A Byte-Erasure Method for Improved Impulse Immunity in DSL Systems using Soft Information from an Inner Code[†]

Dimitris Toumpakaris[*], Wei Yu[**], John M. Cioffi[*], Daniel Gardan[§], and Meryem Ouzzif[§]
[*]Department of Electrical Engineering, 350 Serra Mall, Stanford University, Stanford, CA 94305
[**]Electrical and Computer Engineering Department, 10 King's College Rd, U. of Toronto, Toronto, ON M5S 3G4, Canada
[§]France Télécom R&D, 2 Av. Pierre Marzin, 22 307 Lannion, France
dtouba@stanford.edu, weiyu@comm.utoronto.ca, cioffi@stanford.edu, {daniel.gardan, meryem.ouzzif}@francetelecom.com

***Abstract*—A significant portion of the end-to-end delay in high-rate DSL systems is due to the Impulse Noise protection scheme employed in order to shield those systems against random, non-stationary noise bursts of high energy that appear on the copper lines. Systems are protected from Impulse Noise using a combination of interleaving and Reed-Solomon codes. In order to lower the end-to-end delay without reducing the data rate that is available to the user, one needs to decrease the interleaver depth. This paper presents a way to achieve this reduction without compromising neither the robustness to noise bursts nor the data rate of the system. The proposed algorithm relies on the inner code used by many DSL systems and uses the metric provided by the inner code decoder at the receiver. A DMT-VDSL system is used as a particular example of the achieved reduction of the end-to-end delay.**

I. Introduction

In recent years, broadband services have been experiencing a constantly growing popularity. A large portion of these services is based on some form of the so-called Digital Subscriber Loop (DSL) standards. Early DSL systems have offered dramatically increased rates to the users through the same twisted pairs previously used by voiceband modems. Widely deployed ADSL systems currently offer typical rates of 1 Mbps, sometimes reaching 6 Mbps for customers close to the Central Office. These first-generation systems have used single-user designs, assuming worst-case crosstalk scenarios, and imposing power spectral density limits on the transmitted energy in order to avoid interfering with other DSL systems collocated in the same bundle. Therefore, in that sense they are designed conservatively and don't take full advantage of the capacity of each copper pair. Moreover, large design margins are used to guarantee trouble-free operation.

The demand for even higher rates and the potential of the existing copper network infrastructure have been driving the research for more efficient deployments of DSL systems and advanced signal processing methods. In order to achieve higher rates, more sophisticated designs such as the ones proposed in [1], [2], [3], and [4] are needed. By making use of these methods and by also reducing the loop lengths, new standards, such as VDSL, focus on increasing the data rates available to the customers to sometimes as much as 100 Mbps in order to support Ethernet networks.

For many applications that require such speeds it is also important that the end-to-end delay be kept below some acceptable limits. To a large extent, end-to-end delay is due to interleaving that is applied in order to protect DSL systems from Impulse Noise, *i.e.,* noise bursts of non-stationary nature and relatively high energy that appear on the copper lines. The noise bursts are caused by electromagnetic interference and can be due to physical phenomena, but also to electrical switches, motors and home appliances. In addition to interleaving, Impulse Noise rejection is assisted by the large margins used by current systems. As the margins are lowered to allow better use of the channel, Impulse Noise protection relies increasingly on interleaving.

Lately, there has been an increased interest in the modeling of Impulse Noise by many researchers and operators [5], [6], [7], [8], [9]. Despite the ongoing effort, because of the highly non-stationary nature and the variety of the measured noise bursts, a universally accepted model is yet to be agreed upon.

As is explained in more detail in Section II, Impulse Noise protection is achieved using a combination of Reed-Solomon (RS) codes and interleaving. In order to reduce the end-to-end delay, the interleaver depth needs to be decreased, which translates into either increasing the code redundancy, thus reducing the rate that is available to the user, or lowering the level of Impulse Noise protection. A third way to reduce the required interleaver depth is to use erasure decoding and take advantage of the fact that for a given amount of parity bytes the error-correcting capability of an RS code is doubled when erasures are used.

This property of RS codes has been used in [10] to develop an erasure method based on unused tones in a DMT-based system. In [11], a method that indicates the location of the erasures in a DSL system that satisfies certain conditions is proposed. The method works irrespective of whether an inner code is used by the DSL system, and can also be applied to Single Carrier Modulation DSL systems. This paper deals with the case where an inner code is employed. It will be shown that bytes can be reliably erased by directly using the information provided by the inner code decoder at the receiver.

Section II contains a description of the generic Impulse Noise protection scheme. Section III presents the triangular interleaver proposed in the DSL standards. The proposed method is discussed in Section IV, and its performance is quantified using simulations in Section V. Section VI summarizes the contribution of this study and presents current research topics.

II. System Model

Fig. 1 depicts the Impulse Noise protection scheme used by DSL systems. The data bytes provided by the upper layers of the transmission system are passed to an RS encoder that generates P parity bytes for every K data bytes, thus forming blocks of $N = K + P$ bytes. Reed-Solomon codes are block codes that use symbol alphabets in $GF(2^m)$. In current DSL systems, the used symbols are bytes, so RS codes in $GF(2^8)$ are used, the maximum codeword size in this case

[†]This work was supported in part by France Télécom R&D.

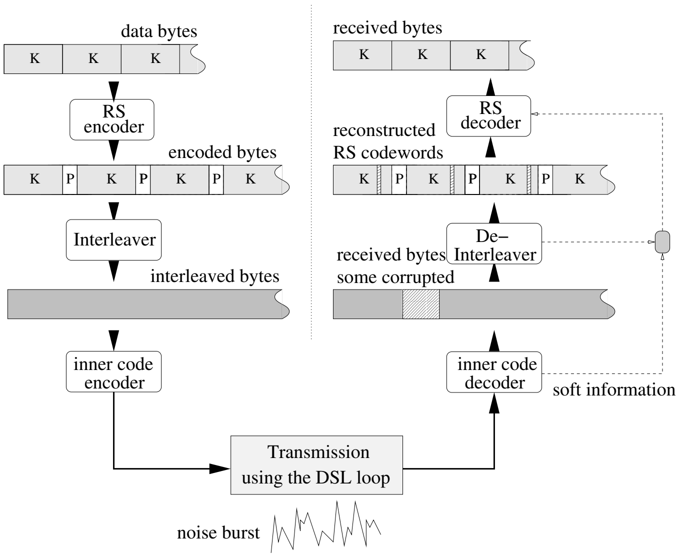

Fig. 1. The system model

being $N = 255$ bytes. Each codeword of an (N, K) RS code will be correctly decoded at the receiver provided that the number of erroneous bytes in the codeword does not exceed $\left\lfloor \frac{P}{2} \right\rfloor$. Moreover, when the location of the errors is known (which is commonly referred to as erasure decoding), up to P bytes can be in error. In general, the number of erasures e and the number of errors with unknown locations f need to satisfy the relation $e + 2f \leq P$ in order for the decoder to be able to reconstruct correctly the original codeword. Reed-Solomon codes are described in detail in error-correcting codes texts such as [12].

The encoded bytes are then interleaved in order to spread error bursts. Provided that the interleaver depth D is large enough compared to the bursts duration, this effectively creates a (DN, DK) code, therefore increasing the burst error-correction capability of the system by a factor equal to D. However, this comes at the expense of increased end-to-end delay that is proportional to D. In order to reduce the delay, D has to be decreased. This can be achieved by providing erasures to the RS decoder that help increase the number of bytes that can be corrected for a given value of parity P, allowing the use of a smaller value for D.

The interleaved bytes are then passed to the DSL system for transmission. In this paper it is assumed that the DSL system uses an inner code. Hence, the interleaved data are re-encoded before being sent to the channel. Inner codes are used to improve the immunity to white noise. For example, Wei's 4-dimensional, 16-state Trellis code that is used in Section V can provide an additional gain of up to 1.6 dB (for Bit Error Probability equal to 10^{-7}) when concatenated with an outer RS code. However, when a noise burst of large amplitude appears, even the inner code will not be able to deter the corruption of some of the data. Nevertheless, one can exploit the fact that the decoder of the inner code that is implemented at the receiver will be able to detect that a noise burst has appeared by the larger-than-usual metric that is calculated at each state transition. This information can then be passed to the outer decoder and used to mark the potentially corrupted bytes as erasures before decoding. The information passed to the RS decoder is depicted by the dotted line in Fig. 1. The deinterleaver mapping is also needed in order to correctly mark the bytes at the input of the RS decoder. The RS decoder uses the erasure information to reconstruct the original RS blocks. Provided that the erasures are correct and that the equation $e + 2f \leq P$ is satisfied, user data are successfully recovered at the receiver.

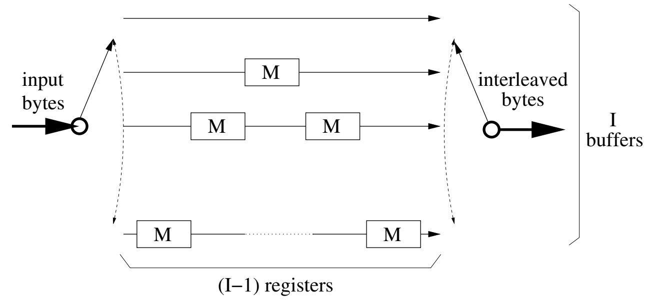

Fig. 2. Generalized Triangular Interleaver

Concatenated inner and outer coding schemes are currently widely used by many transmission and storage systems. Soft information exchange between the inner and the outer decoder is also a well-studied concept and dates back to the work of Forney [13]. The difference with previous work that has focused on evaluating the coding gain for static noise, is that in this study the information exchange between the inner and outer decoder is used to assist recovery from the high-energy, non-stationary noise impulses that appear on DSL systems.

III. Interleaving

A. The generalized triangular interleaver

To illustrate the byte-erasure technique, the generalized triangular interleaver of Fig. 2 that is proposed for VDSL systems is used [14]. The generalized triangular interleaver treats data in blocks of I bytes, where $N = kI$, k an integer, and N is the size of the RS codewords. M is a parameter that determines the interleaver depth.

Each byte of each block of I bytes is sent to a buffer with different delay. It can be seen that the i-th byte of each block is delayed by $(i-1)IM$. Thus, the largest delay is experienced by the I-th byte and is equal to $(I-1)IM$. The largest delay determines the overall delay of the system, and therefore, the end-to-end delay when using the generalized triangular interleaver is equal to $d = (I-1)IM$ bytes. The interleaver depth is equal to $D = IM$.

Typical values that are used in DSL systems are $k = 8$ for $(N, K) = (240, 224)$, which results in blocks of size $I = 30$, and $k = 4$ for $(N, K) = (144, 128)$, which yields blocks of size $I = 36$. An interleaving example is shown in Fig. 3, where $N = 6$, $k = 2$ and $I = 3$ to simplify the diagram.

B. Calculation of the required interleaver depth

Suppose that t, the correction capability of the RS code (t equals P or $\left\lfloor \frac{P}{2} \right\rfloor$ depending on whether erasures are used) sat-

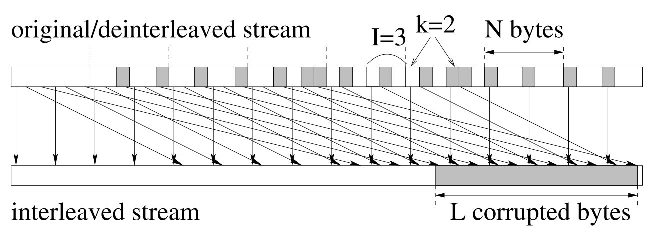

Fig. 3. Interleaving using a generalized triangular interleaver

isfies $t \geq k$, and that the i-th byte of an RS block gets mapped to a location in the interleaved stream that is corrupted by Impulse Noise. The delay of this byte is $((i-1)\ mod\ I)IM$. In the worst case, all the other $k-1$ bytes that experience the same delay will fall on the corrupted area of the interleaved stream and therefore all k bytes will be corrupted. Suppose now that the next group of k bytes in the same RS block whose delay is $(i\ mod\ I)IM$ also falls in the corrupted area of the inteleaved stream. Then, up to $2k$ bytes of the original block may be in error. By induction, it can be seen that if the error-correction capability of the RS code is equal to t bytes, then the system can tolerate up to $\left\lfloor \frac{t}{k} \right\rfloor$ such "groups" of bytes to fall on the corrupted area. Hence, the difference in the delay of bytes that are more than $\left\lfloor \frac{t}{k} \right\rfloor$ groups apart should be larger than the length of the corrupted area. Therefore, if the corrupted area has a length of L bytes, we must have

$$\left\lfloor \frac{t}{k} \right\rfloor IM > L \Rightarrow M > \frac{L}{\left\lfloor \frac{t}{k} \right\rfloor I} \Rightarrow M_{min} = \left\lceil \frac{L}{\left\lfloor \frac{t}{k} \right\rfloor I} \right\rceil.$$

Note that for a certain length L of the noise burst, the value of M that has to be used for a given level of protection depends on I. When the location of the corrupted bytes is unknown, $t = \left\lfloor \frac{P}{2} \right\rfloor$, whereas when erasures are used $t = P$, which approximately halves M, and therefore the end-to-end delay. In practice, it is safer to design a system so that the maximum number of erased bytes per RS block is smaller than P in order to allow for the detection of bytes corrupted by static noise outliers. For example, one could erase up to $P-2$ bytes in a RS codeword, in order to be able to correct one more erroneous byte in an unknown location.

IV. Using the metric of the inner code decoder

In order to demonstrate how reliability information is used to provide erasures in DSL systems, Wei's 4-dimensional, 16-state Trellis code that is used in actual DMT systems is considered [15], [16]. However, as it will become apparent by the discussion of this section, any inner code whose decoder uses some kind of metric to select the most likely paths and/or constellation symbols can be employed.

A. Wei's Trellis Code in DMT systems

DMT standards include the option of using Wei's 4D Trellis code in order to improve the system's performance. The tones are ordered and encoding is done in pairs of tones across each DMT symbol. The implementation details and the Trellis diagram of Wei's code can be found in [16]. The achievable gains for systems using the concatenated scheme consisting of Wei's inner code and an outer RS code have been analyzed in [17], and are typically between 5.2 and 5.5 dB.

In order to decode the bits of each DMT symbol a Viterbi decoder operating along the lines proposed by Wei in [15] is used. At every state transition i which corresponds to a pair of DMT tones, the cost-to-go $c(s, i)$ corresponding to each of the 16 states is evaluated, and the smallest-cost path to each state is updated. At the end of the decoding of a DMT block, each state s has an associated metric $c(I, s)$, where I the total number of pairs of tones with nonzero number of bits. The state with the smallest cost-to-go is selected, and the sequence corresponding to the smallest cost to that state constitutes the decoded data.

B. Impulse Noise

In transmission systems, the available energy is allocated to the transmitted constellation such that the probability of error be kept below some value, typically of the order of 10^{-6} or 10^{-7} for DSL systems. When coding is used, the achieved coding gain lowers the power that is required in order to achieve that error probability. However, inner codes target the isolated outliers of static noise and not the high-energy, non-stationary noise bursts. Therefore, when a system is hit by an impulse, if the impulse duration is large, or if its energy is high, the decoder will not be able to successfully reconstruct the data. This is the main reason why the inner code has to be concatenated with an error-correcting outer code such as RS codes in DSL systems. Nevertheless, although in many cases the inner code will not be able to transmit the data without some of the bits being corrupted, there will be a very reliable indication that part of the data stream has been corrupted because of the large decoder metric. For the case of the Viterbi decoder of Wei's code, if the Impulse Noise is large enough, all the state metrics will have values much larger than when only static noise is present.

In the absence of Impulse Noise, and assuming that the system has been designed to achieve a very small bit error probability, the received constellation points will be close to the originally transmitted ones. Therefore, when evaluating the costs-to-go for each state at each transition, one of them will be clearly smaller than the others. If an unusually large value of static noise gets superimposed on the transmitted symbols, the metric will increase, and the distinction between the states will become less clear. However, since such events happen with very low probability in a well-designed system, the overall metric will still be relatively low compared to the case where a noise burst is superimposed on the transmitted data stream. Moreover, the duration of the disturbance will be small. Such events, and the short bursts that will result after the inner code decoding can be corrected by the outer code without the need to provide erasures.

When a noise burst appears, the distance between received and transmitted symbols will increase significantly and for a longer period of time. By monitoring how the Viterbi metric $c(i, s)$ increases during the decoding process, or using the final value $c(I, s)$ at the end of the DMT symbol, one can conclude with very high reliability if a noise burst has corrupted the symbol. If the metric is relatively small, then nothing unusual has occurred and the decoder of the inner code will either have managed to reconstruct the original data or there will be a short error burst at the output that can be dealt with even if no erasures are used. However, if the metric exceeds some threshold, it is certain that an Impulse Noise has appeared. In this case, the inner code decoder can instruct the outer RS decoder to erase the bytes corresponding to the corrupted DMT symbol (or data block in single-carrier systems).

C. Refining the localization of the erroneous data

In DSL standards [14], [18], [19], noise bursts are assumed to be limited to a maximum length, and systems who seek to provide complete Impulse Noise protection should use interleaver depths that can result to correction of even the worst-case noise. Contrary to the case where no erasures are used, it is very important not only to mark the potentially corrupted bytes, but also to avoid over-marking data, since the total erasures may exceed the correction capability of the RS code. If the assumption that impulses don't exceed a certain length holds, there will be a limit to the number of corrupted blocks

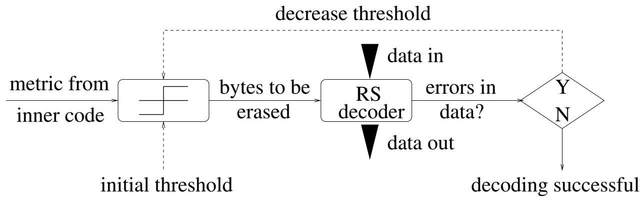

Fig. 4. Refining the metric threshold selection

and over-marking is not a risk when the inner decoder metric is used to provide erasures.

In practice, non-stationary noise protection may be more challenging. There is some indication that there may be impulses with interarrival times that are comparable to the maximum end-to-end delay that can be used. Moreover, as design margins are lowered, interference due HAM radio operators that appears outside the reserved RF bands due to nonlinearities of the circuits may interfere with the transmitted data. The duration of such interferers may well exceed the maximum Impulse Noise duration, and may result in errors even after the outer RS decoder. However, if the localization of the erroneous bytes is refined, leading to fewer erasures, one can expect to be able to improve the system performance without increasing the interleaver depth. In Section IV-B it was assumed that a whole DMT symbol is marked if an impulse appears somewhere along the tones. Correspondingly, for a coded single-carrier system, one would need to consider a sufficiently large window of data to ensure that the impulse/no impulse hypothesis is correctly evaluated. By beginning to examine smaller sections of a DMT symbol, or by reducing the window in a single-carrier system, the risk of not marking corrupted data increases. On the other hand, one cannot know in advance if the disturbance is an impulse limited to the duration specified by the standards. Thus, by using a conservative method to mark the erased bytes it is possible to over-mark the data and therefore cause decoding failure at the outer code because of the number of erasures exceeding the code parity.

One possible way to deal with such cases is to send more information to the outer decoder. More specifically, instead of only sending an erasure/no erasure decision, the inner code could be sending the metric value over a predetermined number of state transitions, or the difference of the metric value after each state transition. This is also interesting from a practical point of view, since usually Viterbi decoders only use a certain depth when calculating the state metrics, and only keep the sum of the metrics over the most recent transitions. The outer RS decoder can then erase in each codeword the bytes associated with the P largest metrics or use an iterative scheme that successively lowers the value of the metric of the bytes to be erased until correction is achieved. A generic algorithm is depicted in Fig. 4.

The choice of a good metric and threshold in order to refine the localization of the corrupted bytes is a topic of current research and still being investigated. Therefore, the simulations of Section V deal with the case where there is only an impulse/no impulse hypothesis that needs to be evaluated, with the impulse known a priori not to exceed a maximum length used for the interleaver design.

V. Simulation Results

A. Interleaver Parameters for the DMT-VDSL case

The remainder of this paper considers the case of DMT-VDSL to simulate the performance of the proposed Impulse Noise protection scheme. In DMT-VDSL, the duration of each symbol is equal to 250 μs. A noise burst of 500 μs, which is the maximum duration specified in the standards [18], [19], will affect up to 3 DMT symbols depending on the alignment of the FFT and the temporal occurrence of the noise burst. In general, the receiver's FFT spreads the impulse so that some tones of the DMT symbol will not be affected. However, some of the impulses measured in actual local loops have very high energies and can corrupt most of the bytes of each symbol. Moreover, there have been impulses measured on the network of France Télécom that exceed 500 μs. Therefore, the simulations use an interleaver depth that allows the outer code to correct noise bursts of length of up to 4 DMT symbols that appear at the deinterleaver input.

Suppose that each DMT symbol carries B bytes. Then, in the worst case, up to $4B$ consecutive bytes (belonging to 4 consecutive DMT symbols) will be corrupted. As shown in Section III, M has to satisfy

$$M > \frac{L}{\left\lfloor \frac{t}{k} \right\rfloor I} = \frac{4B}{\left\lfloor \frac{t}{k} \right\rfloor I} \Rightarrow M_{min} = \left\lceil \frac{4B}{\left\lfloor \frac{t}{k} \right\rfloor I} \right\rceil .$$

Then, the end-to-end delay is $d = MI(I-1) > \frac{4B(I-1)}{\left\lfloor \frac{t}{k} \right\rfloor}$ bytes. To convert the delay to seconds, recall that for DMT systems it takes 250 μs to transmit B bytes. Therefore,

$$d_{min} = \left\lceil \left\lceil \frac{4B}{\left\lfloor \frac{t}{k} \right\rfloor I} \right\rceil \frac{I(I-1)}{B} \right\rceil 250\ \mu s > \frac{(I-1)}{\left\lfloor \frac{t}{k} \right\rfloor} \text{ ms}.$$

B. Achieved reduction of the end-to-end delay

Simulated transmission on a DMT-VDSL system with different Impulse Noise bursts is used to verify the improvement in error correction when using erasures assisted by the inner code decoder. A data buffer is constructed and then encoded using $(N, K) = (240, 224)$, $(144, 128)$ and $(64, 48)$. Then the data are transmitted to the downstream direction using a DMT-VDSL system with parameters shown in Table I. Before being sent to the channel, each DMT symbol is encoded using Wei's code. The transmission is impaired by a deterministic noise burst that is superimposed to the received modulated data. The received data are then demodulated, decoded by an inner decoder and then passed together with the erasures information to the outer RS decoder.

Three RS code scenarios are chosen to examine three different options for the system design.

- $(240, 224)$: The Trellis coding gain ($\approx$ 1.3 dB) is used to extend the system reach. The RS parity overhead is small, leading to a relatively large end-to-end delay.
- $(144, 128)$: moderate RS parity overhead, part of the DSL loop length is exchanged for increased parity, thus reduced end-to-end delay.
- $(64, 48)$: Very large RS overhead is used. In this case, the Trellis coding gain ($\approx$ 0.2 dB) is very small and the loop length is approximately equal to that of a system not using an inner code. In this case the role of the Trellis code is to assist erasures rather than reduce the Bit Error Rate for static noise.

Due to the lack of a generally acceptable model, three representative impulses provided by France Télécom are used for the simulations. The first two represent the worst-case scenario for which the system is designed. They are of long duration and high energy and corrupt a large number of bytes of 4 consecutive DMT symbols. The third impulse lasts for 115 μs and can therefore corrupt the bytes of up to 2 DMT symbols. To facilitate comparisons with [11] and emphasize the fact that the method of this paper leads to the same reduction of the delay, the same impulses as in [11] are used.

Table II compares the end-to-end delay that is required in order to protect the system from each of the three impulses, compared to the delay when erasures are not used. It also gives the corresponding value of M for the target rate of Table I. If the designer allocates parity bytes for the correction of erroneous bytes whose location is not known, the delay will take values between the two extreme cases of no erasures or P erasures.

Clearly, the use of erasures leads to a significant reduction of the delay. In addition to abiding by the 20-ms limit imposed by the VDSL standards, the system is now immune to noise bursts with smaller interarrival times, since the byte spreading due to interleaving is less extensive. Note the lower end-to-end delay of $(144, 128)$ and $(64, 48)$ codes, due to the fact that the parity overhead is larger. In general, the delay can be reduced by allocating a larger portion of the transmitted data to parity. To achieve this goal, one can use the gain provided by the inner code, or the gain due to sophisticated interference cancellation methods such as [1], [2], [3], and [4]. Moreover, as the impulse gets shorter, the improvement due to erasures is less dramatic. This is due to the fact that by erasing a whole symbol, one may mark bytes that do not belong to the corrupted area. However, if one designs for the worst-case impulses, erasures have to be employed in order to be able to achieve a considerable reduction of the end-to-end delay.

VI. Conclusion

In this study, the end-to-end delay of a DSL system is reduced by using information from the decoding of the inner code to erase potentially corrupted bytes before their being passed to the outer RS decoder. The information comes from the metric used by the inner code decoder. As an example, the cost-to-go evaluated by a Viterbi decoder of Wei's 4D code used in DMT-DSL systems was considered. It was shown that the delay can be reduced considerably with respect to the case where no cooperation exists between the inner and outer code, and that the robustness of the system against noise impulses is not compromised.

To deal with the case of lower energy but higher duration noise bursts, as well as RF interference due to HAM operators, a refinement of the presented method which is a topic of current research was briefly described. Such methods are expected to find use as design margins in DSL systems are reduced in order to allow more efficient use of twisted pairs in data transmission.

TABLE I

DMT-VDSL parameters used for the simulations

parameter	value
transmission	Downstream FTTEx
rate (to upper layer)	23.168 Mbps
tones + cyclic prefix	4096 + 320 (complex)
tone spacing	4.3125 kHz
symbol error probability	10^{-7}
band plan	998
PSD mask	Pex.P2.LT.M1 of [19]
maximum power	11.5 dB
margin	6 dB
RS & Wei coding gain	5.2 dB [17]
RS only coding gain coding gain (no inner code)	3.86 dB for $(240, 224)$ code, 4.22 dB for $(144, 128)$ code, 5.02 dB for $(64, 48)$ code.
background static noise	Model E of ETSI

TABLE II

Minimum delay for Impulse Noise protection

noise number	code (N, K)	delay no erasures	M	delay erasures	M
1	(240, 224)	27.75 ms	106	14.75 ms	56
	(144, 128)	16.75 ms	44	9.25 ms	24
	(64, 48)	7.5 ms	25	4.25 ms	14
2	(240, 224)	29.25 ms	112	14.75 ms	56
	(144, 128)	17.75 ms	47	9.25 ms	24
	(64, 48)	8 ms	27	4.25 ms	14
3	(240, 224)	12 ms	46	7.5 ms	28
	(144, 128)	6.25 ms	16	4.75 ms	12
	(64, 48)	2.75 ms	9	2.25 ms	7

References

[1] G. Ginis and J.M. Cioffi, "Vectored Transmission for Digital Subscriber Line Systems," *IEEE J. on Select. Areas Commun. special issue on twisted-pair transmission*, vol. 20, no. 5, pp. 1085-1104, June 2002.

[2] W. Yu, G. Ginis and J.M. Cioffi, "An Adaptive Multiuser Power Control Algorithm for VDSL," *IEEE J. on Select. Areas Commun. special issue on twisted-pair transmission*, vol. 20, no. 5, pp. 1105-1115, June 2002.

[3] M. L. Honig, K. Steiglitz and B. Gopinath, "Multichannel signal processing for data communications in the presence of crosstalk," *IEEE Trans. Commun.*, vol. 38, no. 4, pp. 551–558, April 1990.

[4] M. L. Honig, P. Crespo and K. Steiglitz, "Suppression of near- and far-end crosstalk by linear pre- and post-Filtering," *IEEE J. on Select. Areas Commun.*, vol. 10, no. 3, pp. 614–629, April 1992.

[5] I. Mann, S. McLaughlin, and W. Henkel, "Impulse Generation with appropriate amplitude, length, inter-arrival, and spectral characteristics," *IEEE J. on Select. Areas Commun.*, vol. 20, no. 5, pp. 1–8, May 2002.

[6] R. Kirkby, D. B. Levey and S. McLaughlin, "Statistics of Impulse Noise," ETSI UK TM6, TD18, TD19, TD20, TD21. Edinburgh, UK, 1999.

[7] I. Mann, S. McLaughlin and D. B. Levey, "A new statistics for Impulse Noise measurement," ETSI UK TM6, TD 55. Amsterdam, The Netherlands, 1999.

[8] B. Rolland, D. Bardouil, F. Clérot, D. Collobert, D. Gardan, J. Le Roux, "Classification du bruit impulsif par méthode d'apprentissage non supervisé," Internal report RP/FTR&D/7133, France Télécom, December 2000.

[9] D. Collobert et al., "Proposal for Impulse Noise Classification," ETSI TM6, TD12. Sophia Antipolis, France, April 2002.

[10] F. Sjöberg, "The Zipper Duplex Method in Very High-Speed Digital Subscriber Lines," Ph.D thesis, Luleå University of Technology, April 2000.

[11] D. Toumpakaris, W. Yu, J. M. Cioffi, D. Gardan and M. Ouzzif, "A Simple Byte-Erasure Method for Improved Impulse Immunity in DSL," *accepted for presentation in ICC 2003*, in press.

[12] S. Wicker, *Error Control Systems for Digital Communication and Storage*, Prentice Hall, 1995.

[13] G. D. Forney, "Generalized Minimum Distance Decoding," *IEEE Trans. Inform. Th,* vol.12, no. 2, pp. 125–131, April 1966.

[14] T1E1.4/2000-013R4. Very-high bit-rate Digital Subscriber Lines (VDSL) Metallic Interface, Part 3: Technical Specification of a Multi-Carrier Modulation Transceiver.

[15] L. F. Wei, "Trellis-Coded Modulation with Multidimensional Constellations," *IEEE Trans. Inform. Th.*, vol. 33, no. 4, pp. 483–501, July 1987.

[16] ITU G.992.1 Recommendation (ex. G.dmt): Asymmetrical Digital Subscriber Line (ADSL) Transceivers.

[17] T. N. Zogakis, J. T. Aslanis Jr., and J. M. Cioffi, "Analysis of a Concatenated Coding Scheme for a Discrete Multitone Modulation System," *IEEE MILCOM Conference Record*, vol. 2, pp. 433–437, 1994.

[18] T1E1.4/2000-009R3. VDSL Technical Specification Part 1. Functional Requirements and Common Specification.

[19] ETSI TS 101 270–1 v1.2.1 (1999–10). Transmission and Multiplexing (TM); Access transmission systems on metallic access cables; Very high speed Digital Subscriber Lines (VDSL); Part 1: Functional Requirements.