A Simple Baseband Transmission Scheme for Power Line Channels

Raju Hormis, Inaki Berenguer, and Xiaodong Wang

Abstract—**We propose a simple pulse-amplitude modulation (PAM)-based coded modulation scheme that overcomes two major constraints of power line channels, viz., severe insertion-loss and impulsive noise. The scheme combines low-density parity-check (LDPC) codes, along with cyclic random-error and burst-error correction codes to achieve high-spectral efficiency, low decoding complexity, and a high degree of immunity to impulse noise. To achieve good performance in the presence of intersymbol interference (ISI) on static or slowly time-varying channels, the proposed coset-coding is employed in conjunction with Tomlinson–Harashima precoding and spectral shaping at the transmitter. In Gaussian noise, the scheme performs within 2 dB of unshaped channel capacity at a bit-error rate (BER) of 10^{-11}, even with (3,6)-regular LDPC codes of modest length (1000–2000 bits). To mitigate errors due to impulse noise (a combination of synchronous and asynchronous impulses), a multistage interleaver is proposed, each stage tailored to the error-correcting property of each layer of the coset decomposition. In the presence of residual ISI, colored Gaussian noise, as well as severe synchronous and asynchronous impulse noise, the gap to Shannon capacity of the scheme to a Gaussian-noise-only channel is 5.5 dB at a BER of 10^{-7}.**

Index Terms—**Coded modulation, coset codes, impulse noise, intersymbol interference (ISI), low-density parity-check (LDPC), power line communications, precoding.**

I. Introduction

OVER the past decade, advances in coding, equalization, and very large scale integration (VLSI) design have combined to enable spectacular increases in throughput over power lines cables. This is despite the severe constraints the medium imposes, viz., severe signal attenuation (insertion loss) over long cables, impulse noise, and intersymbol interference (ISI). In particular, impulse noise is a severe impairment, and occurs in the form of time-varying periodic noise synchronized to the line frequency, periodic but asynchronous noise caused by switching power supplies, and asynchronous noise caused by random switching transients in the network (cf. [8], [31], and [32]). Electromagnetic interference (EMI), in the form of narrowband sinusoidal noise from radio and TV sources, is also an impairment. Furthermore, transmission over power lines is constrained by statutory electromagnetic compatibility (EMC) emissions constraints, which restrict the total transmit power and power spectral density (PSD) during transmission.

In this paper, we propose a simple pulse-amplitude modulation (PAM)-based coset-coding scheme to overcome these impairments. For brevity, we focus on static or slowly time-varying channels, which has been shown to be a reasonable assumption in [7]. However, extensions to the rapidly time-varying case are also briefly outlined. The modulation scheme offers high-spectral efficiency, immunity to multiple impulse noise sources, good coding gains, but yet, requires low complexity overall. The idea of coding with cosets within a lattice framework was first generalized in [13] and [14]. An important result, proven in [15] and [29], is that coset codes can achieve the sphere bound—channel capacity without shaping—with simple one-dimensional lattices, and with two or three levels of coset partitioning. Motivated by these ideas, the scheme proposed here for power line channels is based on a three-level coset decomposition with different codes at each layer of the decomposition. Viewing the bottom layer as a Gaussian channel at low signal-to-noise ratio (SNR), the scheme relies on (3,6)-regular low-rate low-density parity-check (LDPC) codes of 1000–2000 bits [22] for steep bit-error rate (BER) reduction. Meanwhile, the middle layer is treated as a binary symmetric channel (BSC) that is coded with hard-decision random and burst error-correcting cyclic codes. In particular, a high-rate Reed–Solomon (RS) code is applied here to protect against random and phased-burst errors. By virtue of its large intra-coset distance, the top layer of the decomposition can be viewed as a BSC channel that is only vulnerable to burst noise. Both RS codes and single-burst-error correcting codes from [23] are investigated. On an additive white Gaussian noise (AWGN)-only channel, the scheme performs within 1.5 dB of the sphere-bound at a BER of 10^{-7}, and within 2 dB at a BER of 10^{-11}. At this BER, the corresponding coding gain over uncoded PAM is about 8.5 dB.

In work related to LDPC-based coded modulation, a coset-coding scheme was proposed in [11] for discrete multitone (DMT) modulation over digital subscriber lines. The authors demonstrated that coset codes can be constructed with relatively short LDPC binary codes, thus keeping latency to a minimum. Independently of this work, one of us proposed a coset-coding scheme for twisted-pair transmission at 10 Gbit/s [25], with LDPC codes on the order of 2000 bits. Here, performance at a BER of 10^{-12} was 2.5 dB from the sphere-bound. As noted by one of the reviewers, a DMT-based LDPC coset-coding scheme was also proposed recently in [3] for ISI-constrained channels; the scheme makes use of RS component codes at high layers due to their low complexity and their well-known construction at high rates. However, the motivation of our proposal is broader, as we view the coset

Manuscript received April 29, 2005; revised December 22, 2005 and February 26, 2006. This work was supported in part by the U.S. National Science Foundation under Grant CCR-0207550 and Grant DMS 0244583. This work was presented in part at the International Communications Conference (ICC'06), Istanbul, Turkey, June 2006.

R. Hormis and X. Wang are with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: raju@ee.columbia.edu; wangx@ee.columbia.edu).

I. Berenguer is with the Laboratory for Communication Engineering, Cambridge University, Cambridge CB3 0FD, U.K. (e-mail: ib226@cam.ac.uk).

Digital Object Identifier 10.1109/JSAC.2006.874429

decomposition with different noise characteristics at each layer. The partition between modulation, coding, and equalization is also different from what is proposed here. For example, in [3], the bottom layer is actually a concatenation of an irregular LDPC code (of length 10^5 bits) and an ensemble of repeat-codes of various rates. This is to support the different capacities of each subcarrier of DMT. The problem is avoided with the PAM-based scheme proposed here, where one designs the coset code for a single AWGN channel without loss of optimality [17]. In other recent work related to coset-coded modulation, bit-interleaved coded modulation (BICM) and multilevel coding (MLC) schemes with LDPC codes were investigated in [21]. With quasi-regular LDPC component codes on the order of 10^6 bits, performance within 0.1–0.2 dB of the sphere-bound was achieved.

To extend the burst error-correction ability of the proposed coding schemes, the well-known technique of interleaving is applied. Uniform interleaving is widely used in concatenated coding where, for example, a soft-decision code-like trellis-coded modulation (TCM) is concatenated with an RS code, separated by a byte-interleaver [26]. The RS code in such schemes corrects burst-errors that are left uncorrected by, or sometimes caused by, the Viterbi decoder. However, the scheme proposed in this paper differs in the sense that interleavers are used at each layer of a coset decomposition independently, each tailored to different properties of the component codes and noise at different layers. Simulation results show good performance with sufficient interleaver depth. In the presence of nonstationary impulse noise, the scheme operates at a gap of 5.5 dB to Gaussian-noise-only channel-capacity at a BER of 10^{-7}. The gap to true channel-capacity in impulse noise could not be computed.

The PAM-based baseband modulation proposed here offers advantages over other modulation schemes that utilize DMT transmission [2], [4], or spread-spectrum techniques [10]. These will be discussed further in Section II-C. In terms of signaling on a static or slowly varying channel with ISI, a well-known technique to mitigate ISI is the combination of ideal decision-feedback equalizer (DFE), spectral shaping at the transmitter, and noise-whitening matched-filter. In Gaussian noise and under the condition of zero excess-bandwidth, this combination of techniques is asymptotically capacity-achieving at high SNR (cf. [12] and [17]). In particular, we employ Tomlinson–Harashima (TH) precoding, which asymptotically approaches the sphere-bound of the channel at high SNR and large constellations. We integrate the proposed coset-coding with these techniques to show that good performance can be achieved on the power line channel, even in the presence of ISI, severe impulse noise, and colored noise. On channels with rapidly time-varying transfer functions, the scheme can be extended to adaptive and turbo-equalized receivers.

The remainder of this paper is organized as follows. In Section II, the power line system model is outlined, along with details of the channel and noise. We also motivate the proposed transmission scheme in this section. Section III elaborates on the design of the proposed coset code in Gaussian noise as a first step to designing for the power line channel. In Section IV, the scheme is augmented to handle synchronous and asynchronous impulse noise. Simulation results are also presented. Section V concludes this paper.

II. System Descriptions

In this section, we review the transmission model and channel conditions for power line communications, focusing on last-mile access over low-voltage lines.

A. Channel Transfer Function

One of the major impairments of PLC channels is its insertion loss (signal attenuation) with increasing distance. The length of typical "last-mile" access power lines is on the order of 150 m, although it varies among countries. A major drawback of power lines, compared with other kinds of cables, is that the cable follows a bus topology, rather than a point-to-point connection. Each power line connecting each house or main to the bus (branch) can have a different terminating impedance. Terminations (e.g., open mains or connected appliances) represent a complex impedance causing reflections (return loss), and consequently, results in a multipath channel at the receiver. The more branched the network is, the larger the number of paths. Moreover, longer paths experience higher attenuations since the signals travel longer distances. Thus, the frequency response of the PLC multipath channel $H(f)$ can be approximated by a sum of N paths [33]. The sum accounts for multipath propagation and frequency-selective fading, viz.

$$H(f) = \sum_{i=1}^{N} \underbrace{g_i}_{\text{weighting}} \underbrace{e^{-(a_0+a_1 f^k)d_i}}_{\text{attenuation}} \underbrace{e^{-j2\pi f \frac{d_i}{v_p}}}_{\text{delay}} \tag{1}$$

where g_i represents a weighting factor along path i with distance d_i; a_0, a_1 are attenuating parameters; k is the exponent of the attenuation, usually in the interval 0.2 to 1. The last term represents the propagation delay, with v_p denoting the velocity of propagation. Typical values of a_0, a_1, and k are given in [27] and [33].

In this paper, we consider values of g_i, d_i, N, a_0, a_1, and k that represent a typical reference channel for last-mile access based on three-phase underground distribution grids using PVC isolated cables, whose parameters are given in [27, channel 3] and are based on real measurements in Germany [33]. Channel 3 represents a hostile channel consisting of a 210 m line with eight branches, and hence multiples sources of reflected signal power. The impulse response lasts on the order of 10 μs. The frequency response of this channel is shown in Fig. 1, along with a more benign channel of length 100 m with no branches ([27, channel 1]).

In general, most channels exhibit long-term variations in the transfer function. In [7], the channel transfer function was observed to also exhibit small variations that were periodic with line frequency; however, the authors showed that even these channels can be modeled as a sequence of static channels. Hence, we focus on slowly time-varying channels in this paper. An important characteristic that we depend on is the symmetry of the transfer function (cf. [5]), which holds true when the terminating impedances of the transceivers are identical at both ends of the link. This property permits transmitter-side techniques, as will be shown in Section II-C. Techniques for rapidly time-varying channels are also discussed.

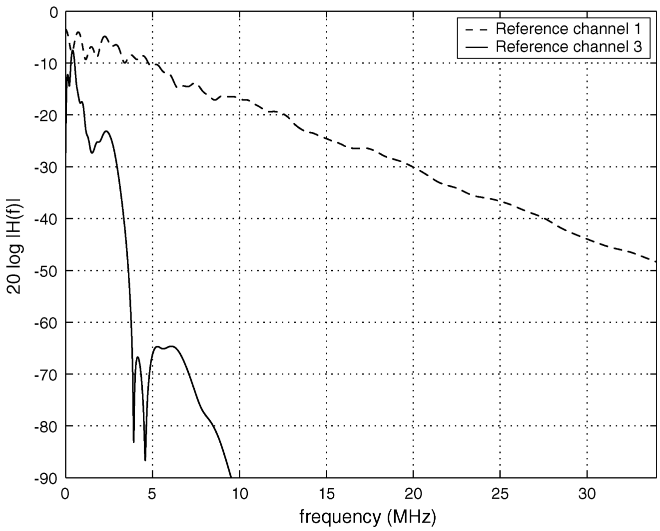

Fig. 1. Frequency response of PLC access reference channels 1 and 3.

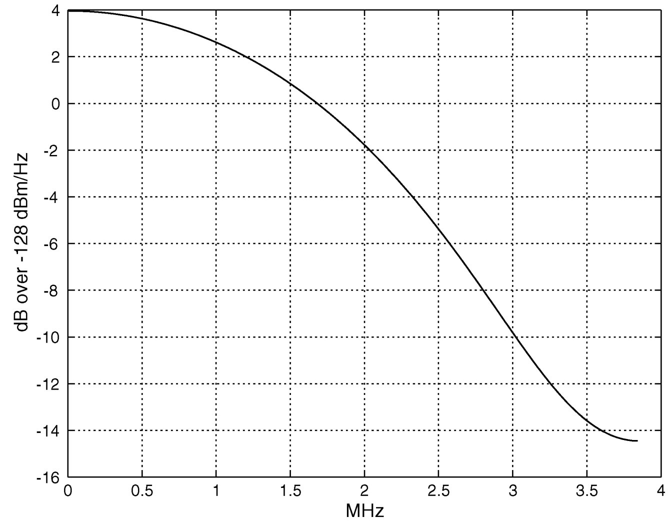

Fig. 2. PSD mask of colored Gaussian noise.

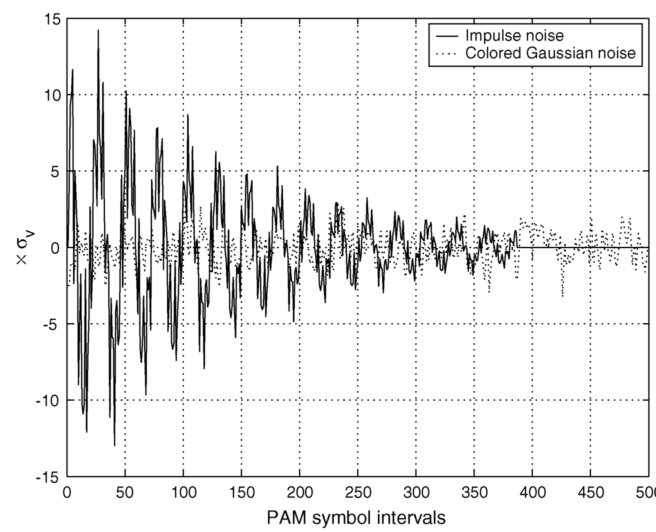

Fig. 3. Realizations of impulse noise and colored Gaussian noise. Amplitudes relative to σ_v of Gaussian noise.

B. Additive Noise

A comprehensive analysis in [32] characterized the noise sources that afflict power lines. The authors showed that the various noise sources can be classified broadly into the following categories.

a) Colored Gaussian noise.
b) Narrowband sinusoidal noise (EMI) that originates from commercial AM, FM, and ham radio sources.
c) Periodic impulsive noise that is synchronous to the mains (i.e., every AC cycle) originated by transients in appliances connected to the power lines.
d) Periodic impulse noise that is asynchronous to the mains, caused by switching power supplies.
e) Asynchronous and aperiodic impulsive noise usually caused by random switching transients.

The noise sources d) and e) are highly time-varying, with their properties changing in microseconds. In this paper, we consider all noise sources mentioned above, except b). While narrowband EMI is a significant constraint in practical schemes, for this paper, we assume that the narrowband interference can be mitigated with a frequency notch, using a combination of spectral shaping at the transmitter and noise-whitening matched-filter at the receiver. Meanwhile, the Gaussian noise is assumed to be strongly colored, with higher energy at low frequencies [31]. The spectral shape of the colored noise, relative to a level of -128 dBm/Hz, is shown in Fig. 2.

1) Synchronous Impulse Noise: It has been measured that a high percentage of the impulsive noise occurs periodically and synchronously to the mains. In general, the impulsive noise consists of a collection of damped sinusoids [8], with higher content in the low frequencies. The periodic impulses can be modeled as a collection of I_s damped sinusoids

$$n_s(t) = \sum_{i=1}^{I_s} A_i \sin\left(2\pi f_i (t - t_{\mathrm{arr},s}) + \alpha_i\right) \times e^{-\frac{t-t_{\mathrm{arr},s}}{\tau_i}} \Pi\left(\frac{t - t_{\mathrm{arr},s}}{t_{w,s}}\right) \quad (2)$$

where f_i is the "pseudofrequency" of the sinusoid, and α_i the phase, of the ith damped sinusoid. A plot of a single burst from such an impulse train is shown in Fig. 3. $\Pi(t)$ is defined as a square pulse of duration $t_{w,s}$ s, with constant amplitude in the interval $0 < t \leq 1$ and zero elsewhere. $t_{\mathrm{arr},s}$ is the periodic arrival time, and A_i denotes the amplitude of the ith sinusoid. We assume $A_i \sim \mathcal{N}(0, G_i \sigma_v^2)$, $i = 1 \dots I_s$, where G_i represents the increase over the variance of Gaussian background noise σ_v^2, and can range from 20–30 dB. The gain G_i of sinusoids at higher pseudofrequencies is selected to match the typical low-frequency content observed in impulsive noise measurements, usually below 1 MHz. The term τ_i denotes the damping factor. Meanwhile, the pulse amplitude equals the standard deviation of the background noise, i.e.,

$$\Pi\left(\frac{t - t_{\mathrm{arr},s}}{t_{w,s}}\right)\Bigg|_{t=t_{\mathrm{arr},s}} = \sigma_v. \quad (3)$$

In [32], impulses of approximately $t_{w,s} = 50\ \mu$s have been measured, and this value is used in the simulations. In [8], pseudofrequencies were characterized from 500 KHz to 3 MHz. In this paper, we consider three component sinusoids $(I_s = 3)$, with pseudofrequencies of 300 KHz, 2 and 3.5 MHz.

2) Asynchronous Impulse Noise: The combination of all impulsive noise sources that are asynchronous to the main frequency can be modeled as a sum of damped sinusoids as in (2), but where arrival time $t_{\mathrm{arr},a}$ is modeled as a random variable [32]. The asynchronous bursts are usually caused by switching transients. Let $t_{IAT,a} = t_{\mathrm{arr},a}^{(p)} - t_{\mathrm{arr},a}^{(p-1)}$ denote the interarrival time between consecutive bursts of asynchronous impulse noise, viz., burst p and $p-1$. Then, as discovered in [32], $t_{IAT,a}$ can be modeled with an exponential distribution. In the simulations, we select $t_{IAT,a}$ to be exponentially distributed with mean of 100 ms. We assume the impulse width $t_{w,a}$ to be constant, approximately 100μs. However, the amplitudes of the sinusoids A_i, $i = 1 \ldots I_a$ remain Gaussian distributed, as in Section II-B1.

C. Transmission Model

Nearly Static Channels: We employ a simple baseband PAM-based scheme in this paper, with an emphasis on static or slowly time-varying channels. In stationary Gaussian noise and under the condition of zero excess-bandwidth, a PAM-based scheme—when combined with ideal DFE, spectral shaping at the transmitter, and noise-whitening matched-filter—is asymptotically capacity-achieving at high SNR (cf. [17]). On the other hand, the impulse noise statistics are time-variant on the order of a few microseconds [32]; this makes it difficult to compute even the capacity of such a channel.

To simplify the design of a transmitter in impulse noise, we take a decidedly suboptimal approach. First, the shaping transmit-filter, equalizer and matched-filter are computed with well-known methods for an ISI-constrained Gaussian channel. For slow variations of the channel, periodic training sequences can be transmitted to update the equalizer and shaping filters via an adaptive update algorithm (cf. [19]). For very gradually changing channels, decision-directed updates should suffice, obviating the need for a training sequence. In summary, these techniques present a flat AWGN channel to a channel decoder, which greatly simplifies the design of a coding scheme (Section III). The code is then augmented to protect against nonstationary impulse noise (Section IV).

Rapidly Time-Varying Channels: For channels with short-term variation, a combination of adaptive and iterative (turbo) equalization is warranted, as transmitter-side shaping and pre-equalization are impractical. Periodic adaptation of the equalizer is needed to account for the channel variation. Furthermore, to improve the performance of the equalizer, iterative schemes have been proposed (cf. [30]). The soft-output extrinsic information from a channel-decoder can be used to update the equalizer, and vice versa, in an iterative manner. However, for brevity, we do not investigate these techniques in this paper.

Modulation: The PAM-based baseband scheme employed here offers advantages over other power line schemes that utilize DMT transmission [2], [4], or spread-spectrum techniques [10]. Besides simplicity and low latency, PAM constellations have low peak-to-average (PAR) ratios compared to multicarrier schemes. This eases the design of the analog front-end of transceivers, and also eases EMC compliance. Further, the channel-shortening needed for multicarrier schemes via time-domain equalization (TEQ) is obviated [26].

In terms of baseband modulation compared with carrier modulation, a baseband scheme has the advantage of operating in the frequency region exhibiting least insertion loss over a cable (cf. Fig. 1). Furthermore, at high frequencies, a power line cable that is designed to operate at 50–60 Hz starts to behave like an inefficient antenna [20]. A baseband scheme would naturally occupy the spectral region of least electromagnetic leakage, enhancing EMC compliance. The choice of baseband spectrum also minimizes exposure to external EMI from TV and radio sources. However, a disadvantage is that the signals are exposed to severe impulse noise due to line currents operating at 50–60 Hz and at harmonics thereof. Overall, the proposal resembles the PAM-based schemes used for T1/E1 lines in North America and Europe (via symmetric DSL, g.SHDSL [1]).

Constellation Shaping: For simplicity, we do not address constellation shaping, although shaping schemes can be applied to obtain additional shaping gain in Gaussian channels. Here, the optimal N-dimensional shaping region is the well-known N-sphere, which can yield up to $\pi e/6$ (1.53 dB) of shaping gain as $N \to \infty$ for large constellations [17]. However, in this paper, both the coset code and the pre-equalization are designed to approach only the sphere-bound of the channel, discounting gains that can be achieved by shaping. Furthermore, for impulse-noise channels, the exact shaping loss is not known.

Transmitter Spectral Shaping: Let x_n represent a sequence of transmitted PAM symbols with power spectral density $\mathcal{S}_{xx}(f)$. Let $T(f)$ represent a spectral-shaping filter, designed to achieve the optimal water-filling spectrum for the power line channel $H(f)$. The latter was defined in (1). In the system proposed here, $T(f)$ is used to shape x_n prior to transmission, although this might not be practical for an actual power line system (due to EMC compliance requirements). The filter also inserts a spectral null at half-baud rate to ensure zero excess-bandwidth. Since $H(f)$ decreases steeply at high frequencies, this can be done with simple first-order filters with negligible loss in capacity. Since the system is transformer-coupled to the power network, we also insert a spectral null at DC to minimize power loss. Ideal low-pass filtering for anti-aliasing and noise-rejection is assumed at the receiver. Let $P(f)$ denote the pulse response of the combined system, i.e.,

$$P(f) = T(f)H(f) \tag{4}$$

and let p_n denote the time-domain impulse response. Then, we can write the signal model for the proposed scheme quite simply as

$$y_n = p_n \otimes x_n + v_n \tag{5}$$

where $\otimes$ denotes convolution and v_n represents the additive colored Gaussian noise only. Let the PSD of the latter be denoted

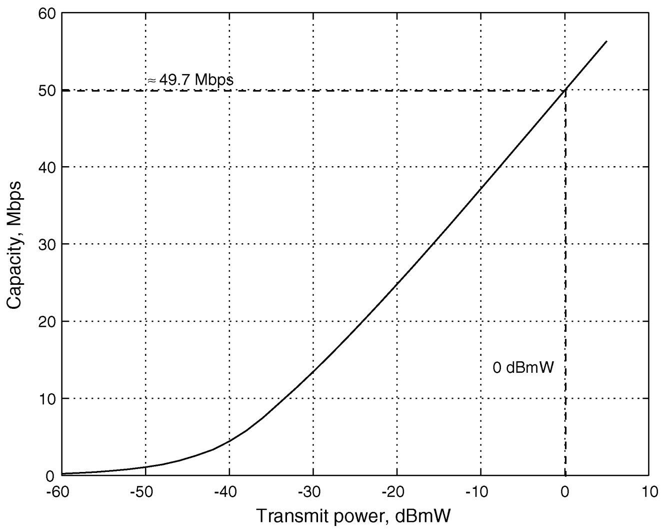

Fig. 4. Capacity of "channel 3" under optimum water-filling transmit spectrum.

by $\mathcal{S}_{vv}(f)$. Then, given a total power constraint P_T such that

$$\int_{\mathcal{B}} \mathcal{S}_{xx}(f)\, |T(f)|^2\, df \leq P_T \text{ mW} \tag{6}$$

we can write optimal water-filling spectrum [17] as

$$T(f) = \begin{cases} K - \frac{|H(f)|^2}{\mathcal{S}_{vv}(f)}, & f \in \mathcal{B} \\ 0 & , \quad f \notin \mathcal{B} \end{cases} \tag{7}$$

where $\mathcal{B}$ is a capacity-achieving region that must be computed, and K is a constant chosen such that (6) is satisfied.

Channel Capacity: Given capacity-achieving region $\mathcal{B}$, and noise PSD $\mathcal{S}_{vv}(f)$, the capacity of the frequency-selective channel, considering *only* colored Gaussian noise, is now given by

$$C = \frac{1}{2} \int_{\mathcal{B}} \log_2 \left(1 + \frac{\mathcal{S}_{xx}(f)\, |T(f)H(f)|^2}{\mathcal{S}_{vv}(f)} \right) df \text{ b/dim.} \tag{8}$$

Consider "channel 3" of Fig. 1, with colored Gaussian noise PSD of Fig. 2. The optimum water-filling capacity of this channel is computed with (8) and is shown in Fig. 4 for various values of transmit power. For a nominal transmit power of 0 dBmW, the graph shows a capacity of ≈49.7 Mbit/s, achieved over a frequency region $\mathcal{B} : 0 \leq f \leq 3.85$ MHz.

Equalization: The well-known minimum mean-squared error (MMSE) DFE is used to mitigate ISI. Let $B(z)$ denote the feedback filter of the DFE that cancels post-cursor ISI, assuming perfect decision-feedback. Due to the difficulty of combining DFEs with block codes, the proposed scheme makes use of the well-known TH precoding [12], [17]. This entails using $B(z)$ in a feedback loop at the transmitter to mitigate post-cursor ISI *a priori*, as shown in Fig. 8. The symmetry of the channel transfer function proven in [5] permits precoding. The TH-precoding induces a small transmitter power-penalty; for an M-PAM constellation, the penalty has been shown to be $(M^2/M^2 - 1)$, which is asymptotically negligible for large constellations. However, the TH-precoder also causes shaping loss by up to $\pi e/6 \approx 1.53$ dB if constellation shaping were employed (which is not the case in this paper). To avoid the shaping loss in order to achieve Shannon capacity, a practical alternative is Laroia precoding, which is asymptotically capacity-achieving at high SNRs with large constellations [17].

A receiver filter, $W(z)$ in Fig. 8, denotes the noise-whitening matched-filter of the DFE. This filter also mitigates precursor ISI. Since the system is constrained to occupy zero excess-bandwidth, the receiver is invariant to sampling phase and $W(z)$ can be a simple baud-spaced equalizer. Under this assumption, we are now left with almost-Gaussian residual ISI and whitened Gaussian noise at the input to the channel decoder [16]. This motivates our approach of designing the coding scheme for the AWGN channel first.

III. Coding in Gaussian Noise

In this section, we propose a coset-coding technique that combines both bandwidth-efficiency and near-sphere-bound performance in the presence of Gaussian noise. The latter assumption holds true under the conditions mentioned in the previous section. However, in Section IV, we also consider impulse noise. The coding scheme proposed in this section is based on the notion of sphere-bound-achieving coset codes, which were investigated in [15] and [29].

A. Code Structure

We start with a brief summary of lattices and multilevel coset-codes, and we refer the reader to the comprehensive treatments in [6], [13], and [14]. An N-dimensional *lattice* Λ can be viewed as an infinite set of uniformly spaced points in Euclidean space $\mathbb{R}^N$. A *sublattice* Λ' is a uniformly-spaced subset of the points of Λ. The sublattice Λ' is said to induce a partition, $\Lambda|\Lambda'$, of the infinite lattice Λ. A *partition chain*, $\Lambda|\Lambda'|\Lambda'' \cdots$, is a sequence of lattices such that each is a sublattice of the previous one, i.e., $\Lambda \supseteq \Lambda' \supseteq \Lambda'' \cdots$. Lastly, the schemes in this paper use block codes, which will require m-dimensional Cartesian products of lattices. This is denoted by $\Lambda^m \triangleq \Lambda \otimes \Lambda \otimes \ldots \otimes \Lambda$. Cartesian products of sublattices, $(\Lambda')^m$, $(\Lambda'')^m$, and so on, are defined similarly.

The modulation scheme of this paper uses PAM constellations, which can viewed as subsets of lattices. Formally defined, the PAM constellation can be viewed as a finite set of points belonging to a translate of the N-dimensional lattice Λ and bounded by a rectangular shaping region $\mathcal{S}$. The constellation can be expressed as $(\Lambda + \Omega) \cap \mathcal{S}$, where $\Omega \in \mathbb{R}^N$ is a translation vector. Ω is selected to center a constellation symmetrically around the origin. The sublattices of a lattice constellation are similarly bounded by $\mathcal{S}$. Consider a sublattice Λ' of Λ. A coset of Λ' can be defined as a translation of Λ' by λ, such that

$$\Lambda' + \lambda \triangleq \{\mathbf{x} = \mathbf{u} + \lambda | \mathbf{x} \in \Lambda, \quad \mathbf{u} \in \Lambda', \quad \lambda \in [\Lambda|\Lambda']\}. \tag{9}$$

$[\Lambda|\Lambda']$ represents the set of translates that satisfies (9). The coset partitions $\Lambda'|\Lambda'', \Lambda''|\Lambda''', \cdots$ of a multilevel partition chain can be defined similarly.

We can now define coset codes formally. Consider a three-level lattice partition chain $\Lambda|\Lambda'|\Lambda''$. Let $\mathbf{G}_{\Lambda|\Lambda'}$ and $\mathbf{G}_{\Lambda'|\Lambda''}$ denote generator matrices of block codes that, respectively, generate codewords $\mathbf{c}_{\Lambda|\Lambda'}$ and $\mathbf{c}_{\Lambda'|\Lambda''}$ over alphabets $[\Lambda|\Lambda']$ and $[\Lambda'|\Lambda'']$. A third matrix $\mathbf{G}_{\Lambda''}$ generates codewords $\mathbf{c}_{\Lambda''}$ that selects m points from a subset $\{(\Lambda''+\Omega)\cap\mathcal{S}\}$. Now, a coset-code $\mathbf{L}$ can be defined as a set of codewords selected such that

$$\mathbf{L} \triangleq \{\mathbf{x} = \mathbf{c}_{\Lambda|\Lambda'} + \mathbf{c}_{\Lambda'|\Lambda''} + \mathbf{c}_{\Lambda''} | \mathbf{x} \in \Lambda^m\}. \quad (10)$$

The examples in this paper use a three-level coset partition over $\mathbb{Z}$. The bits mapped on $\{(\Lambda'' + \Omega) \cap \mathcal{S}\}$ are left uncoded in Gaussian noise, but coded for impulse noise. If the rates of each component code of the coset decomposition are $R_{\Lambda|\Lambda'}$, $R_{\Lambda'|\Lambda''}$ and $R_{\Lambda''}$, it is easy to see that the coding rate of $\mathbf{L}$ is given by

$$R(\mathbf{L}) = \frac{1}{N}\left[R_{\Lambda|\Lambda'}\log_2|\Lambda|\Lambda'| + R_{\Lambda'|\Lambda''}\log_2|\Lambda'|\Lambda''| + R_{\Lambda''}\log_2|(\Lambda''+\Omega)\cap\mathcal{S}|\,\right] \text{ b/dim.} \quad (11)$$

B. Code Construction in Gaussian Noise

In this section, we discuss the code construction, coset decomposition, and choice of coding rate for a slowly time-varying power line channel.

Capacity Considerations and Rate Allocation: Let $\mathrm{C}(\mathbf{L})$ denote the capacity of a coset code $\mathbf{L}$ over a lattice partition $\Lambda|\Lambda'|\Lambda''$, and let $C_{\Lambda|\Lambda'}$, $C_{\Lambda'|\Lambda''}$ and $C_{\Lambda''}$ denote the capacities of each layer of the coset decomposition. A key result proved in [29] is that $\mathrm{C}(\mathbf{L})$ can be achieved by any combination of coding rates, provided $R_{\Lambda|\Lambda'} + R_{\Lambda'|\Lambda''} + R_{\Lambda''} = \mathrm{C}(\mathbf{L})$. In particular, apportioning $\mathrm{C}(\mathbf{L})$ by matching coding rate to partition capacity, i.e.,

$$R_{\Lambda|\Lambda'} := \mathrm{C}_{\Lambda|\Lambda'}, \quad R_{\Lambda'|\Lambda''} := \mathrm{C}_{\Lambda'|\Lambda''}, \quad R_{\Lambda''} := \mathrm{C}_{\Lambda''} \quad (12)$$

has an important benefit in terms of reducing complexity. This choice of rate-allocation allows soft-decision multistage decoding to be used without loss of optimality, assuming capacity-achieving component codes are used. This rate allocation strategy is used in this paper, but for simplicity, we use hard-decision decoding.

To compute the rate allocation, consider a nominal transmit power of 0 dBmW. The water-filling capacity analysis in Section II-C revealed a capacity of approximately 49.7 Mbit/s for “channel 3” over a frequency band $0 \le f \le 3.85$ MHz. We implement a zero excess-bandwidth PAM-based scheme, operating at 7.7 MHz, leading to transmission of 6.44 b/symbol. On the other hand, the capacity of a multilevel decomposition of 128-PAM in Gaussian noise is shown in Fig. 5. The figure shows that a capacity of 6.44 b/dim can be achieved at a minimum SNR of about 38.9 dB at the input to a coset code demodulator. For slowly time-varying channels, the rate-allocation must be recomputed periodically with coordination from the receiver.

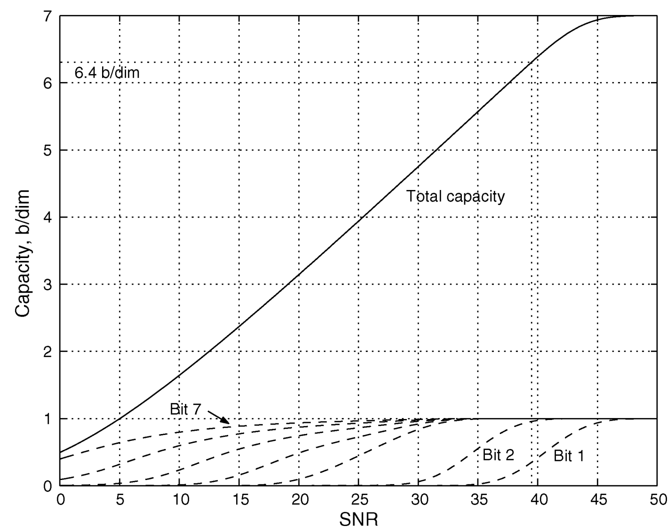

Fig. 5. Capacity analysis of three-level coset partition for 128-PAM, as in [29].

With this approach in mind, the analysis of Fig. 5 shows that rate $R_{\Lambda|\Lambda'} \approx 0.5$ is optimal to code the $\Lambda|\Lambda'$ partition, while $R_{\Lambda'|\Lambda''} \approx 0.9$ is optimal to code $\Lambda'|\Lambda''$. The analysis also shows that the bits mapped onto $\{(\Lambda''+\Omega)\cap\mathcal{S}\}$ can be transmitted at full rate in Gaussian noise. This justifies the choice of a simple three-level coset partition.

Selection of Component Codes: To obtain a steep reduction in BER, we select $\mathbf{G}_{\Lambda|\Lambda'}$ as short (3,6)-regular LDPC codes from [22], which are known to have good performance at low rates. The codes have length 1000–2000 bits which results in low complexity. For $\mathbf{G}_{\Lambda'|\Lambda''}$, we consider two choices: a rate-0.9 regular LDPC code from [9], and a rate-0.9 RS code. It can be shown that a relatively weak algebraic code $\mathbf{G}_{\Lambda'|\Lambda''}$ is sufficient for the $\Lambda'|\Lambda''$ partition, under hard-decision multistage decoding.

Let $P_e(\Lambda|\Lambda')$ denote the bit-error probability on $\Lambda|\Lambda'$, $P_e(\Lambda'|\Lambda'')$ the corresponding probability on $\Lambda'|\Lambda''$, and so on. Then, assuming hard-decision multistage decoding, we can write the error probability of a PAM symbol in Λ as

$$P_e(\Lambda) = P_e(\Lambda|\Lambda') + [1 - P_e(\Lambda|\Lambda')]\,P_e(\Lambda'|\Lambda'') + [1 - P_e(\Lambda|\Lambda')]\,[1 - P_e(\Lambda'|\Lambda'')]\,P_e(\Lambda''). \quad (13)$$

When $P_e(\Lambda|\Lambda') \nrightarrow 0$, $P_e(\Lambda)$ is dominated by errors in $\Lambda|\Lambda'$ and the resulting error-propagation in subsequent layers. When $P_e(\Lambda|\Lambda') \to 0$ at the bottom of the turbo-cliff region, (13) can be approximated by

$$P_e(\Lambda) \approx P_e(\Lambda'|\Lambda'') + [1 - P_e(\Lambda'|\Lambda'')]\,P_e(\Lambda''). \quad (14)$$

For a code on the $\mathbb{Z}|2\mathbb{Z}|4\mathbb{Z}$ lattice partition, $P_e(\Lambda|\Lambda') \to 0$ implies that $\mathbf{G}_{\Lambda'|\Lambda''}$ operates on the correct coset of Λ' in Λ with high probability. However, notice that constellation points on Λ' have 6 dB higher intra-coset separation than points in Λ.

Fig. 6 compares the performance of the rate-0.5 LDPC code $\mathbf{G}_{\Lambda|\Lambda'}$ and a rate -0.9 RS code $\mathbf{G}_{\Lambda'|\Lambda''}$ with 2-PAM modulation.

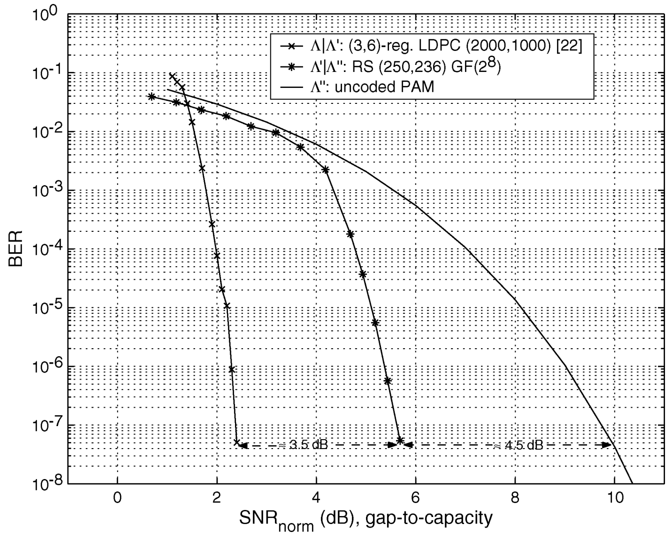

Fig. 6. A comparison between the RS code $\mathbf{G}_{\Lambda'|\Lambda''}$ and LDPC code $\mathbf{G}_{\Lambda|\Lambda'}$.

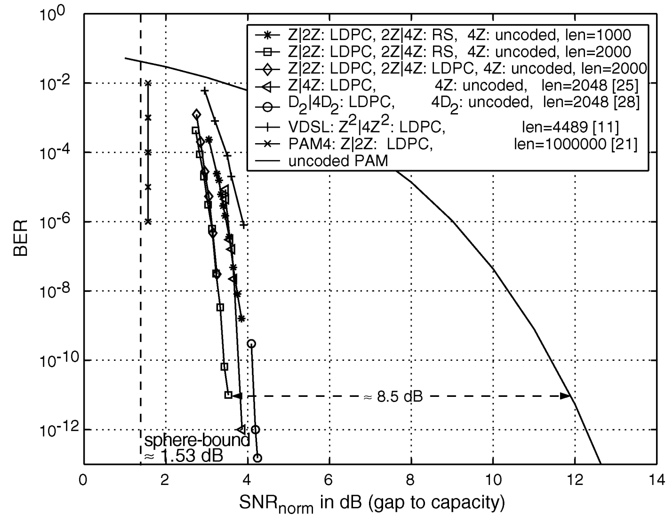

Fig. 7. Comparison of various regular-LDPC coded modulation schemes.

The performance is shown relative to gap-to-capacity [26] or normalized SNR [17], which can be defined as

$$\mathrm{SNR}_{\mathrm{norm}} \triangleq \frac{\mathrm{SNR}}{2^{2R}-1}\,\mathrm{dB} \tag{15}$$

where R is the rate of the code. A capacity-achieving code operates at $\mathrm{SNR}_{\mathrm{norm}} = 0$ dB, while a sphere-bound achieving code operates at $\mathrm{SNR}_{\mathrm{norm}} = 1.53$ dB. Though the RS code is considerably weaker than $\mathbf{G}_{\Lambda|\Lambda'}$, the operating point of $\mathbf{G}_{\Lambda'|\Lambda''}$ is a constant 6 dB ahead of the operating point of $\mathbf{G}_{\Lambda|\Lambda'}$, provided that $\mathbf{G}_{\Lambda|\Lambda'}$ is decoded correctly. This effectively makes $P_e(\Lambda'|\Lambda'') \ll P_e(\Lambda|\Lambda')$, at least within the SNR region shown in Fig. 6. In this range, (14) reduces to $P_e(\Lambda) \approx P_e(\Lambda'')$. Hence, a weak algebraic code like an RS code suffices on the $\Lambda'|\Lambda''$ partition, up to a point.

To estimate the asymptotic performance, we assume that the BER of Fig. 6 reduces at the same rate as shown, as SNR increases. The difference in the BER slopes implies that, at some SNR, say α_0 dB, we will have $P_e(\Lambda|\Lambda') \approx P_e(\Lambda'|\Lambda'')$. This marks the error-floor region of the code, since errors in the region $\mathrm{SNR} > \alpha_0$ are now dominated by $\Lambda'|\Lambda''$ as in (14). To ensure that overall $P_e(\Lambda)$ is sufficiently low asymptotically, $\mathbf{G}_{\Lambda'|\Lambda''}$ must be designed so that α_0 is sufficiently high. For example, the RS (250,236) code shown in Fig. 6 exhibits a factor of BER reduction of 10 per 0.25 dB, while $\mathbf{G}_{\Lambda|\Lambda'}$ shows a decrease in BER of 10 per 0.1 dB. At a $\mathrm{BER} \approx 10^{-7}$, the coding gain between $\mathbf{G}_{\Lambda|\Lambda'}$ and $\mathbf{G}_{\Lambda|\Lambda'}$ is about 3.5 dB. Assuming that the trend continues, the difference increases to 6 dB when $P_e(\Lambda|\Lambda') \approx P_e(\Lambda'|\Lambda'') \approx 10^{-22}$. As explained earlier, this is the "crossover" point, and can be viewed as the error floor region of the code. As SNR increases, $P_e(\Lambda)$ is dominated by errors in $\Lambda'|\Lambda''$. In this example, we see that the choices of $\mathbf{G}_{\Lambda|\Lambda'}$ and $\mathbf{G}_{\Lambda'|\Lambda''}$ are sufficient to keep the error floor low enough for power line communications.

Estimating of the Code Length: As shown in [29], the Gallager random-coding exponent [18] can be used to estimate the block lengths needed to achieve a certain probability of block-error. As pointed out by one of the reviewers, these code lengths must be viewed as lower bounds, as neither LDPC codes nor RS codes have been proved to achieve the random-coding exponent. In particular, we are interested in block-error rates below 10^{-7} to be competitive with performance on digital subscriber lines. The random-coding analysis shows that a code length of 1000–2000 bits at each level of the partition is sufficient, in theory, to achieve the required error rate.

C. Discussion of Simulation Results

Based on the code construction discussed previously, two coset codes are investigated here: one making use of LDPC codes, and the other a combination of LDPC and RS codes. Both schemes use hard-decision multistage decoding. By the capacity analysis of Fig. 5, both schemes use the same (3,6)-regular rate-0.5 regular LDPC code from [22] for $\mathbf{G}_{\Lambda|\Lambda'}$. Furthermore, the top layer is left uncoded in both schemes. However, the schemes differ in the choice of encoder matrix $\mathbf{G}_{\Lambda'|\Lambda''}$: one is a rate-0.94 regular LDPC code from [9], while the other is a binary expansion of an RS code of same rate. Codes around 1000–2000 bits in length are used in all cases. To our knowledge, the family of LDPC codes in [9] exhibit the best performance among regular high-rate codes.

The performance of the schemes is shown in Fig. 7, relative to gap-to-capacity or normalized SNR. The scheme combining LDPC and RS codes lies within 2 dB of the sphere-bound—3.5 dB of Shannon capacity—at a BER of 10^{-11}, measured with simulations over 2.5×10^{12} bits. There is almost no difference in performance to the LDPC-only example proposed above for BERs measured up to 10^{-7}. A second example—with a (3,6)-regular code of 1000 bits for $\mathbf{G}_{\Lambda|\Lambda'}$—shows similar performance to recent proposals for 10 G-Base-T Ethernet [25], [28] which uses (6,32)-regular codes of 2048 bits. However, the complexity of the schemes proposed in this paper is far lower due to the small degree of the nodes in $\mathbf{G}_{\Lambda|\Lambda'}$.

The results can be compared with other LDPC-based coded modulation schemes. In [21], for example, two-level coding schemes over $\mathbb{Z}|2\mathbb{Z}$ with a total rate of 1 b/sym were analyzed. The schemes considered irregular and quasi-regular LDPC codes for $\mathbf{G}_{\Lambda|\Lambda'}$ and $\mathbf{G}_{\Lambda'|\Lambda''}$ with length 10^6 bits each. Here, a gap to the sphere-bound of 0.2–0.3 dB was observed. In [11], the authors investigated LDPC coset-coding over DMT modulation on $\mathbb{Z}^2|4\mathbb{Z}^2$ with QAM constellations for VDSL applications, with component codes on the order of 2000–4000 bits. The scheme is about 1 dB away in performance from proposals in this paper. As mentioned in Section I, the authors of [3] proposed a combination of LDPC and RS codes for coset-coding over DMT on $\mathbb{Z}^2|2\mathbb{Z}^2|4\mathbb{Z}^2$. In the presence of AWGN, the gap to Shannon capacity of the scheme in [3] was found to be about 2.3 dB (0.8–1.2 dB gap to the sphere-bound). The justification for RS codes was to reduce complexity and to ease code selection, which is only part of the rationale of our paper. Meanwhile, to handle different coding rates on each subchannel, a concatenation of LDPC and various repeat-codes are used on the $\Lambda|\Lambda'$ partition.

Soft-Decision and Iterative-Decoding Considerations: As noted by one of the reviewers, it is possible to gain further improvements in the proposed scheme via soft-decision multistage decoding, rather than hard-decisions. Soft-decision decoding also permits iterative equalization, which is useful on time-varying channels that preclude pre-equalization. For the RS code $\mathbf{G}_{\Lambda'|\Lambda''}$, it is well known that a soft-decision binary decoder shows $\approx$1.7 dB of coding gain over a hard-decision binary decoder of the same rate [24]. Analysis in [29] has shown this to be less—about 0.9 dB—for an 8-PAM scheme. In particular, it was shown that there was no significant advantage to soft-decisions beyond the bottom layer $\Lambda|\Lambda'$ in terms of gap to capacity. However, we expect the improvement in soft-decision decoding to improve the error-floor behavior of the code, since a coding gain in $\mathbf{G}_{\Lambda'|\Lambda''}$ only increases α_0, the SNR at which errors in $\Lambda'|\Lambda''$ start to dominate.

Further improvements in performance are possible by feedback of soft extrinsic information from sublattice Λ'' back to $\Lambda|\Lambda'$, and so on iteratively. As also pointed out in [29], feedback from high layers of a decomposition—if properly decorrelated by interleaving—reduces the multiple mappings of bits to symbols on $\Lambda|\Lambda'$. In other words, the decoder on $\Lambda|\Lambda'$ eventually operates on a reduced constellation (2-PAM for a binary lattice decomposition), after sufficient iterations. Similar arguments can be drawn for every layer of the coset decomposition. To take advantage of iterative decoding, a good soft-decision code $\mathbf{G}_{\Lambda'|\Lambda''}$ on $\Lambda'|\Lambda''$ is necessary as such a code can be viewed as "amplifying" extrinsic information, i.e., a code that can accurately compute soft *a posteriori* bit-probabilities, given soft extrinsic information *a priori*. At first glance, it would appear that an LDPC code is better suited for $\mathbf{G}_{\Lambda'|\Lambda''}$ than a soft-decision RS code. However, in asymmetric DSL modems [26], near-capacity performance has been shown with iterative soft-RS decoding concatenated with trellis-coded modulation (TCM). In this example, extrinsic information is iteratively exchanged via a byte-interleaver. Hence, we speculate that the gap to the sphere-bound can be reduced further with soft-decision RS decoding on $\Lambda'|\Lambda''$, combined with interleaving and iterative decoding. However, the improvement is difficult to quantify without further simulations.

IV. Coding in Gaussian and Impulse Noise

In this section, we investigate LDPC-based coset-coding under the simultaneous constraints of colored Gaussian noise and impulse noise. Both synchronous and asynchronous impulse noise models from Section II-B2 are considered, and different coding schemes are compared. Detailed results are presented in Section IV-C. The schemes proposed in this section also incorporate interleaving to mitigate against long bursts of impulse noise without sacrificing coding rate.

A. Error-Correction Schemes Across $\Lambda|\Lambda'|\Lambda''$

- *Coding on* $\Lambda|\Lambda'$. This layer is affected by both Gaussian noise and burst impulse noise. However, the coding scheme here makes use of the same LDPC codes $\mathbf{G}_{\Lambda|\Lambda'}$ used in the AWGN case of Section III-C. The goal is to obtain a steep reduction in BER. The LDPC construction is based on [22], and at low coding rates (only), seems to perform well even in the presence of impulse noise.
- *Coding on* $\Lambda'|\Lambda''$. The capacity analysis of Fig. 5 in Gaussian noise showed that the $\Lambda'|\Lambda''$ partition could be coded at a high rate due to the smaller impact of Gaussian noise at this layer. However, the effect of impulse noise is to cause bursts of errors, either in the form of a single-burst or as multiple phased-bursts. Viewing $\Lambda'|\Lambda''$ as a binary symmetric channel (BSC) for simplicity, we are now interested in a code $\mathbf{G}_{\Lambda'|\Lambda''}$ that can correct both random and burst errors efficiently, with little loss in rate. The burst-correction efficiency, η, of an (n,k) code can be defined [24] by the amount of redundancy required to correct all error bursts of length l-bits or less, viz.,

$$\eta = \frac{2l}{n-k}. \tag{16}$$

A code that can correct all bursts of length l-bits or less with an efficiency of $\eta = 1$ is said to achieve the *Reiger bound*. Though not optimum, we rely on the random and phased-burst error-correcting properties of RS cyclic codes for $\Lambda'|\Lambda''$. A binary expansion of a t-error correcting (n,k) RS code over $GF(2^q)$ can correct the following.
— Any combination of t or fewer random bit errors.
— A single burst of length $l = (t-1)q + 1$ bits or less.
— Any combination of

$$\frac{t}{1 + \lfloor (l+q-2)/q \rfloor} \tag{17}$$

separate bursts of length l, in bits [24].
These properties follow from the fact that a $GF(2^q)$ RS code operates on q-bit symbols. Notice that, as $q \to \infty$, $\eta = 1$. Hence, RS codes are asymptotically optimum. For comparison, LDPC codes are also investigated for $\Lambda'|\Lambda''$.
- *Single-Burst and Phased-Burst Error-Correction of* Λ''. Depending on the sources of impulse noise, the errors of this layer can be either be dominated by long single-burst errors or multiple phased-bursts. There are virtually no errors due to Gaussian noise. Viewing Λ'' as a BSC channel, we investigate RS codes, as well as simple cyclic codes

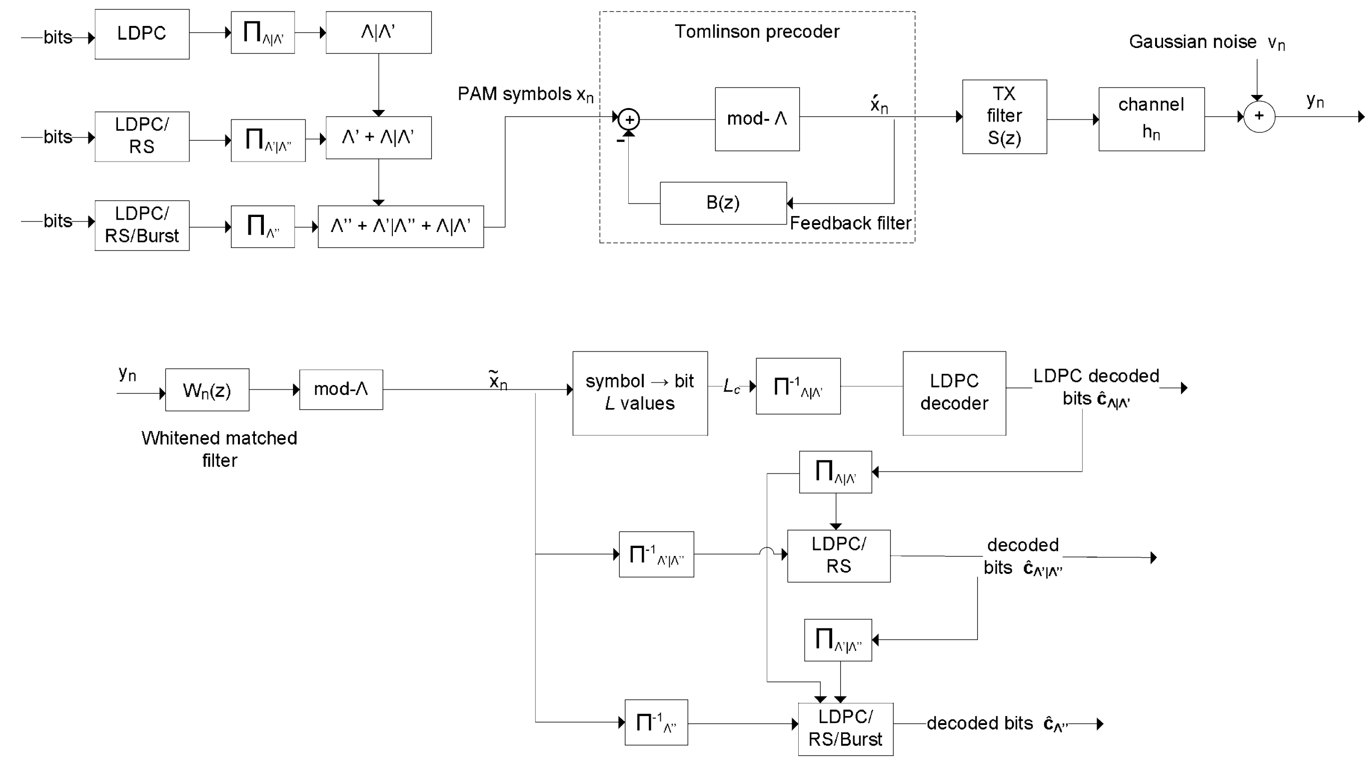

Fig. 8. LDPC coset-coding combined with TH precoding.

over $GF(2)$ optimized for single-burst error-correction [23]. For comparison, LDPC codes are also considered. The main attraction of single-burst error-correcting codes is their low-complexity (lengths on the order of 100–200 bits) coupled with high-efficiency η. In particular, a (195,182) code from [23] is considered, with $\eta \approx 0.77$. To generate codes that fit the coset codeword length m, longer burst-correcting codes can be constructed by interleaving. Simple code-shortening can be used to align the codes on coset codeword boundaries.

B. Multilevel Interleaving

To extend the burst error-correction ability of the coding schemes above, uniform interleavers are used at each layer of a coset decomposition. The proposed scheme is depicted in Fig. 8 and consists of uniform interleavers Π_{Λ}, $\Pi_{\Lambda'}$, and $\Pi_{\Lambda''}$, respectively. In the presence of nonstationary impulsive noise, simulation results (Section IV-C3) will show that interleaved LDPC coding is not sufficient. On the other hand, a coset-coding scheme with a combination of LDPC and cyclic codes is seen to perform well. For the combination of LDPC and cyclic codes, the scheme can be summarized as follows.

1) *Uniform Bit-Interleaving of LDPC Code Bits with* $\Pi_{\Lambda|\Lambda'}$: Let $d_H(G_\Lambda)$ denote the minimum Hamming distance of the LDPC code, and let $\Psi_{\mathbf{L}}$ denote the interleaver depth of coset code $\mathbf{L}$. Then, with optimal decoding, a well-known result from coding theory [24] allows one to correct $\lfloor \Psi_{\mathbf{L}} d_H(G_\Lambda)/2 \rfloor$ burst errors by uniform interleaving with depth $\Psi_{\mathbf{L}}$.
2) $GF(2^q)$ *Symbol-Interleaving of RS Codewords with* $\Pi_{\Lambda'|\Lambda''}$: Now, by GF-symbol interleaving with depth $\Psi_{\mathbf{L}}$, a uniformly interleaved cyclic code over $\mathrm{GF}(2^q)$ can correct any single burst of $\Psi_{\mathbf{L}}((t-1)q+1)$ bit errors. This motivates the use of a q-bit symbol interleaver $\Pi_{\Lambda'|\Lambda''}$ at the $\Lambda'|\Lambda''$ partition.
3) $GF(2^p)$ *Symbol-Interleaving of Cyclic Codewords with* $\Pi_{\Lambda''}$: For a $\mathrm{GF}(2^p)$ cyclic code used on $(\Lambda'' + \Omega) \cap \mathcal{S}$, similar arguments can be drawn to motivate interleaving with $\Pi_{\Lambda''}$ on p-bit boundaries. If the single-burst correcting short codes of [23] are used, a uniform bit-interleaver over $GF(2)$ is sufficient. Since these are already constructed with interleaving to fit the coset codeword length (say, with depth $\Psi_{\Lambda''}$), the effective interleaver depth is $\Psi_{\mathbf{L}}\Psi_{\Lambda''}$.

Due to the multilevel nature of the interleaver, notice from Fig. 8 that the decoded bits from one stage have to be re-interleaved in order to be used as the coset labels of the next.

C. Simulation Results With Colored Noise and Impulse Noise

In this section, we present performance analysis of the proposed coding scheme under the conditions of ISI, colored noise, and impulse noise. In particular, we augment the PAM-based scheme designed in Section III-C to withstand impulse noise. This is difficult to do in an optimal sense due to the nonstationary nature of the impulse noise, which makes it hard to analyze. Our approach is to define a worst case condition, and then design the coding and interleaving scheme accordingly to handle this case. This is clearly a suboptimal approach. Furthermore, even the worst case scenario is a simplified assumption that does not always hold true, as will be explained further. However, to gauge the performance of the scheme, the results—in impulse and Gaussian noise—are compared with channel-capacity with

TABLE I
COMBINATION OF CODING SCHEMES INVESTIGATED UNDER REALISTIC NOISE CONDITIONS

Coset code	$R^b(\mathbf{L})$	$\mathbf{G}_{\Lambda\|\Lambda'}$	$\mathbf{G}_{\Lambda'\|\Lambda''}$	$\mathbf{G}_{\Lambda''}$	$\Psi_{\mathbf{L}}$
$\mathbf{L}_1$	47.7	(3,6)-reg., [22]	(6,32)-reg., rate-0.87 [9]	(7,80)-reg., rate-0.99 [9]	1, 24
$\mathbf{L}_2$	46.3	"	(6,32)-reg., rate-0.86 [9]	(7,80)-reg., rate-0.93 [9]	"
$\mathbf{L}_3$	47.1	"	RS GF (2^8), rate-0.88	RS GF (2^{10}), rate-0.95	"
$\mathbf{L}_4$	46.4	"	RS GF (2^8), rate-0.88	Single-Burst [23], rate-0.93	"

only Gaussian noise. This provides a bound on the gap to the true channel-capacity in the presence of impulse noise.

Table I briefly outlines the combination of coding schemes that were investigated to this end. The channel under consideration is "channel 3," as depicted in Fig. 1. $R^b(\mathbf{L})$ refers to the final transmission rate after additional coding to handle bursts of impulse noise.

1) A "Worst Case" Scenario: Consider the example of Section III-C, which operates at zero-excess bandwidth at a baud-rate of 7.7 MHz. Since we have a binary lattice partition $\mathbb{Z}|2\mathbb{Z}|4\mathbb{Z}$, each coset codeword in $\mathbf{L}$ consists of 2000 128-PAM symbols, as length of $\mathbf{G}_{\Lambda|\Lambda'} = 2000$. This implies a PAM symbol duration of $T_s \approx 0.13\ \mu$s and a frame duration $T_f = 2000 \cdot T_s \approx 260\ \mu$s. Notice that this is much smaller than the periodicity of channel variations observed in [7], which is on the order of 10–20 ms. Consider the European electricity network and assume six synchronous noise impulses per 50 Hz AC cycle. To design for maximum number of burst errors during the cycle period, we assume the impulses are equally spaced in time. Then, an impulse event occurs every 3.3 ms, or approximately every 12.7 coset codewords. We assume that the duration of each synchronous impulse noise burst is $t_{w,s} \approx 50\ \mu$s. Since the peak amplitudes of the synchronous bursts follow a Gaussian distribution $\mathcal{N}(0, G_i\sigma_v^2)$ in (2), we assume that, in the worst case, all symbols exposed to this burst would result in incorrectly decoded bits in the absence of coding. The synchronous burst spans, in the worst case, 385 PAM symbols, as depicted in Fig. 3.

Recall also that the duration of each asynchronous noise burst is assumed to be $t_{w,a} = 100\ \mu$s, which corresponds to 770 PAM symbols. The average interarrival time of the asynchronous bursts is $\tau_{IAT,a} = 100$ ms. Since the asynchronous impulses follow a Poisson arrival process, any number of asynchronous impulse bursts can arrive within a given interval. In particular, the probability of two or more such asynchronous impulses within a codeword interval is given by

$$P\{\geq 2 \text{ Poisson arrivals},\ t = T_f\} = 1 - e^{-\frac{t}{\tau_{IAT,a}}}\left(1 + \frac{t}{\tau_{IAT,a}}\right) \approx 3.4 \times 10^{-6}. \quad (18)$$

Consider an example of an interleaved scheme with $\Psi_{\mathbf{L}} = 24$, and hence, $t = 24T_f$ s. Then, $P\{\geq 2 \text{ Poisson arrivals}\} \approx 2\times 10^{-5}$. When $\Psi_{\mathbf{L}} = 100$, the probability of two asynchronous arrivals is $\approx 3.3 \times 10^{-4}$. These probabilities are small, but of course, not negligible. For now, we assume that two asynchronous impulses will not occur within a $\Psi_{\mathbf{L}}T_f$ time interval; if it does, the error will be detected and corrected by a different means of error control discussed in Section IV-C3.

Hence, the "worst case" scenario can be stated as follows: one asynchronous impulse and a commensurate number of synchronous impulses occur within an interval $\Psi_{\mathbf{L}}T_f$ s. The burst-lengths $T_{w,s}$ and $T_{w,a}$ of the impulses are spanned completely within $\Psi_{\mathbf{L}}T_f$. Let $l^b_{\Lambda|\Lambda'}$, $l^b_{\Lambda'|\Lambda''}$, and $l^b_{\Lambda''}$ denote the total phased-burst lengths, in bits, on $\Lambda|\Lambda$, $\Lambda'|\Lambda''$, and Λ'', respectively. Let l^b_Λ denote the burst length in terms of PAM symbols in Λ. Then

$$l^b_\Lambda = \underbrace{\left(\frac{\Psi_{\mathbf{L}}T_f}{t_{\text{arr},s}} \cdot \frac{T_{w,s}}{T_s}\right)}_{\#\text{ of synchronous impulses} \times \text{length}} + \underbrace{\frac{T_{w,a}}{T_s}}_{\text{asynchronous length}} \quad \text{symbols} \quad (19)$$

$$l^b_{\Lambda|\Lambda'} = \log_2|\Lambda|\Lambda'| \times l^b_\Lambda = l^b_\Lambda \quad \text{bits} \quad (20)$$

$$l^b_{\Lambda'|\Lambda''} = \log_2|\Lambda'|\Lambda''| \times l^b_\Lambda = l^b_\Lambda \quad \text{bits} \quad (21)$$

$$l^b_{\Lambda''} = \log_2|(\Lambda'' + \Omega) \cap \mathcal{S}| \times l^b_\Lambda = 5l^b_\Lambda \quad \text{bits}. \quad (22)$$

These burst-lengths are used in subsequent sections. We also use $\hat{l}^b_{\Lambda|\Lambda'}$, $\hat{l}^b_{\Lambda'|\Lambda''}$, and $\hat{l}^b_{\Lambda''}$ to denote the *correctable* total phased-burst lengths, in bits, on $\Lambda|\Lambda$, $\Lambda'|\Lambda''$, and Λ'', respectively.

2) Tradeoff Between Interleaver Depth and Burst-Error Coding: Our goal is to find a good combination of interleaving and burst-error correction needed to correct all burst-errors in the worst case scenario. To correct one asynchronous burst and a commensurate number of synchronous bursts within a period $\Psi_{\mathbf{L}}T_f$, we first estimate the rates by which $\mathbf{G}_{\Lambda|\Lambda'}$, $\mathbf{G}_{\Lambda'|\Lambda''}$, and $\mathbf{G}_{\Lambda''}$ have their coding rates lowered to correct only the burst-errors. Computing this rate reduction can be very complex, since some codes can correct both random and burst errors simultaneously, and some errors due to Gaussian noise may overlap with burst-errors. To simplify this analysis and to obtain just an estimate of the reduction in rate needed, we assume that all the code redundancy created by rate reduction is available for *optimal* phased-burst-error correction. Then, the coding rates $R_{\Lambda|\Lambda'}$, $R_{\Lambda'|\Lambda''}$, and $R_{\Lambda''}$ are reduced by factors $\rho_{\Lambda|\Lambda'}$, $\rho_{\Lambda'|\Lambda''}$, and $\rho_{\Lambda''}$, respectively, where

$$\rho_{\Lambda|\Lambda'} = 2\frac{l^b_{\Lambda|\Lambda'}}{m\Psi_{\mathbf{L}}} \text{ b/sym}, \quad \rho_{\Lambda'|\Lambda''} = 2\frac{l^b_{\Lambda'|\Lambda''}}{m\Psi_{\mathbf{L}}} \text{ b/sym},$$
$$\rho_{\Lambda''} = 2\frac{l^b_{\Lambda''}}{5m\Psi_{\mathbf{L}}} \text{ b/sym}. \quad (23)$$

TABLE II
COMPONENT CODES USED IN THE SIMULATIONS, ALONG WITH ESTIMATES OF WORST CASE BURST-LENGTHS l^b, AND CORRECTABLE LENGTHS $\hat{l}^b$. ρ DENOTES THE RATE-REDUCTION CHOSEN TO SUPPORT ADDITIONAL BURST-ERROR CORRECTION

Component Code	b/sym	$\Psi_{\mathbf{L}}=1$ bits	bits	$\Psi_{\mathbf{L}}=24$ bits	bits
$\mathbf{G}_{\Lambda\vert\Lambda'}$	$\rho_{\Lambda\vert\Lambda'}$	$l^b_{\Lambda\vert\Lambda'}$	$\hat{l}^b_{\Lambda\vert\Lambda'}$	$l^b_{\Lambda\vert\Lambda'}$	$\hat{l}^b_{\Lambda\vert\Lambda'}$
(2000,1000), rate-0.5 [22]: $\mathbf{L}_{1,\ldots,4}$	0.000	1.1e+03	0.0e+00	1.5e+003	0.0e+00
$\mathbf{G}_{\Lambda'\vert\Lambda''}$	$\rho_{\Lambda'\vert\Lambda''}$	$l^b_{\Lambda'\vert\Lambda''}$	$\hat{l}^b_{\Lambda'\vert\Lambda''}$	$l^b_{\Lambda'\vert\Lambda''}$	$\hat{l}^b_{\Lambda'\vert\Lambda''}$
(6,32)-reg, rate-0.87 [9] : $\mathbf{L}_1$	0.065	1.1e+03	7.7e+02	1.5e+03	1.5e+03
(6,32)-reg, rate-0.86 [9] : $\mathbf{L}_2$	0.078	1.1e+03	9.4e+02	1.5e+03	1.9e+03
(250, 214) RS - GF(2^8) : $\mathbf{L}_3, \mathbf{L}_4$	0.064	1.1e+03	7.7e+02	1.5e+03	1.5e+03
$\mathbf{G}_{\Lambda''}$	$\rho_{\Lambda''}$	$l^b_{\Lambda''}$	$\hat{l}^b_{\Lambda''}$	$l^b_{\Lambda''}$	$\hat{l}^b_{\Lambda''}$
(7,80)-reg, rate-0.99 [9] : $\mathbf{L}_1$	0.014	5.7e+03	8.6e+02	7.5e+03	1.7e+03
(7,80)-reg, rate-0.93 [9] : $\mathbf{L}_2$	0.069	5.7e+03	4.2e+03	7.5e+03	8.3e+03
(1000, 948) RS – $GF(2^{10})$: $\mathbf{L}_3$	0.052	5.7e+03	3.1e+03	7.5e+03	6.2e+03
(195, 182) Single-Burst [23] : $\mathbf{L}_4$	0.067	5.7e+03	4.0e+03	7.5e+03	8.0e+03

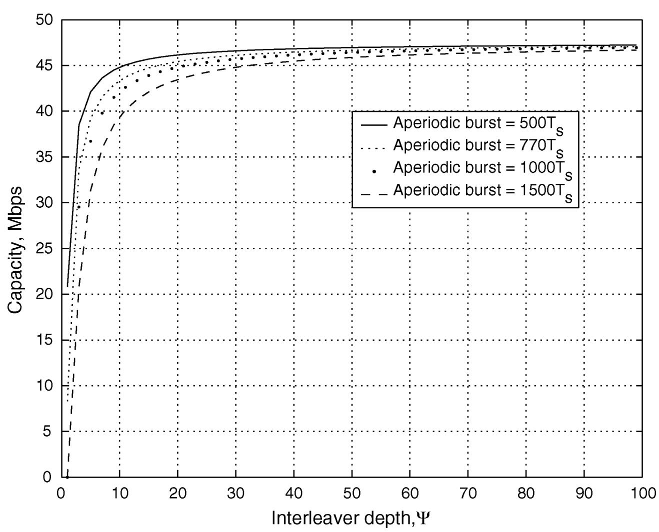

Fig. 9. Estimation of transmission rates on "channel 3" in the presence of burst noise, assuming a combination of interleaving and coding $P_T = 0$ dBmW.

The final rate is therefore

$$R^b(\mathbf{L}) = \frac{1}{T_s}\left[R_{\Lambda|\Lambda'}(1-\rho_{\Lambda|\Lambda'}) + R_{\Lambda'|\Lambda''}(1-\rho_{\Lambda'|\Lambda''}) + R_{\Lambda''}(1-\rho_{\Lambda''})\right] \text{ bit/sec} \quad (24)$$

which is sketched in Fig. 9 for various values of $\Psi_{\mathbf{L}}$ and $T_{w,a}$. Naturally, it can be seen that the performance increases with interleaver depth $\Psi_{\mathbf{L}}$. However, the full Gaussian capacity cannot be reached, irrespective of $\Psi_{\mathbf{L}}$, since

$$\lim_{\Psi_{\mathbf{L}}\to\infty} \rho_{\Lambda|\Lambda'} = \lim_{\Psi_{\mathbf{L}}\to\infty} \rho_{\Lambda'|\Lambda''} = \lim_{\Psi_{\mathbf{L}}\to\infty} \rho_{\Lambda''} = \frac{1}{m}\frac{T_f}{t_{\text{arr},s}}\frac{T_{w,s}}{T_s} \quad \text{b/sym.} \quad (25)$$

The limit is determined by the fraction of PAM symbols affected by synchronous noise, which is independent of $\Psi_{\mathbf{L}}$.

Suppose we are provided with a 2 dBmW budget in transmit power, or corresponding SNR, for protection against burst noise. From the optimal water-filling capacity graph of Fig. 4, this corresponds to 2.6 Mbit/s decrease in information rate to be used for impulse-noise protection. From Fig. 9, we estimate that $\Psi_{\mathbf{L}} = 24$ is a good choice for error-free transmission at $R^b(\mathbf{L}) \approx 49.7 - 2.6$ Mbit/s. With these rate and interleaver-depth estimates, our goal is to design a practical scheme with LDPC, RS, or burst-error correcting codes. Our design choice is to leave the LDPC code $\mathbf{G}_{\Lambda|\Lambda'}$ unchanged since the code has been observed to perform well in the presence of impulse noise, even without reducing the rate. Table II shows the various component codes used in the simulations, along with their reduction in rate $\rho_{\Lambda|\Lambda'}$, $\rho_{\Lambda'|\Lambda''}$, and $\rho_{\Lambda''}$, as well as the worst case burst-lengths experienced at the corresponding layers. To our knowledge, the LDPC codes selected from [9] are the best high-rate regular LDPC codes in AWGN channels.

3) Discussion of Simulation Results: The coset codes of Table I were tested under the channel and noise conditions described in earlier sections. The details of the component codes are listed in Table II. All the coset codes, $\mathbf{L}_1, \ldots, \mathbf{L}_4$ use the same regular LDPC code from [22] for $\mathbf{G}_{\Lambda|\Lambda'}$. The comparative performance of the schemes with $\Psi_{\mathbf{L}} = 1$ is shown in Fig. 10. The performance in all the figures is shown in terms of gap-to-capacity in Gaussian noise (since capacity in impulse-noise is difficult to compute). It can be seen that the impulse noise, coupled with lack of interleaving, has a devastating impact on BER performance for all schemes, irrespective of code type. This is can be seen from Table II, where it is not possible to have $n - k \geq 2l^b$ for any of the coding rates selected when $\Psi_{\mathbf{L}} = 1$; this makes it impossible to correct the noise bursts experienced across the coset decomposition.

With interleaver depth of $\Psi_{\mathbf{L}} = 24$, the BER performance is also shown in Fig. 10. It can be seen that code $\mathbf{L}_3$ exhibits best performance, with no apparent error floor. The gap to capacity for rate $R^b(\mathbf{L}_3)$ is about 5.5 dB at a BER of 10^{-7}. The code $\mathbf{L}_4$ also shows good performance until a BER of 3×10^{-6}, at

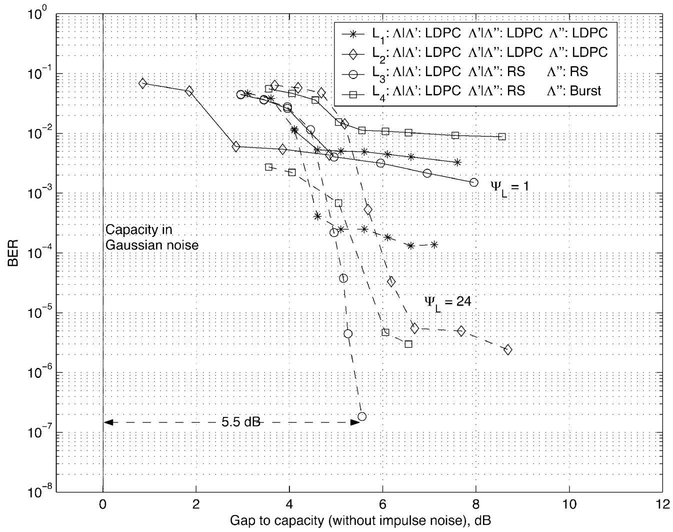

Fig. 10. Performance of coding schemes on "channel 3," in the presence of colored noise, residual ISI, synchronous and asynchronous impulse noise. Interleaver depth $\Psi = 1$ and $\Psi = 24$.

which point an error floor is seen. The advantage of phased-burst error protection afforded by RS codes is evident. However, the complexity of a 1000-symbol RS code over $\mathrm{GF}(2^{10})$ is vastly higher than an interleaved (195,182) binary cyclic code. The LDPC-only scheme $\mathbf{L}_2$ performs well until a $\mathrm{BER} \approx 5 \times 10^{-6}$, but then exhibits an error floor. $\mathbf{L}_1$ still exhibits poor performance. This suggests that low-rate LDPC codes are more capable of handling burst-errors.

Handling Scenarios Worse Than the "Worst Case": Although the simulations results show good performance with code $\mathbf{L}_3$ and $\Psi_{\mathbf{L}} = 24$, we have not adequately characterized the system at BERs below 10^{-7} due to simulation complexity. In particular, conditions worse than the "worst case" scenario of Section IV-C1 occur with a probability $P_b \triangleq P\{\geq 2 \text{ asynchronous impulse arrivals}, t = 24T_f\} \approx 2 \times 10^{-5}$; i.e., on average, the worst case is exceeded every 1.2×10^6 coset codewords. Since these are uncorrectable errors with high probability, we would expect to see an error floor around $\mathrm{BER} \approx 10^{-7}$, which can be inferred from block-error rates already measured.

To solve this anticipated error-floor, we briefly summarize a well-known solution in analytical form; due to the length of the simulations, we could not provide supporting results. An effective solution is to combine the proposed scheme with an automatic repeat-request (ARQ) protocol. Such schemes maintain coding rate—without sacrificing noise immunity—by using forward error-correction (FEC) until, say, probability of block-error $< 10^{-3}$. In the event of an uncorrectable error, a retransmit of the interleaved set of codewords is performed. There exists several variations on this theme, cf. type-I hybrid-ARQ protocols [24]. A well-known technique, viz., type-I with selective-ARQ, yields an effective error-free transmission rate of

$$R^b_{type-I}(\mathbf{L}_3) = R^b(\mathbf{L}_3)\,[1 - P_b.(1 - P_u)] \quad \text{bit/sec} \qquad (26)$$

where $R^b(\mathbf{L}_3)$ is the rate of code $\mathbf{L}_3$, and P_u is the probability of an undetected error for a codeword in $\mathbf{L}_3$. P_u can be made very small with cyclic-redundancy codes, with little loss in rate [24]. Since $P_b \approx 10^{-5}$, it is easy to achieve error-free transmission at nearly the full rate $R^b(\mathbf{L}_3)$.

V. Conclusion

A simple LDPC-based coset-coding scheme for power line channels was investigated. The scheme combines LDPC and cyclic codes to achieve near-capacity performance in Gaussian noise, and to correct burst-errors in impulse noise. At a BER of 10^{-11}, the gap to unshaped channel capacity is about 2 dB in Gaussian noise (corresponding to a coding gain of 8.5 dB over uncoded PAM). The component codes are based on simple regular LDPC codes of small length. To mitigate impulse noise, RS and burst-error-correcting cyclic codes were investigated. An interleaving scheme is also proposed, consisting of distinct interleaving stages tailored to each level of the coset code. This results in increased immunity to burst noise caused by impulses. To mitigate ISI, the coding scheme is investigated on slowly varying power line channels with TH-precoding. In the presence of colored Gaussian noise, synchronous and asynchronous impulse noise and residual ISI, the gap to channel-capacity of a Gaussian-noise-only channel is about 5.5 dB at a $\mathrm{BER} \approx 10^{-7}$.

Acknowledgment

The authors thank the editor, Dr. Galli, and the anonymous reviewers for their meticulous feedback, which greatly helped to improve the quality of this paper.

References

[1] ITU, Single-pair high speed digital subscriber line (SHDSL) transceivers, ITU Recommendation G.991.2, Feb. 2001.

[2] J. Abad, A. Badenes, J. Blasco, J. Carreras, V. Dominguez, C. Gomez, S. Iranzo, J. C. Riveiro, D. Ruiz, L. M. Torres, and J. Comabella, "Extending the power line LAN up to the neighborhood transformer," *IEEE Commun. Mag.*, vol. 41, no. 4, pp. 64–70, Apr. 2003.

[3] M. Ardakani, T. Esmailian, and F. R. Kschischang, "Near-capacity coding in multi-carrier modulation systems," *IEEE Trans. Commun.*, vol. 52, no. 11, pp. 1880–1889, 2004.

[4] S. Baig and N. D. Gohar, "A discrete multitone transceiver at the heart of the PHY layer of an in-home power line communication local area network," *IEEE Commun. Mag.*, vol. 41, no. 4, pp. 48–70, Apr. 2003.

[5] T. C. Banwell and S. Galli, "On the symmetry of the power line channel," in *Proc. Int. Symp. Power Lines Commun.*, 2001, pp. 325–330.

[6] J. H. Conway and N. J. A. Sloane, *Sphere Packings, Lattices, and Groups*. New York: Springer-Verlag, 1999.

[7] J. A. Cortes, F. J. Cañete, L. Diez, and J. T. Entrambasaguas, "Characterization of the cyclic short-time variation of indoor power-line channels," in *Proc. Int. Symp. Power Line Commun. Appl.*, Apr. 2005, pp. 326–330.

[8] V. Degardin, M. Lienard, A. Zeddam, F. Gauthier, and P. Degauque, "Classification and characterization of impulsive noise on indoor power lines used for data communications," *IEEE Trans. Consum. Electron.*, vol. 48, no. 4, pp. 913–918, Nov. 2002.

[9] I. Djurdjevic, J. Xu, A.-G. Khaled, and S. Lin, "A class of low-density parity-check codes constructed based on Reed–Solomon codes with two information symbols," *IEEE Commun. Lett.*, vol. 7, no. 7, pp. 317–319, Jul. 2003.

[10] K. M. Dostert, "Frequency-hopping spread-spectrum modulation for digital communications over electrical power lines," *IEEE J. Sel. Areas Commun.*, vol. 8, no. 4, pp. 700–710, May 1990.

[11] E. Eleftheriou and S. Ölçcer, "Low-density parity-check codes for digital subscriber lines," in *Proc. Int. Conf. Commun.*, Apr. 2002, vol. 3, no. 28, pp. 1752–1757.

[12] R. F. H. Fischer, *Precoding and Signal Shaping for Digital Transmission*. New York: Wiley, 2002.
[13] G. D. Forney, "Coset codes—Part I: Introduction and geometrical classification," *IEEE Trans. Inf. Theory*, vol. 34, no. 5, pp. 1123–1151, Sep. 1988.
[14] ——, "Coset codes—Part II: Binary lattices and related codes," *IEEE Trans. Inf. Theory*, vol. 34, no. 5, pp. 1152–1187, Sep. 1988.
[15] G. D. Forney, M. D. Trott, and S.-Y. Chung, "Sphere-bound-achieving coset codes and multilevel coset codes," *IEEE Trans. Inf. Theory*, vol. 46, no. 3, pp. 820–850, May 2000.
[16] G. D. Forney and M. V. Eyubŭglu, "Combined equalization and coding using precoding," *IEEE Commun. Mag.*, vol. 29, no. 12, pp. 25–34, Dec. 1991.
[17] G. D. Forney and G. Ungerboeck, "Modulation and coding for linear Gaussian channels," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2384–2415, Oct. 1998.
[18] R. G. Gallager, *Information Theory and Reliable Communication*. New York: Wiley, 1968.
[19] R. D. Gitlin, J. F. Hayes, and S. B. Weinstein, *Data Communications Principles*. New York: Plenum Press, 1992, 10 013.
[20] M. Gotz, M. Rapp, and K. Dostert, "Power line channel characteristics and their effect on communication system design," *IEEE Commun. Mag.*, vol. 41, no. 4, pp. 78–86, Apr. 2004.
[21] J. Hou, P. H. Siegel, L. B. Milstein, and H. D. Pfister, "Capacity-approaching bandwidth-efficient coded modulation schemes based on low density parity-check codes," *IEEE Trans. Inf. Theory*, vol. 49, no. 9, pp. 2141–2155, Sep. 2003.
[22] X.-Y. Hu, E. Eleftheriou, and D. M. Arnold, "Regular and irregular progressive edge-growth Tanner graphs," *IEEE Trans. Inf. Theory*, vol. 51, no. 1, pp. 386–398, Jan. 2005.
[23] T. Kasami, "Optimum shortened cyclic codes for burst-error correction," *IEEE Trans. Inf. Theory.*, vol. 9, no. 2, pp. 105–109, Apr. 1963.
[24] S. Lin and D. J. Costello, *Error Control Coding: Fundamentals and Applications*. Upper Saddle River, NJ: Prentice-Hall, 2004.
[25] S. Rao, R. Hormis, and E. Krouk, The 4D PAM-8 proposal for 10 G-Base-T, IEEE 802.3an 10 G-Base-T Task Force, Nov. 2003. [Online]. Available: http://grouper.ieee.org/groups/802/3/tutorial/ nov03/rao_1_11_03.pdf
[26] T. Starr, J. M. Cioffi, and P. J. Silverman, *Understanding Digital Subscriber Line Technology*. Englewood Cliffs, NJ: Prentice-Hall, 1999.
[27] D. Toumpakaris, J. M. Cioffi, D. Gardan, and M. Ouzzif, "The capacity of typical powerline reference channels and strategies for system design," in *Proc. 5th Int. Symp. Power Line Commun.*, Malmo, Sweden, 2001.
[28] G. Ungerboeck, 10 G-base-T modulation and coding, set of fixed precoders, and start-up IEEE 802.3an 10 G-Base-T Task Force, Nov. 2004. [Online]. Available: http://www.ieee802.org/3/an/public/nov04/ungerboeck_1_1104.pdf
[29] U. Wachsmann, R. F. H. Fischer, and J. B. Huber, "Multilevel codes: Theoretical concepts and practical design rules," *IEEE Trans. Inf. Theory*, vol. 45, no. 5, pp. 1361–1391, Jul. 1999.
[30] X. Wang and H. V. Poor, "Iterative (turbo) soft interference cancellation and decoding for coded CDMA," *IEEE Trans. Commun.*, vol. 47, no. 7, pp. 1046–1061, Jul. 1999.
[31] M. Zimmermann and K. Dostert, "An analysis of the broadband noise scenario in powerline networks," in *Proc. Int. Symp. Power Line Commun.*, Lancaster, U.K., 1999.
[32] ——, "Analysis and modeling of impulsive noise in broadband powerline communications," *IEEE Trans. Electromag. Compat.*, vol. 44, no. 1, pp. 249–258, Feb. 2002.
[33] ——, "A multipath model for the powerline channel," *IEEE Trans. Commun.*, vol. 50, no. 4, pp. 553–559, Apr. 2002.

Raju Hormis is working towards the Ph.D. degree in electrical engineering at Columbia University, New York.

His interests are in wireline communications, error-control coding, video coding, and video signal conditioning.

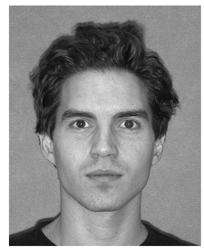

Inaki Berenguer received the M.S. and Ph.D. degrees in engineering from Cambridge University, Cambridge, U.K., in 2002 and 2004, respectively, and was a two-year graduate Fulbright Scholar at Columbia University, New York.

From July 1999 to January 2001, he was with Hewlett Packard, Brussels, Belgium. In January 2001, he joined the Wireless Division of STMicroelectronics, San Diego, CA, where he worked on a W-CDMA project until August 2001. In Summer 2004, he worked for Intel Corporation and from September to December 2004, and in Summer 2005, he was a Research Associate at NEC Laboratories America, Princeton, NJ. Since December 2005, he has been with McKinsey & Company.

Dr. Berengue is the recipient of the British Council, La Caixa, Cambridge-MIT Institute, and Leonardo Da Vinci Fellowships.

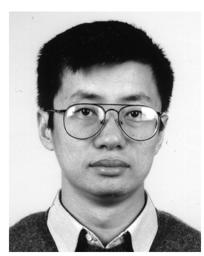

Xiaodong Wang received the B.S. degree in electrical engineering and applied mathematics (with the Highest Honor) from Shanghai Jiao Tong University, Shanghai, China, in 1992, the M.S. degree in electrical and computer engineering from Purdue University, West Lafayette, IN, in 1995, and the Ph.D. degree in electrical engineering from Princeton University, Princeton, NJ, in 1998.

From July 1998 to December 2001, he was on the faculty of the Department of Electrical Engineering, Texas A&M University. In January 2002, he joined the faculty of the Department of Electrical Engineering, Columbia University. Among his publications is a recent book *Wireless Communication Systems: Advanced Techniques for Signal Reception*. His current interests include wireless communications, statistical signal processing, and genomic signal processing. His research interests fall in the general areas of computing, signal processing, and communications. He has published extensively in these areas.

Dr. Wang received the 1999 NSF CAREER Award, and the 2001 IEEE Communications Society and Information Theory Society Joint Paper Award. He currently serves as an Associate Editor for the IEEE TRANSACTIONS ON COMMUNICATIONS, the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS, the IEEE TRANSACTIONS ON SIGNAL PROCESSING, and the IEEE TRANSACTIONS ON INFORMATION THEORY.