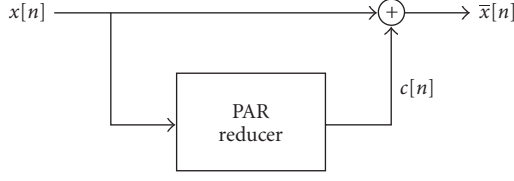
FIGURE 2: A PAR reduction signal, $c[n]$, is added to the data signal $x[n]$ to counteract the peaks in $x[n]$. The signal $c[n]$ is a function of the data signal $x[n]$, and is constructed from a small subset of tones.

reduce the data capacity of the system unnecessarily [19]. Thus, it is of interest to have, already at an early stage in a system design, some knowledge about the reduction capabilities for a certain number of tones in order to balance this tradeoff.

In matrix form, we define the signal model as

$$\bar{\mathbf{x}}_L = \mathbf{x}_L + \mathbf{c}_L = \mathbf{x}_L + \check{\mathbf{Q}}_L \check{\mathbf{C}}. \tag{2}$$

We let the length NL vector $\mathbf{x}_L$ denote the data signal of one symbol block and $\mathbf{c}_L$ the reduction signal of the same length. The FFT size is denoted by N and the number L represents the oversampling factor, which we introduce in order to have a better control of the continuous-time PAR. The construction of $\mathbf{c}_L$ from the reserved tones is written as $\mathbf{c}_L = \check{\mathbf{Q}}_L \check{\mathbf{C}}$, where $\check{\mathbf{Q}}_L$ is an $NL \times 2U$ matrix of sine and cosine basis vectors with frequencies specified by the U reserved tones $t_1, \ldots, t_U$ [12].

2.3. Optimization criteria

Having defined the reduction model in (2) above, we now formulate what the reduction algorithm should aim at. The most common approach in PAR reduction is basically to reduce the signal as much as possible. In the real-valued baseband environment of a DSL system, this can be formulated as a linear programme [2]:

$$\begin{aligned} &\underset{\check{\mathbf{C}}}{\text{minimize}} \ \gamma \\ &\text{subject to} \ \left| \mathbf{x}_L + \check{\mathbf{Q}}_L \check{\mathbf{C}} \right| \leq \gamma \sigma. \end{aligned} \tag{3}$$

The inequality compares each vector element to the right-hand side scalar $\gamma\sigma$, which is the level of the highest signal peaks.

When assigning a certain target PAR level to the system, the algorithm can be told not to put any efforts in reducing the peak level further down than this level. Since what is important in practice is to avoid signal clipping through overloading the line driver or clipping in the D/A converter, there is no reason to reduce the PAR of already acceptable symbols. Additionally, reducing the peak level further may cause the power on the reduction tones to increase to undesired levels. We can then define the optimisation criterion as minimising the added reduction power given a certain target crest factor

γ_{target}, or target PAR level γ_{target}^2 [15]:

$$\begin{aligned} &\underset{\check{\mathbf{C}}}{\text{minimize}} \ \check{\mathbf{C}}^T \check{\mathbf{C}} \\ &\text{subject to} \ \left| \mathbf{x}_L + \check{\mathbf{Q}}_L \check{\mathbf{C}} \right| \leq \gamma_{\text{target}} \sigma. \end{aligned} \tag{4}$$

This quadratic programme does not always have feasible solutions since it cannot be guaranteed that the target PAR level γ_{target}^2 is achievable. Thus, we choose to minimise the PAR level if the target PAR is not reached.

To take the maximum allowed PSD level into consideration, a set of quadratic constraints can be added to (3) and (4):

$$\check{C}_{l,\sin}^2 + \check{C}_{l,\cos}^2 \leq A_{l,\max}^2, \tag{5}$$

where $A_{l,\max}$ denotes the maximum magnitude for reduction tone t_l, with $l \in \{1 \cdots U\}$, and $\check{C}_{l,\sin}$ and $\check{C}_{l,\cos}$ denote the sine and cosine weights on a certain reduction tone. This introduction of a set of quadratic constraints will result in (3) and (4) no longer being linear or quadratic programmes. However, they will still be quadratically constrained problems, which still are convex and thereby reasonably easy to solve.

What will be studied now is the combination of the problem formulations in (3) and (4).

Step 1. First, try to solve (4) with the additional constraints in (5) to obtain the target PAR with as little added power as possible.

Step 2. If this fails, solve (3) and (5) to minimize the PAR level.

A suitable algorithm for practical implementation is the active-set algorithm [9, 12, 20]. Solving for the minimum PAR, this algorithm will converge to a solution close to the optimal solution already after a few iterations, and the PSD constraints can easily be incorporated [12].

2.4. Expression for PAR for an unreduced signal

In order to analyse what is achievable with the PAR reduction algorithm, the distribution of the PAR for an unreduced signal has to be derived. We here focus on the symbol clip probability, defined as the probability that the maximum sample value in a full DMT frame is above a certain level. This reflects the probability that a symbol is distorted during transmission. Our choice of definition is commonly used in the literature, although it would also be possible to view the clip probability on a per-sample basis.

Assuming that the signal after the IFFT is Gaussian IID, $x[n] \in N(0, \sigma^2)$, the sample clip probability at level $\gamma\sigma$ is

$$\mathcal{P}_1(\gamma) = \text{Prob}\left(\left|x[n]\right| > \gamma\sigma\right) = 2Q(\gamma), \tag{6}$$

where the $Q(\cdot)$ function denotes the tail probability for a Gaussian random variable:

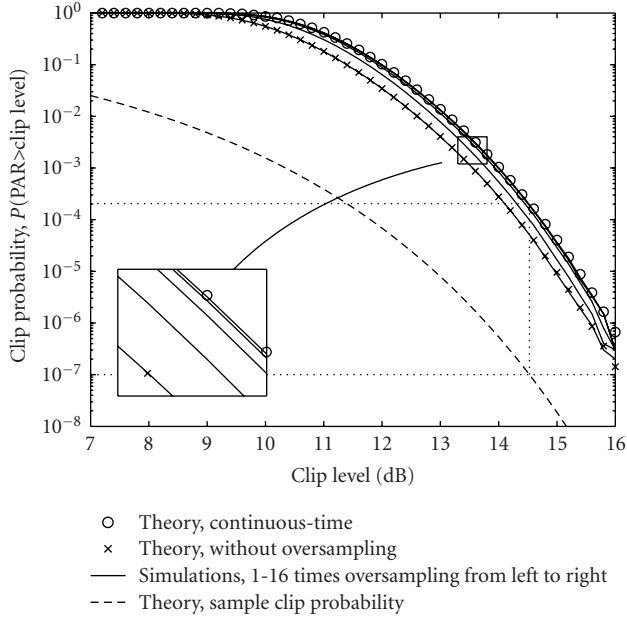$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-x^2/2} \, dx. \tag{7}$$

FIGURE 3: Symbol clip probability for an unreduced signal of length $N = 512$, evaluated at different levels of oversampling by sinc interpolation. The cross-marks show the theoretical values for a white Gaussian signal without oversampling and the circles show the calculated values for a lowpass continuous-time white signal using Rice's formula. The solid lines show, from left to right, the simulated clip probability for a signal without oversampling, and with 2, 4, 8, and 16 times oversampling, respectively. After four to eight times oversampling, the calculated results for the continuous-time signal describes the data signal closely. The dashed line shows the sample clip probability, which is the same for continuous-time signals as for signals without oversampling. The dotted lines show the translation between sample and symbol clip probabilities. Starting at the sample clip probability 10^{-7}, we see that this corresponds to a clip level of $\gamma_{\text{unred}} = 14.5$ dB, which for the continuous-time signal corresponds to the symbol clip probability $p_{\text{symclip}} = 2 \cdot 10^{-4}$.

The symbol clip probability for the symbol of N IID samples is straightforward to calculate [2]

$$\mathcal{P}_N(\gamma) = 1 - \text{Prob (all } |x[n]| < \gamma\sigma) = 1 - \left(1 - \mathcal{P}_1(\gamma)\right)^N$$
$$= 1 - \left(1 - 2Q(\gamma)\right)^N. \tag{8}$$

For a signal oversampled using sinc interpolation, the sample clip probability is identical to the critically sampled case. The derivation of the symbol clip probability is not as easy to obtain, due to that the signal no longer is IID. However, assuming that the signal is a continuous-time band-limited Gaussian noise gives us a possibility of deriving the clip probability using Rice's formula. The derivation is given in Appendix B, and the resulting expression is

$$\mathcal{P}_S(\gamma) = \text{Prob(clip at } \gamma\sigma) = 1 - \exp\left(-\frac{N}{\sqrt{3}}e^{-\gamma^2/2}\right). \tag{9}$$
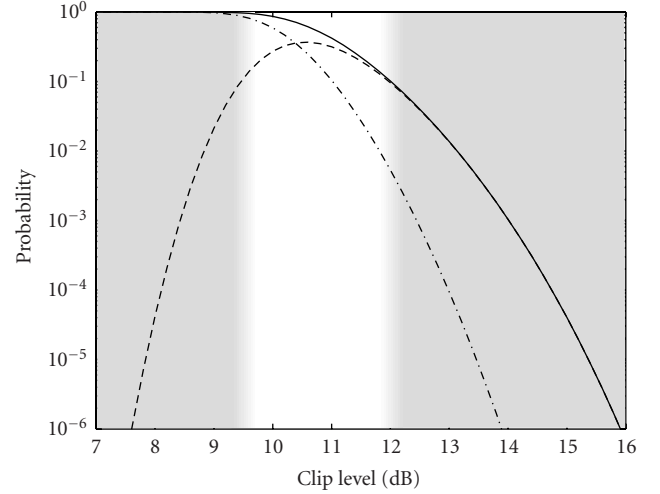

FIGURE 4: Probabilities of different number of clips, for a symbol length of $N = 512$. In the shaded area to the left, we will almost always have one or many clips. In the rightmost shaded area, there will most often be one or zero clips.

Figure 3 shows the sample and symbol clip probabilities for different levels of oversampling, from $L = 1$ to $L = 16$, for a system with a symbol length of $N = 512$. As seen from the figure, with increasing oversampling, the expression (9) gives an excellent match.

In Figure 4 we plot some qualitative results of when PAR reduction is needed. The solid line is the same as the theoretical values marked with circles in Figure 3, showing the probability of at least one clip in a symbol from (9). Using (B.12) in Appendix B, we plot two more lines showing the probabilities of at least two clips, and exactly one clip, respectively. We see that in the leftmost shaded region, the reduction algorithm will have to be active for almost every symbol. In the rightmost region, there will very seldom be more than one peak exceeding a given level. We will use these results to judge when the bounds developed in Section 3 are reasonably tight.

Moving on, while we focus on the symbol clip probability, the ADSL standard [16] is based on unreduced signals and sets the limit on clip probability on a per-sample basis. The signal is restricted to be clipped no more than the fraction 10^{-7} of the time. Translating this to a symbol clip probability is not straight-forward, since there can be one or many clips in a clipped symbol. However, considering an unreduced signal, we can use the expression for the sample clip probability in (6) to get the acceptable clip level:

$$\gamma_{\text{unred}} = Q^{-1}\left(\frac{10^{-7}}{2}\right) = 5.33 \quad (14.5 \text{ dB}). \tag{10}$$

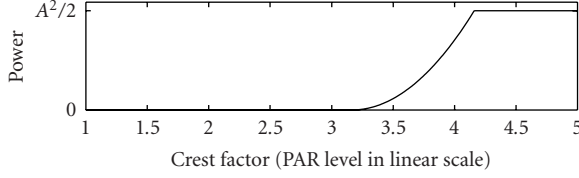Then, we can use (9) with this clip level to get a value of the

FIGURE 5: Illustration of (12), showing minimum instantaneous power on a reduction tone as a function of peak magnitude, with a target PAR level of 10 dB, $\gamma_{\text{target}} = 3.16$. The axes are shown in linear scale. Aiming at the target PAR level of 10 dB, the reduction algorithm will not be active below this value. Above $\gamma_{\text{target}} + UA/\sigma = 4.16$ (12.4 dB), the algorithm will output the maximum power $A^2/2$.

symbol clip probability:

$$p_{\text{symclip}} = 1 - \exp\left(-\frac{N}{\sqrt{3}} e^{-5.33^2/2}\right), \tag{11}$$

approximately $2 \cdot 10^{-4}$ for $N = 512$. This translation is also shown with the dotted lines in Figure 3. This γ_{unred} value will be used as a reference in Figures 6, 7, 8, 11, and 12, and the p_{symclip} is shown as the baseline reference in Figure 9.

3. PERFORMANCE PREDICTION WITH BOUNDS

Our aim is to see, or rather predict, what PAR level we can achieve with the tone reservation approach and PSD restrictions. We present a material allowing to do this analytically, without having to resort to extensive system simulations, which are often quite complicated when it comes to PAR reduction. Based on the optimisation criteria and distribution of the unreduced signal from the previous section, bounds for the achievable PAR level given a certain system using a certain number of tones will be derived. For easy practical use of these bounds, the outcome of this section will be summarised in Section 4.

3.1. Limitations imposed by a maximum average PSD

We will now analyse what can be done under the -10 dB per tone PSD limitation on the reduction signal compared to the data signal. Most of the symbols do not have very high signal peaks, and thus do not need a large reduction signal to pass under the clip level. We will assign a target PAR level to the algorithm, and not put any efforts in reducing the signal further down. We will also define a maximum magnitude per reduction tone. It can be expected that having a high target PAR level, only a few symbols will need reduction and the maximum tone magnitude can be set high, still having an average PSD below the limit. On the other hand, lowering the target PAR will demand reduction of more symbols, and the maximum reduction magnitude has to be kept lower, to not violate the average PSD constraint. This lower reduction magnitude will, in turn, make it difficult to achieve the target PAR. We will predict the best balance of the target PAR level γ_{target}^2 (how often reduction is used) and allowed reduction magnitude A (to stay below the PSD limit).

3.1.1. Allowed magnitude

We start by predicting the average power needed to reduce the signal down to the target PAR level γ_{target}^2. Then we can assign a maximum tone magnitude to avoid a too high average power. For an individual symbol, we use the following lower bound on reduction tone magnitude. Assume that the signal has a maximum peak with magnitude $x_{\max}$. To reduce the signal down to a PAR of γ_{target}^2 requires at least a total reduction tone magnitude of $x_{\max} - \gamma_{\text{target}}\sigma$.

Proof. Only looking at the highest peak of the signal, we see that letting all reduction tones be in phase at this sample, with a total magnitude of $x_{\max} - \gamma_{\text{target}}\sigma$ in the counter-phase direction, will reduce the signal down to $\gamma_{\text{target}}\sigma$, with PAR γ_{target}^2. All other scenarios, such as taking into consideration other samples in the original signal, and the possibility of generating new peaks, would need a larger reduction signal. □

For an individual symbol with peak magnitude $x_{\max}$, we need at least the total magnitude $x_{\max} - \gamma_{\text{target}}\sigma$, spread over the U tones. With the same PSD constraints on all tones, we would like to have the maximum tone magnitude as low as possible, which means having the same magnitude $(x_{\max} - \gamma_{\text{target}}\sigma)/U$ on all reduction tones. This could also be seen as a lower bound, or perhaps rather a best case, for a PSD-friendly average reduction power spread over the tones. Depending on the peak magnitude $x_{\max}$, the instantaneous reduction power on a certain tone will at least be

$$g(x_{\max})$$

$$= \begin{cases} 0 & x_{\max} \leq \gamma_{\text{target}}\sigma, \\ \frac{1}{2}\left(\frac{x_{\max} - \gamma_{\text{target}}\sigma}{U}\right)^2 & \gamma_{\text{target}}\sigma < x_{\max} \leq \gamma_{\text{target}}\sigma + UA, \\ A^2/2 & x_{\max} > \gamma_{\text{target}}\sigma + UA, \end{cases} \tag{12}$$

where A denotes the maximum allowed reduction magnitude per tone, see also Figure 5. Following Step 2 of the optimisation criterion given in Section 2.3, for peak levels above $\gamma_{\text{target}}\sigma + UA$, the algorithm can only output this large reduction signal, and the target PAR will thereby not be achieved.

To evaluate the minimum average PSD for a tone, we calculate the expected value of (12). For this, we need the density function $f(\gamma)$ for the normalised peak magnitude $x_{\max}/\sigma$, based on the clip probability from (9):

$$f(\gamma) = \frac{\partial F(\gamma)}{\partial \gamma} = \frac{\partial(1 - \mathcal{P}_S(\gamma))}{\partial \gamma}$$

$$= \frac{\partial}{\partial \gamma} \exp\left(-\frac{N}{\sqrt{3}} e^{-\gamma^2/2}\right) \tag{13}$$

$$= \frac{N\gamma}{\sqrt{3}} \exp\left(-\frac{N}{\sqrt{3}} e^{-\gamma^2/2}\right) e^{-\gamma^2/2}.$$
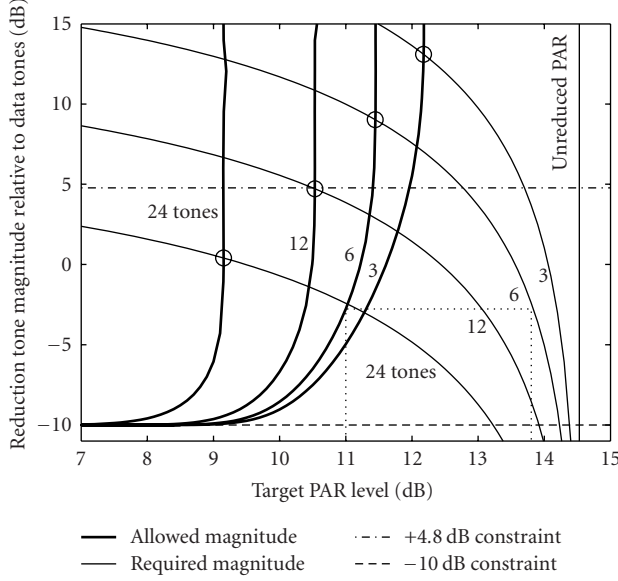
FIGURE 6: The thick solid lines starting from the bottom-left corner and bending upwards show the maximum reduction tone magnitude without exceeding the average PSD level, when aiming at the target PAR level shown on the horizontal axis. The thin solid lines starting from the bottom right show what reduction tone magnitude is needed to possibly achieve the target PAR level on the horizontal axis. The two horizontal lines at +4.8 dB and −10 dB show the limitations when having a maximum peak PSD, as described in Section 3.2.

Then we can calculate a lower bound on the average reduction power on each tone as the expected value of the minimum instantaneous power:

$$
\begin{aligned}
\text{Power}_{\text{red}} &\geq \int_0^\infty g(\gamma\sigma) f(\gamma) d\gamma \\
&= \int_{\gamma_{\text{target}}}^{\gamma_{\text{target}}+UA/\sigma} \frac{1}{2} \left(\frac{\gamma\sigma - \gamma_{\text{target}}\sigma}{U} \right)^2 f(\gamma) d\gamma \\
&\quad + \int_{\gamma_{\text{target}}+UA/\sigma}^\infty A^2/2 f(\gamma) d\gamma \\
&= \frac{1}{2U^2} \int_{\gamma_{\text{target}}}^{\gamma_{\text{target}}+UA/\sigma} (\gamma\sigma - \gamma_{\text{target}}\sigma)^2 f(\gamma) d\gamma \\
&\quad + \frac{A^2}{2} \mathcal{P}_S \left(\gamma_{\text{target}} + \frac{UA}{\sigma} \right).
\end{aligned}
\tag{14}
$$

In addition, to conform to the PSD limitation of a reduction tone PSD of 10 dB below the data-tone PSD, we can calculate the allowed average power $\text{Power}_{\text{red}}$ on a reduction tone as

$$
\text{Power}_{\text{red}} \leq 10^{-10/10} \text{Power}_{\text{data}} = 10^{-1} \frac{\sigma^2}{U_0 - U}, \tag{15}
$$

where U_0 is the number of tones originally available for data transmission before introducing tone reservation. The average power on the data tones, $\text{Power}_{\text{data}}$, is obtained by dividing the total signal power σ^2 with the number of tones used for data transmission. We assume that the power on the data tones fills the PSD mask completely. It is now possible
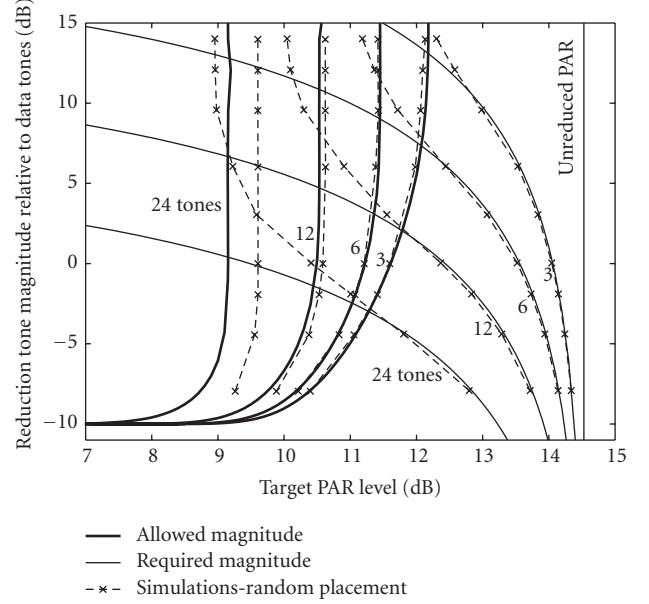

FIGURE 7: The bounds from Figure 6, shown together with simulations with a random placement of tones. Since the bounds only calculate with a single signal peak, the simulations results move somewhat upwards and to the right.

to use the two $\text{Power}_{\text{red}}$ expressions in (14) and (15) to solve for the maximum value of A. The thick solid lines, starting in the bottom left and bending upwards, in Figure 6 show the resulting values of the maximum magnitude A as a function of a certain target PAR level. The four lines correspond, from right to left, to systems with 3, 6, 12, and 24 tones, respectively. Choosing a target PAR level, we can read out the highest value we could set the maximum tone magnitude A to in order not to exceed the PSD limit on average. Thus, the allowed area is to the right, or below, the thick solid lines.

For very low target PAR levels in the system, almost all symbols will need to be reduced. In order to not exceed the constraint on the average PSD, we have to put a strict constraint on the maximum tone amplitude, A. For low target PAR levels, it can be seen that the thick solid lines for the peak constraint A converge to the average constraint at −10 dB.

On the other hand, if we aim at a not too low PAR, fewer symbols will need reduction (cf. Figure 4). Then the reduction signal can be allowed to be much stronger on occasional symbols without violating the average PSD limit. In Figure 6, the curves are bending strongly upwards at certain PAR levels. Choosing a target PAR level a bit above this value will allow a strong reduction signal, and thus a high probability of attaining the target.

3.1.2. Required magnitude

The curves developed in Section 3.1.1 and shown in Figure 6 tell us how much power we can put on the reduction tones without exceeding the PSD mask on average. However, the thick lines in Figure 6 give only a bound on an

allowable region for combinations of the target PAR level γ_{target}^2 and maximum reduction tone magnitude. They say nothing about what is achievable, that is, they do not guarantee that the target PAR level can, or will be, reached. Here we can reuse the lower bound on reduction tone magnitude from Section 3.1.1. Based on this bound and the distribution of the unreduced signal peak from (9), we can calculate what reduction tone magnitude at least is needed for a certain level of reduction at a certain clip probability:

$$A \geq \frac{(\gamma_{\text{unred}} - \gamma_{\text{target}})\sigma}{U}, \tag{16}$$

where γ_{unred} is the crest factor for the unreduced signal at the clip probability $p_{\text{symclip}} = 2 \cdot 10^{-4}$, as described in Section 2.4.

The thin solid lines starting from the bottom right in Figure 6 show this tone magnitude as a function of target PAR level, for 3, 6, 12, and 24 tones. The rightmost vertical line shows the PAR of the unreduced signal, γ_{unred}^2, which is about 14.5 dB at the clip probability $2 \cdot 10^{-4}$. To be able to reduce the signal level down to the target PAR on the horizontal axis, we need at least the amount of tone magnitude specified by the thin solid lines.

For each number of tones we then have two different lower bounds. The limit on average PSD gives a maximum allowed tone magnitude, and (16) gives a value of the minimum needed magnitude. Both bounds give a lower limit on PAR level and the allowable area is to the right of both curves. If it would be attainable, the best point would be given by the intersection of the bounds, for each number or tones marked with a circle in the figure.

The solid lines in Figure 6 show the allowed and required magnitude based on derivations only assuming one peak. Here Figure 4 comes to our aid. There we can see that the one-peak-only assumption almost always holds above 12 dB. To evaluate the derivations, simulation results are shown in Figure 7 with thin dashed lines next to the bounds. Looking at the thick curves for allowed tone magnitude, we see that for a low number of tones, the simulations closely follow the bound. For a higher number of tones, we cannot allow ourselves as high magnitude as the bound suggests. The reason is that with a high number of tones, we can work at low PAR levels, which means that we may have several peaks (cf. the leftmost grey area in Figure 4). Then the reduction algorithm has to spend power on reducing many peaks, and not one single peak, which is the situation the bound describes. Also, for the thin bound lines we see that the bounds are tighter at higher PAR levels. Reduction to low levels will generate signals with many peaks, again deviating from the basic case the bound is based on. Considering the simulations, the intersection points move upwards, to a higher magnitude per tone, rather than rightwards, to a higher PAR. The bound gives a good indication about the PAR, while it may show a too low required tone magnitude. The bounds will be further evaluated with simulations in Section 5.
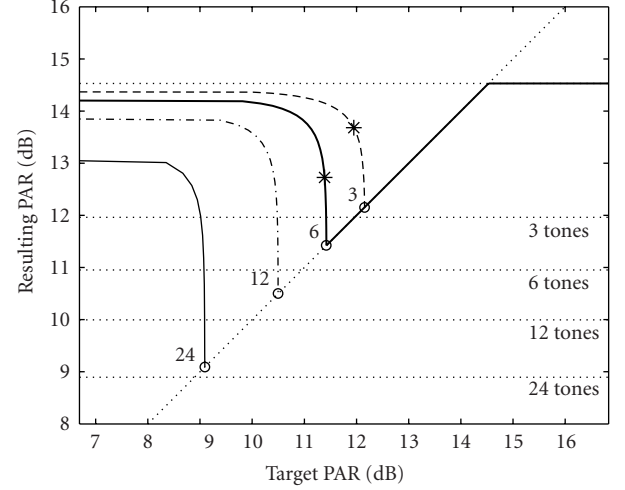

FIGURE 8: Relationship between target PAR and the effective resulting PAR level (given that the average PSD level is enforced), for reduction with 3, 6, 12, and 24 tones. The curve for 6 tones is a bold solid line. The dotted horizontal lines show the simulated value for reduction without PSD constraints, for the same number of tones. The asterisks show the combined -10 dB average PSD bound and $+4.8$ dB peak PSD bound based on FEXT calculations as described in Section 3.2. The 12 and 24 tone cases are not constrained by the $+4.8$ dB peak PSD bound.

3.1.3. Influence of target PSD

The balancing of target PAR level and allowed maximum amplitude is difficult. As an example, consider a system with six reduction tones. From Figure 6 we see that the optimal point is 11.4 dB. Let us assume that the designer chooses a target PAR level of 11.0 dB, which is below this point. Following the dotted line at 11.0 dB in Figure 6 up to the thick line corresponding to the allowed magnitude, we see that the maximum allowed reduction tone magnitude is -2.8 dB compared to the data tones. If the algorithm really aims at the 11.0 dB level, the PSD mask will be violated. Otherwise, the reduction tone magnitude has to be limited to -2.8 dB. Following the dotted line along this level to the thin 6-tone curve to the right shows us that we at maximum can reach a resulting PAR of 13.8 dB, or a crest factor of 4.90.

This relationship between target and resulting PAR while conforming to the PSD mask is shown in Figure 8. As described in the previous example, we can get the maximum A value based on a certain target PAR. Then we can use (16) to get the bound on the resulting PAR:

$$\gamma \geq \gamma_{\text{unred}} - \frac{UA}{\sigma}, \tag{17}$$

shown on the vertical axis for a clip probability of $2 \cdot 10^{-4}$. From the figure we see that there is a sharp optimum for each choice of number of tones. Aiming a little bit lower than this optimum, the resulting PAR level is severely increased (or in practice, the average PSD would be violated). The upper horizontal line shows the PAR when having no reduction at all, and the four lower lines show simulation results of what
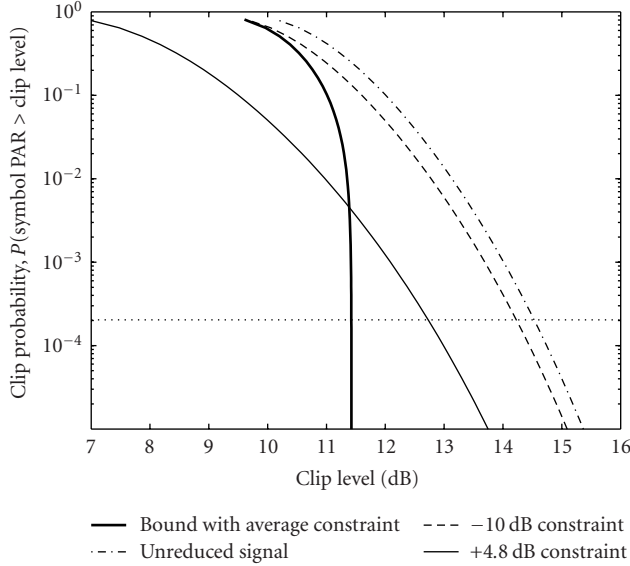
FIGURE 9: Clip probability curve for reduction with 6 tones. The thick line shows the bound for reduction under the average constraint for any parameter choice. The rightmost dash-dotted curve shows the PAR for an unreduced signal according to (9), and the dashed curve shows the bound when using no averaging. The thin solid line shows the bound when using the +4.8 dB constraint based on the FEXT calculations in Section 3.2.

could be achieved when not having any PSD constraints. The optimum points discussed above have higher PAR than these levels, which shows that the system performance is primarily limited by the constraint in PSD levels. Having no PSD constraints gives results dependent on the number and placement of the reduction tones.

Evaluating the bound in Figure 8 for different clip probabilities will give new values of the optimum points. These optimum points can then be plotted in a clip probability curve, commonly shown in PAR reduction papers; see the thick line in Figure 9.

Worth to notice in Figure 9 is that we get a minimum obtainable PAR that is almost independent on what clip probability you look at. Also, since the curve represents the best solution given different clip probabilities, a fixed design cannot get arbitrarily close to this bound at all values. We can only expect a system with a certain parameter choice to get close to the bound at the very specific target PAR level the system is designed for.

3.1.4. Combination of allowed and required magnitudes

We are now ready to define this optimum point in terms of achievable PAR level for a system with a certain number of tones. First we combine the two Power$_{red}$ expressions in (14) and (15) to get the relationship between A and U. To get a set of best design choices of U and γ_{target} (although maybe not reachable), we let the maximum allowed magnitude A be equal to the minimum needed required value from (16). After eliminating the dependence on σ^2, the resulting

expression is

$$10^{-1} \frac{1}{U_0 - U} \geq \frac{1}{2U^2} \int_{\gamma_{target}}^{\gamma_{unred}} (\gamma - \gamma_{target})^2 f(\gamma) d\gamma$$
$$+ \frac{1}{2U^2} (\gamma_{unred} - \gamma_{target})^2 \underbrace{\int_{\gamma_{unred}}^{\infty} f(\gamma) d\gamma}_{= p_{symclip} \text{ from (11)}}, \quad (18)$$

where we note that the second integral is the tail probability $p_{symclip}$ at level γ_{unred}. Let us interpret this equation. To the left, we have the allowed power on a tone, from the limit on the average PSD. To the right, we have a sum of two terms. The first is the power per tone needed to reduce from the peak level down to the target level. This corresponds to the first item in the optimisation formulation, "reduce down to the target PAR level using as little added power as possible." The second expression corresponds to the cases when the allowed instantaneous power is insufficient for reducing the peak down to the target PAR level. For these cases, we reduce with the maximum available instantaneous tone magnitude A, which from (16) is equal to $\sigma(\gamma_{unred} - \gamma_{target})/U$. This corresponds to the second item in the optimisation formulation, "minimise the PAR level if the target is unreachable." Compared to the first integral, the second term is very small, since it includes the small probability $p_{symclip}$ of a signal larger than γ_{unred}. Note that (18) can be solved easily with regards to U and γ_{target}. The relationship between U and the target PAR level is shown in Figures 11 and 12 in the upcoming section.

3.2. Limitations imposed by a maximum peak PSD

The derivations so far concerned only a limit on average PSD level. In Section 2.1 was discussed that a maximum instantaneous level may be needed as well, which will affect the reachable PAR level according to (17). This is the case considered in [12, 13], which we elaborate on here to give more detailed guidelines of where to put the constraint. We will consider the most restrictive case first, when the peak value is constrained at −10 dB and we do not use any averaging at all, and then extend to a scenario based on crosstalk calculations.

3.2.1. Constraint without averaging

For the case when no averaging is done between different symbols, all symbols need to comply with the −10 dB limitation and we cannot take advantage of the fact that most of the symbols need no or a very small amount of reduction. Thus, no power can be saved for the cases when it is needed to reduce a strong peak.

Two horizontal dashed lines are shown in Figure 6. The lower one is placed at −10 dB and corresponds to this limitation. In addition to being on the right-hand side of the bounds discussed before, we now also have to be below this horizontal line. As can be seen from the figure, this severely affects the reduction performance. For example, using only 6 PAR reduction tones, only about 0.3 dB reduction can be achieved. In this case, the PAR reduction will most likely not be worth the effort.

In Figure 8, the curves move right and down with increasing target PAR level. Not using any averaging corresponds to having so low target PAR that all symbols need reduction. What may be possible to achieve here is thus shown by the values to the left in the figure. Having this strict limitation on the reduction signal makes the PAR reduction only usable if there is a high number of tones available for PAR reduction, for instance when there are many tones unavailable to carry data due to low SNR.

3.2.2. Constraint based on FEXT calculations

The two previous extreme cases show very different reduction results, due to the difference in maximum allowed magnitude. It can be expected that there should be a maximum instantaneous tone power somewhere in between, determined to not cause harmful interference to other users. If we consider the *far-end crosstalk* (FEXT), we can come up with a reasonable point to put the PSD constraint at.

Consider a situation where our modem in question is responsible for one quarter of the FEXT to a certain neighbouring modem, a fairly pessimistic case. As our reduction tones are at -10 dB of our data tones, it means that the neighbouring modem in question on the reduction tones are expecting a FEXT level of three quarters of the normal level. ADSL is designed to use a SNR margin of 6 dB. If that user has a SNR margin of 6 dB, how much can we then increase our disturbance, without more than half of this SNR margin being lost? Adding 3 dB to the FEXT level means adding (during an instant only) as much FEXT as the user already has from all the other users, that is, matching the three quarters of the nominal FEXT level. This is three times the one quarter we have on the data-carrying tones, which means a peak level of +4.8 dB on our reduction tones as compared to our data-carrying tones.

This level is shown as the upper horizontal dash-dotted line in Figure 6. We also need to be below this line, but compared to the previous, much more restrictive, case, the system now has more reduction capabilities. The 24-tone system is not affected at all by this peak constraint, and the 12-tone system has just enough magnitude to avoid being affected. However, the systems with 3 and 6 tones are still limited by the constraint, which is also shown with the asterisks in Figure 8.

4. NUMERICAL RECIPE

To summarize the previous section, we will now give a description of how to, based on a certain system environment, generate an estimate of PAR reduction performance under PSD constraints. A description and typical values of the system parameters are given in Table 1, and the design parameters we want to obtain values on are given in Table 2.

To get a lower bound on the achievable PAR under an average PSD constraint dB_{avg}, typically -10 dB, use for example Matlab's quad function to evaluate the following function,
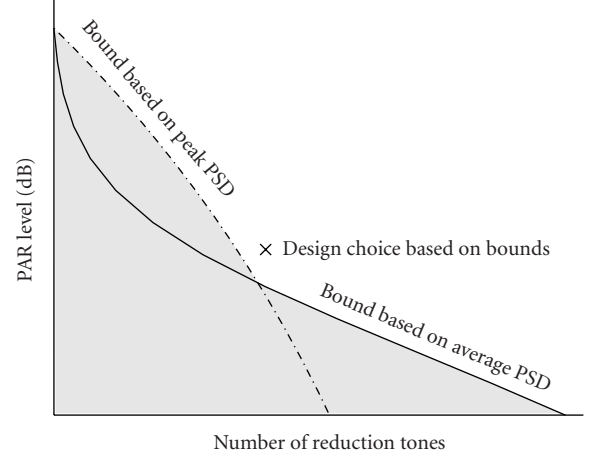

FIGURE 10: Schematic figure of two bounds on achievable PAR reduction, shown as a function of number of reduction tones. The two bounds based on maximum average and maximum peak PSDs define an unreachable region. For system design, a target point with a certain safety margin is desired.

directly following from (18):

$$\frac{U^2}{U_0 - U} 2 \cdot 10^{dB_{avg}/10} \geq \int_{\gamma_{target}}^{\gamma_{unred}} (\gamma - \gamma_{target})^2 f(\gamma) d\gamma \qquad (19)$$
$$+ \mathcal{P}_S(\gamma_{unred})(\gamma_{unred} - \gamma_{target})^2,$$

where

$$f(\gamma) = \frac{N\gamma}{\sqrt{3}} \exp\left(-\frac{N}{\sqrt{3}} e^{-\gamma^2/2}\right) e^{-\gamma^2/2},$$
$$\mathcal{P}_S(\gamma) = 1 - \exp\left(-\frac{N}{\sqrt{3}} e^{-\gamma^2/2}\right). \qquad (20)$$

This equation can be plotted for U (number of reduction tones) as a function of γ_{target} (target PAR level), showing the minimum number of reduction tones needed to achieve a certain PAR level. An example of this bound is plotted as the solid line in Figure 10.

A lower bound on achievable PAR under a peak PSD constraint dB_{peak}, typically +4.8 dB, is from (17) given by

$$\gamma \geq \gamma_{unred} - 10^{dB_{peak}/20} \frac{U\sqrt{2}}{\sqrt{U_0 - U}}. \qquad (21)$$

Such a bound has been plotted with the dash-dotted line in Figure 10. Together with the previous bound, it defines an unreachable region marked with grey. When designing a system using randomly chosen tones, choosing a point 1 dB outside this area should be sufficient as a tentative design choice. The 1 dB margin will be motivated in the following section.

5. EFFECTS ON SYSTEM PERFORMANCE

Using the method in Section 4, we now combine the bounds from Section 3 to evaluate what effect these will have on system performance, first for an ADSL2 system, then also for an ADSL2+ system.
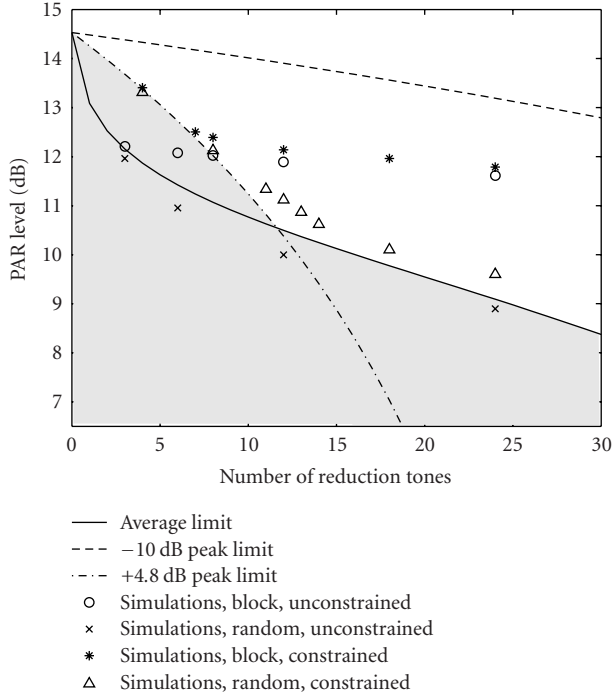
FIGURE 11: Bounds on achievable PAR as a function of a number of reduction tones, for an ADSL2 system with $N = 512$, and a symbol clip probability of $2 \cdot 10^{-4}$. The solid line shows the bound based on a maximum average PSD. The dash-dotted line shows the bound based on a peak PSD set by FEXT calculations, and the upper dashed line is the very restrictive case when no averaging is used at all. The marker symbols show simulated reduction performance, for random or block-placed tones, with or without PSD constraints.

5.1. ADSL2 system with combined peak-and-average constraints

Looking at the bounds for a different number of tones at a certain clip probability, we obtain a plot of the achievable PAR as a function of the number of reduction tones, see Figure 11. These curves show what performance could at best be achieved when having a certain number of reduction tones evaluated at a certain probability.

The solid line shows the bound for reduction under the -10 dB average PSD limit from (18) in Section 3.1. As seen, the bound decreases very fast in the beginning, then decreases slower at the higher number of reduction tones. With only a few tones, we can effectively reduce the highest peaks in the signal. These high peaks occur very seldom, so significant reduction can be achieved without increasing the average PSD much. Using more tones, we can start working with lower peaks in the signal as well, but cannot save the reduction power to a few, very high, peaks.

The two other lines show the peak PSD constraints from Section 3.2, using (17). The upper dashed line shows the restrictive -10 dB PSD constraint based on no averaging. Since this constraint admits only a very small reduction signal, we would have to use a very large number of tones to get any significant PAR reduction. More interesting is the $+4.8$ dB

constraint imposed by the FEXT calculations, shown by the dash-dotted line. The larger allowed signal for this case will give a significantly better performance. We can see that this peak-PSD bound crosses the bound based on average PSD after 11 reduction tones. The good performance for only a few tones indicated by that bound demanded a very large reduction signal. Having a constraint on this as well, we see that the combined bound, with both the solid and dash-dotted lines, will show a better yield per tone in the beginning up to 11 tones. After this point, the average PSD will be the limiting factor, and the performance will increase slower.

The lines only represent bounds on achievable performance based on PSD limitations. The performance will also depend on what tones are chosen as reduction tones. Without PSD-constraints, tones spread out over the frequency band in an unstructured manner will give better results than regularly or block-placed tones [2, 15, 21]. However, spreading out tones over the whole frequency band is not too attractive in wireline systems as the SNR for the lower part of the spectrum generally is much better than for the highest tones. We will consider two extreme cases: tones randomly placed over the frequency band and tones allocated as a contiguous block of the highest tones. In practice, a combination of these extreme cases may be a good choice.

Simulation results for reduction with and without PSD constraints are also shown in Figure 11. The simulations are done with 8 times oversampling, and a 32-sided linear approximation [12] of the quadratic power constraint in (5). For the simulations without PSD constraints, we see that the performance for the block-placed tones increases very slowly with the number of reduction tones, compared to the random placement. Most of the reduction performance is achieved after only a few reduction tones. Nevertheless, when the PSD is constrained, the peak PSD bound sets the limit on performance up to about 8 tones. This suggests that the tone placement is of minor importance for this low number of tones.

Since the bound for the peak constraints decreases faster than the bound for the average constraint, we are interested in the point where the two bounds meet. From Figure 11 this can be seen to be between 11 and 12 tones. Simulations performed with randomly scattered tones are shown with the triangles. The bounds give us a good hint about what the resulting performance will be. However, with a high number of tones, we do not end up as close to the bounds as we do with a low number of tones. The difference is around one half dB. This could be seen already in Figure 6. There, the simulations did not allow as much reduction magnitude as the bound suggested, due to the bound only calculating with one peak.

The simulated PAR for the number of tones where the bounds cross was around 11.3 dB. Translating to linear scale, we have a crest factor of about 3.7. It is interesting to know what this means in volts on the line if this is to be used, for example, to design the supply voltage to the line driver. If we have a transmit power of $P = 20$ dBm (100 mW), and a load of $R = 100 \, \Omega$, the RMS value of the transmit signal is $U = \sqrt{PR} = \sqrt{0.1 \, \text{W} \cdot 100 \, \Omega} = 3.16$ V. This means that instead of
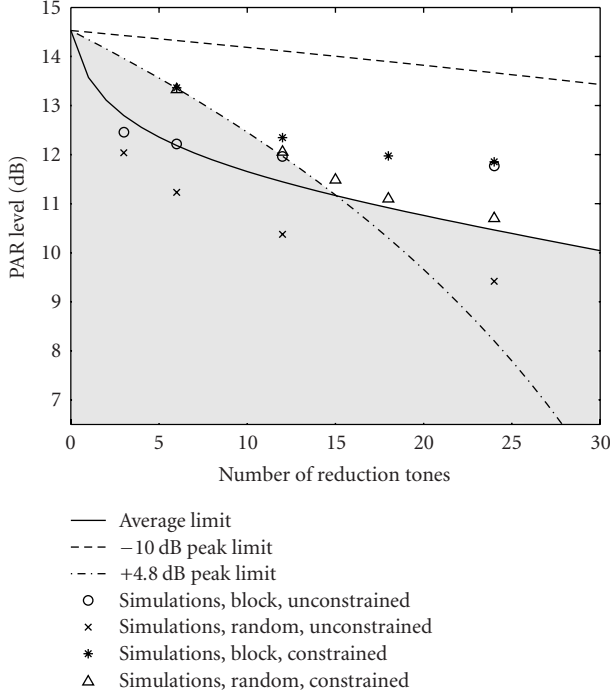
FIGURE 12: Bounds on achievable PAR as a function of number of reduction tones. This figure is like Figure 11, but shows an ADSL2+ system with $N = 1024$, and a symbol clip probability of $4 \cdot 10^{-4}$.

$3.16\,\text{V} \cdot 5.33 \approx 17\,\text{V}$ we now have a peak of $3.16\,\text{V} \cdot 3.7 \approx 12\,\text{V}$. (In practice, this calculation is more complicated, involving higher loads, step-up transformers, etc.)

The reduction in signal span and supply voltage gives a good indication on power consumption in the line driver. However, the exact relationship between the supply voltage and power consumption is not straightforward, since it depends on how capacitive the load is [22]. The rules of thumb for power consumption given by designers of line drivers vary between linear and quadratic functions of the supply voltage. Reduction from $17\,\text{V}$ to $12\,\text{V}$ would then mean a power reduction between 30% and 50%.

5.2. Extension to ADSL2+

The derivations so far have described a system with an FFT length of 512 in the downstream direction, such as ADSL2. ADSL2+, the extension of ADSL2, uses more downstream bandwidth, which means that the downstream FFT length here instead is 1024, and all tones in the newly added top half of the spectrum are used for downstream transmission.

We have seen that the limits on achievable PAR reduction performance are strongly dependent on the number of PAR reduction tones. Increasing the system FFT size to 1024, where we use a flat transmit PSD for all data tones, moves the intersection point between the two bounds up to about 15 tones, as shown in Figure 12. The simulation results are here closer to the bound than in the ADSL2 case. We see that also in this case, we can expect a faster PAR decrease down

TABLE 1: System parameters for calculation of reduction bounds.

Symbol	Description	Typical value
N	FFT size	ADSL2: 512 ADSL2+: 1024
γ_{unred}	Crest factor (PAR in linear scale) for the unreduced signal	5.33 (14.5 dB)
U_0	Number of data tones available before PAR	ADSL2: 223 ADSL2+: 479
dB_{avg}	Average PSD constraint in dB	-10
dB_{peak}	Peak PSD constraint in dB	$+4.8$

TABLE 2: Design parameters for calculation of reduction bounds.

Symbol	Description
U	Number of reduction tones
γ_{target}	Target crest factor

to about 11 dB. Below that level, a higher number of tones is needed for the same gain.

If we compare Figures 11 and 12, we see that the bounds have moved to the right for the ADSL2+ case with the double number of subcarriers. The success in PAR reduction depends on how large reduction signal we can create, in comparison to the size of the data signal. Thus, what is important for good reduction under a PSD constraint is often the *relative* number of reduction tones. On the contrary, the reduction performance of the nonconstrained reduction is more dependent on the *absolute* number of reduction tones, since what effects the reduction performance here is how well we can create an impulse in counter-phase to the signal peak.

6. CONCLUSIONS

Applying tone reservation to DSL systems can reduce the system PAR, with algorithms possible to implement in current standards. For implementation, low-cost active-set-based algorithms may be a good choice [12]. However, designing the algorithms, care must be taken not to exceed the PSD levels set up by the standards. Introducing PSD constraints will significantly alter the achievable performance of the reduction systems.

Using the requirements specified in the ITU standards and extending this with engineering assumptions, we have derived bounds on achievable performance. These bounds are applicable to most DMT systems, such as ADSL2 and ADSL2+. Simulations searching for the optimal solution confirm that the bounds give a good indication of realistic system performance.

We have demonstrated how these bounds can be used to predict system performance for varying parameter choices, and we have exemplified how they can be used to tailor PAR reduction to different systems. Thus, the bounds can be used to quickly determine if tone reservation PAR reduction is a

worthwhile technology to be included in a system. Using the bounds, this can be done in an afternoon. After a positive indication, the system design could then proceed with the large task to create a simulation chain in order to fine-tune the settings for the PAR reduction algorithm.

APPENDICES

A. PSD CONSTRAINTS FORMULATED IN THE STANDARDS

The ADSL1 standard [16] defines the PSD constraints as:

> For the subcarriers with ($b_i = 0$ and $g_i = 0$), the ATU-C transmitter should and is recommended to transmit no power on those subcarriers. The ATU-R receiver cannot assume any particular PSD levels on those subcarriers. The transmit PSD levels of those subcarriers with $g_i = 0$ will be at least 10 dB below the sync symbol reference transmit PSD level if the subcarrier is below the lowest used subcarrier (lowest i with $b_i > 0$) and will be below the sync symbol reference transmit PSD level if the subcarrier is above the lowest used subcarrier.

In ITU standards, the word *will* defines a mandatory requirement, *should* a recommendation, and *may* an option. The ADSL1 standard thus recommends to not use these subcarriers for PAR reduction, but allows the same PSD as for the data tones, except for a -10 dB limitation on the tones below the lowest data subcarrier.

In the newer ADSL2 standard [17], the formulation is:

> For the subcarriers not in the MEDLEYset, the ATU will transmit no power on the subcarrier (i.e., $Z_i = 0$, see Section 8.8.2) if the subcarrier is below the first used subcarrier index or if the subcarrier is in the SUPPORTED-set and in the BLACKOUTset. Otherwise, the ATU may transmit at a discretionary transmit PSD level on the subcarrier (which may change from symbol to symbol), not to exceed the maximum transmit PSD level for these subcarriers. The maximum transmit PSD level for each of these subcarriers will be defined as 10 dB below the reference transmit PSD level, fine tuned by RMSGI dB (see Section 8.5) and limited to the transmit spectral mask.

Below the lowest data subcarrier, no PAR reduction tones are allowed. However, there is no recommendation to not use the other subcarriers as reduction tones, as long as they are at least 10 dB below the data tone PSD.

Summarizing the standards, we see that although the recommendation is to not put any power at all on the reduction tones, the ADSL1 standard allows much higher power, compared to the ADSL2 standard. Thus, the more well-defined ADSL2 formulation sets a stricter limit on reduction performance.

B. DERIVATION OF THE PAR FOR AN OVERSAMPLED SYSTEM

The derivation is based on the signal viewed as a Gaussian process [14]. This can also be intuitively motivated by the central limit theorem when adding many subcarriers with independent data. We model the signal as Gaussian noise with constant spectral density in the frequency interval up to f_1:

$$R(f) = \begin{cases} \dfrac{\sigma^2}{2f_1} & \text{if } -f_1 \le f \le f_1, \\ 0 & \text{otherwise.} \end{cases} \tag{B.1}$$

Rice's formula for a Gaussian stationary process [23–25] states that the intensity of upcrossings, that is, the expected number of upcrossings of the level $\gamma\sigma$ in an interval of length 1, is

$$\mu^+(\gamma\sigma) = E\left(N^+_{[0,1](x,\gamma\sigma)}\right) = \frac{1}{2\pi}\sqrt{\frac{\lambda_2}{\lambda_0}}e^{-\gamma\sigma^2/(2\lambda_0)}, \tag{B.2}$$

where λ_0 and λ_2 are functions of the covariance function $r(\tau)$:

$$\lambda_0 = r(0), \qquad \lambda_2 = -r''(0). \tag{B.3}$$

If we consider the case when the level $\gamma\sigma$ is high enough for the crossing times to be spread out and independent, we can model the time between each crossing as exponentially distributed. Then the number of upcrossings during the interval of length 1 follows a Poisson process with intensity $\mu^+(\gamma\sigma)$. The corresponding process describing the number of downcrossings of the level $-\gamma\sigma$ will have identical intensity, due to the symmetry and zero-mean of the signal. We are interested in the intensity of crossing either $\gamma\sigma$ or $-\gamma\sigma$, and this intensity could then be described by the sum of the two single-sided intensities.

The probability that we will have no crossings of the level $\gamma\sigma$ in a time interval of length T is

$$\text{Prob}\left(\max(|x|) < \gamma\sigma\right) = e^{-T2\mu^+(\gamma\sigma)}, \tag{B.4}$$

and the clip probability is thereby given by

$$\mathcal{P}_S(\gamma) = \text{Prob}(\text{clip at level } \gamma\sigma) = 1 - e^{-T2\mu^+(\gamma\sigma)}. \tag{B.5}$$

To solve this for our signal, we start with calculating the covariance function:

$$r(\tau) = \int_{-\infty}^{\infty} e^{j2\pi f\tau}R(f)df = \sigma^2\frac{\sin(2\pi\tau f_1)}{2\pi\tau f_1}, \tag{B.6}$$

which shows the variance of the signal, $r(0) = \sigma^2$. Differentiating $r(\tau)$ gives

$$r'(\tau) = \sigma^2\left(\frac{\cos(2\pi\tau f_1)}{\tau} - \frac{1}{2}\frac{\sin(2\pi\tau f_1)}{\pi\tau^2 f_1}\right),$$

$$r''(\tau) = \sigma^2\left(-2\frac{\sin(2\pi\tau f_1)\pi f_1}{\tau}\right.$$
$$\left. -2\frac{\cos(2\pi\tau f_1)}{\tau^2} + \frac{\sin(2\pi\tau f_1)}{\pi\tau^3 f_1}\right). \tag{B.7}$$

Then the intensities λ_0 and λ_2 are given by

$$
\begin{aligned}
\lambda_0 &= r(0) = \sigma^2, \\
\lambda_2 &= -r''(0) = -\lim_{\tau \to 0} r''(\tau) \\
&= 4\pi^2 f_1^2 \sigma^2 \\
&\quad - \lim_{\tau \to 0} \sigma^2 \Bigg[-\frac{2}{\tau^2} \left(1 - \frac{1}{2}(4\pi^2 \tau^2 f_1^2) + O(\tau^4) \right) \\
&\qquad + \frac{1}{\pi\tau^3 f_1} \left((2\pi\tau f_1) - \frac{1}{6}(8\pi^3 \tau^3 f_1^3) + O(\tau^5) \right) \Bigg] \\
&= 4\pi^2 f_1^2 \sigma^2 - \lim_{\tau \to 0} \sigma^2 \left[+4\pi^2 f_1^2 - \frac{8}{6}\pi^2 f_1^2 + O(\tau^2) \right] \\
&= \frac{4}{3}\pi^2 f_1^2 \sigma^2.
\end{aligned} \tag{B.8}
$$

Inserting this into (B.2) and (B.5), we get

$$
\begin{aligned}
\mathcal{P}_S(\gamma) &= 1 - \exp\left(-T\frac{2}{2\pi}\sqrt{\frac{(2\pi f_1)^2}{3} e^{-\gamma^2 \sigma^2/(2\lambda_0)}} \right) \\
&= 1 - \exp\left(-2Tf_1 \frac{1}{\sqrt{3}} e^{-\gamma^2/2} \right).
\end{aligned} \tag{B.9}
$$

The factor Tf_1 is dependent on the symbol length. If the sample rate is F_s, and the whole band up to the Nyquist frequency is used, then $f_1 = F_s/2$. At the same time, T corresponds to N samples, each $1/F_s$ in time. Thereby,

$$
Tf_1 = \frac{N}{F_s}\frac{F_s}{2} = \frac{N}{2}. \tag{B.10}
$$

The resulting clip probability is then

$$
\mathcal{P}_S(\gamma) = \text{Prob}(\text{clip at level } \gamma\sigma) = 1 - \exp\left(-\frac{N}{\sqrt{3}} e^{-\gamma^2/2} \right). \tag{B.11}
$$

Notable with this approximation is that the clip probability for $\gamma = 0$ is not exactly one but instead $1 - \exp(-N/\sqrt{3})$. As mentioned, when modelling the crossings with a Poisson process, the model is not applicable when we have a too low clip level, that is, close to 0. However, this low region is not of interest. The related problem of deriving clip probability for a complex-valued continuous-time OFDM signal was addressed in [26].

Using the Poisson process model, we can also calculate the probability for exactly a certain number of clips. In particular, the probability for one single clip during a symbol interval is

$$
T2\mu^+(\gamma\sigma)e^{-T2\mu^+(\gamma\sigma)} = \frac{N}{\sqrt{3}} e^{-\gamma^2/2} \exp\left(-\frac{N}{\sqrt{3}} e^{-\gamma^2/2} \right). \tag{B.12}
$$

ACKNOWLEDGMENTS

REFERENCES

[1] A. Gatherer and M. Polley, "Controlling clipping probability in DMT transmission," in *Proceedings of 31st Asilomar Conference on Signals, Systems, and Computers*, vol. 1, pp. 578–584, Pacific Grove, Calif, USA, November 1997.

[2] J. Tellado-Mourelo, *Peak to average power reduction for multi-carrier modulation*, Ph.D. thesis, Stanford University, Stanford, Calif, USA, 1999.

[3] M. Friese, "Multitone signals with low crest factor," *IEEE Transactions on Communications*, vol. 45, no. 10, pp. 1338–1344, 1997.

[4] D. J. G. Mestdagh and P. M. P. Spruyt, "A method to reduce the probability of clipping in DMT-based transceivers," *IEEE Transactions on Communications*, vol. 44, no. 10, pp. 1234–1238, 1996.

[5] B. M. Popovic, "Synthesis of power efficient multitone signals with flat amplitude spectrum," *IEEE Transactions on Communications*, vol. 39, no. 7, pp. 1031–1033, 1991.

[6] R. W. Bauml, R. F. H. Fischer, and J. B. Huber, "Reducing the peak-to-average power ratio of multicarrier modulation by selected mapping," *Electronics Letters*, vol. 32, no. 22, pp. 2056–2057, 1996.

[7] D. L. Jones, "Peak power reduction in OFDM and DMT via active channel modification," in *Proceedings of 33rd Asilomar Conference on Signals, Systems, and Computers*, vol. 2, pp. 1076–1079, Pacific Grove, Calif, USA, October 1999.

[8] P. O. Börjesson, H. G. Feichtinger, N. Grip, et al., "A low-complexity PAR-reduction method for DMT-VDSL," in *Proceedings of 5th International Symposium on Digital Signal Processing for Communication Systems (DSPCS '99)*, pp. 164–169, Perth, Australia, February 1999.

[9] B. S. Krongold and D. L. Jones, "An active-set approach for OFDM PAR reduction via tone reservation," *IEEE Transactions on Signal Processing*, vol. 52, no. 2, pp. 495–509, 2004.

[10] W. Henkel and V. Zrno, "PAR reduction revisited: an extension to Tellado's method," in *Proceedings of 6th International OFDM-Workshop (InOWo '01)*, pp. 31-1–31-6, Hamburg, Germany, September 2001.

[11] J. Tellado and J. M. Cioffi, *Further results on peak-to-average ratio reduction*, ANSI Document, T1E1.4 no. 98-252, August, 1998.

[12] N. Andgart, B. S. Krongold, P. Ödling, A. Johansson, and P. O. Börjesson, "PSD-constrained PAR reduction for DMT/OFDM," *EURASIP Journal on Applied Signal Processing*, vol. 2004, no. 10, pp. 1498–1507, 2004.

[13] N. Petersson, A. Johansson, P. Ödling, and P. O. Börjesson, "A performance bound on PSD-constrained PAR reduction," in *Proceedings of IEEE International Conference on Communications (ICC '03)*, vol. 5, pp. 3498–3502, Anchorage, Alaska, USA, May 2003.

[14] S. Wei, D. L. Goeckel, and P. E. Kelly, "A modern extreme value theory approach to calculating the distribution of the peak-to-average power ratio in OFDM systems," in *Proceedings of IEEE International Conference on Communications (ICC '02)*, vol. 3, pp. 1686–1690, New York, NY, USA, April–May 2002.

[15] N. Petersson, "Peak and power reduction in multicarrier systems," Licentiate thesis, Lund University, Lund, Sweden, 2002.

[16] ITU-T, *Asymmetric digital subscriber line (ADSL) transceivers*, Recommendation G.992.1, June 1999.

[17] ITU-T, *Asymmetric digital subscriber line (ADSL) transceivers—2 (ADSL2)*, Recommendation G.992.3, July 2002.

[18] ITU-T, *Asymmetric digital subscriber line (ADSL) transceivers—Extended bandwidth ADSL2 (ADSL2+)*, Recommendation G.992.5, May 2003.

[19] P. Ödling, N. Petersson, A. Johansson, and P. O. Börjesson, "How much PAR to bring to the party?" in *Proceedings of IEEE Nordic Signal Processing Symposium (NORSIG '02)*, Tromsø-Trondheim, Norway, October 2002.

[20] D. G. Luenberger, *Linear and Nonlinear Programming*, Addison-Wesley, Reading, Mass, USA, 2nd edition, 1984.

[21] N. Petersson, A. Johansson, P. Ödling, and P. O. Börjesson, "Analysis of tone selection for PAR reduction," in *Proceedings of 3rd International Conference on Information, Communications and Signal Processing (ICICS '01)*, Singapore, October 2001.

[22] N. Larsson and K. Werner, "Signal peak reduction for power amplifiers with active termination," Master's thesis, Lund Institute of Technology, Lund, Sweden, 2002.

[23] G. Lindgren, *Lectures on Stationary Stochastic Processes*, Mathematical Statistics, Lund University, Lund, Sweden, 2002.

[24] S. O. Rice, "Mathematical analysis of random noise," *Bell Systems Technical Journal*, vol. 23, pp. 282–332, 1944.

[25] S. O. Rice, "Mathematical analysis of random noise," *Bell Systems Technical Journal*, vol. 24, pp. 46–156, 1945.

[26] H. Ochiai and H. Imai, "On the distribution of the peak-to-average power ratio in OFDM signals," *IEEE Transactions on Communications*, vol. 49, no. 2, pp. 282–289, 2001.

Niklas Andgart was born in Hässleholm, Sweden, in 1975. He received the M.S.E.E. degree in 2000, the Licentiate in Engineering degree in 2002, and the Ph.D. degree in signal processing in 2005, all from Lund University. During the fall of 1999, he was with the Vehicle and Dynamics Laboratory at the University of California at Berkeley, and in early 2004, he was visiting the Department of Electrical and Electronic Engineering at the University of Melbourne. Currently, he is with the Department of Information Technology at Lund University. His research is within signal processing for communication systems and he works with DSL research in cooperation with Ericsson AB in Stockholm.

Per Ödling was born 1966 in Örnsköldsvik, Sweden. He received an M.S.E.E. degree in 1989, a Licentiate in Engineering degree in 1993, and a Ph.D. degree in signal processing in 1995, all from Luleå University of Technology, Sweden. In 2000, he was awarded the Docent degree from Lund Institute of Technology, and in 2003, he was appointed Full Professor there. From 1995, he was an Assistant Professor at Luleå University of Technology, serving as Vice-Head of the Division of Signal Processing. In parallel, he consulted for Telia AB and ST Microelectronics, developing an OFDM-based proposal for the standardization of UMTS/IMT-2000 and VDSL for standardization in ITU, ETSI, and ANSI. Accepting a position as Key Researcher at the Telecommunications Research Center, Vienna, in 1999, he left the Arctic North for historic Vienna. There he spent three years advising graduate students and industry. He also consulted for the Austrian Telecommunications Regulatory Authority on the unbundling of the local loop. He is, since 2003, a Professor at Lund Institute of Technology, stationed at Ericsson AB, Stockholm. He also serves as an Associate Editor for the IEEE Transactions on Vehicular Technology. He has published more than fourty journal and conference papers, thirty five standardization contributions, and a dozen patents.

Albin Johansson was born 1968 in Stockholm, Sweden. He received an M.S.E.E. degree in 1993 from Royal Institute of Technology in Stockholm and is now pursuing a Ph.D. degree at Lund Institute of Technology. Since 2004, he has been holding a position at Ericsson AB as an Expert Broadband Access within Wireline, being responsible for system architecture in Ericsson's wireline broadband access products. He has been actively involved in the development of the standardization of ADSL within ETSI, ANSI, ITU-T, and ADSL Forum. He has been Editor for ITU-T G.997.1 and Chair in one of ADSL Forums subcommittees. In addition, from 1992 to 1995, he was teaching undergraduate students at Royal Institute of Technology. Since 2001, he has been a Member of the Signal Processing Group at Lund Institute of Technology. He has published 6 conference papers, numerous standardization contributions, and holds 7 patents.

Per Ola Börjesson was born in Karlshamn, Sweden in 1945. He received his M.S. degree in electrical engineering in 1970 and his Ph.D. degree in telecommunication theory in 1980, both from Lund Institute of Technology (LTH), Lund, Sweden. In 1983, he received the degree of Docent in telecommunication theory. Between 1988 and 1998, he was Professor of signal processing at Luleå University of Technology. Since 1998, he is Professor of signal processing at Lund University. His primary research interest lies in high-performance communication systems, in particular, high-data-rate wireless and twisted-pair systems. He is presently researching signal processing techniques in communication systems that use orthogonal frequency-division multiplexing (OFDM) or discrete multitone (DMT) modulation. He emphasizes the interaction between models and real systems, from the creation of application-oriented models based on system knowledge to the implementation and evaluation of algorithms.