

A Message-Passing Receiver for BICM-OFDM Over Unknown Clustered-Sparse Channels

Philip Schniter

Abstract—We propose a factor-graph-based approach to joint channel-estimation-and-decoding (JCED) of bit-interleaved coded orthogonal frequency division multiplexing (BICM-OFDM). In contrast to existing designs, ours is capable of exploiting not only sparsity in sampled channel taps but also clustering among the large taps, behaviors which are known to manifest at larger communication bandwidths. In order to exploit these channel-tap structures, we adopt a two-state Gaussian mixture prior in conjunction with a Markov model on the hidden state. For loopy belief propagation, we exploit a “generalized approximate message passing” (GAMP) algorithm recently developed in the context of compressed sensing, and show that it can be successfully coupled with soft-input soft-output decoding, as well as hidden Markov inference, through the standard sum-product framework. For N subcarriers and any channel length $L < N$, the resulting JCED-GAMP scheme has a computational complexity of only $\mathcal{O}(N \log_2 N + N|\mathcal{S}|)$, where $|\mathcal{S}|$ is the constellation size. Numerical experiments using IEEE 802.15.4a channels show that our scheme yields BER performance within 1 dB of the known-channel bound and 3–4 dB better than soft equalization based on LMMSE and LASSO.

Index Terms—Belief propagation, blind equalizers, channel estimation, decoding, message passing, orthogonal frequency-division multiplexing (OFDM), ultra-wideband communication.

I. INTRODUCTION

WHEN designing a digital communications receiver, it is common to model the effects of multipath propagation in discrete time using a convolutive linear channel that, in the slow-fading scenario, can be characterized by a fixed impulse response $\{x_j\}_{j=0}^{L-1}$ over the duration of one codeword. When the communication bandwidth is sufficiently low, the “taps” $\{x_j\}_{j=0}^{L-1}$ are well modeled as independent complex Gaussian random variables, resulting in the “uncorrelated Rayleigh-fading” and “uncorrelated Rician-fading” models that have dominated the wireless communications literature for many decades [1]. For receiver design, the Gaussian tap assumption is very convenient because the optimal estimation

scheme is well known to be linear [2]. As the communication bandwidth increases, however, the channel taps are no longer well-modeled as Gaussian nor independent. Rather, they tend to be heavy-tailed or “sparse” in that only a few values in $\{x_j\}_{j=0}^{L-1}$ have significant amplitude [3]–[6]. Moreover, groups of large taps are often clustered together in lag j . These behaviors are both a blessing and a curse: a blessing because, of all tap distributions, the independent Gaussian one is most detrimental to capacity [7], but a curse because optimal channel estimation becomes nonlinear and thus receiver design becomes more complicated.

Recently, there have been many attempts to apply breakthrough nonlinear estimation techniques from the field of “compressive sensing” [8] (e.g., LASSO [9], [10]) to the wireless channel estimation problem. We refer to this approach as “compressed channel sensing” (CCS), after the recent comprehensive overview [11]. The CCS literature generally takes a *decoupled* approach to the problem of channel estimation and data decoding, in that pilot-symbol knowledge is first exploited for sparse-channel estimation, after which the channel estimate is used for data decoding. However, this decoupled approach is known to be suboptimal [12].

The considerations above motivate a *joint* approach to structured-sparse-channel-estimation and decoding that offers both near-optimal decoding performance and low implementation complexity. In this paper, we propose exactly such a scheme. In particular, we focus on orthogonal frequency-division multiplexing (OFDM) with bit-interleaved coded modulation (BICM), and propose a novel factor-graph-based receiver that leverages recent results in “generalized approximate message passing” (GAMP) [13], soft-input/soft-output (SISO) decoding [14], and structured-sparse estimation [15]. Our receiver assumes a clustered-sparse channel-tap prior constructed using a two-state Gaussian mixture with a Markov model on the hidden tap state. The scheme that we propose has only $\mathcal{O}(N \log_2 N + N|\mathcal{S}|)$ complexity, where N denotes the number of subcarriers and $|\mathcal{S}|$ denotes the constellation size, facilitating large values of N and channel length $L < N$ (e.g., we use $N = 1024$ and $L = 256$ for our numerical results). For rich non-line-of-sight (NLOS) channels generated according to the IEEE 802.15.4a standard [16], our numerical experiments show bit error rate (BER) performance within 1 dB of the known-channel bound and 3–4 dB better than soft equalization based on LMMSE and LASSO.

We now place our work in the context of existing factor-graph designs. Factor-graph based joint channel-estimation and decoding (JCED) was proposed more than a decade ago (see, e.g., the early overview [17]). To calculate the messages passed

Manuscript received January 25, 2011; revised May 20, 2011, August 29, 2011; accepted September 07, 2011. Date of publication September 23, 2011; date of current version November 18, 2011. This work was supported in part by the National Science Foundation under Grant CCF-1018368, in part by the Defense Advanced Research Projects Agency/Office of Naval Research under Grant N66001-10-1-4090, and an allocation of computing time from the Ohio Supercomputer Center. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Andrew Singer.

The author is with the Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH 43210 USA (e-mail: schniter@ece.osu.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JSTSP.2011.2169232

among the nodes of the factor graph, first instincts suggest to apply the standard “sum-product algorithm” (SPA) [18]–[20]. Exact SPA on the JCED factor graph is computationally infeasible, however, and so it must be approximated. For this, there are many options, since many well-known iterative inference algorithms can themselves be recognized as SPA approximations, e.g., the expectation–maximization (EM) algorithm [21], particle filtering [22], variational (or “mean-field”) techniques [23], and even steepest descent [24]. Moreover, because the JCED factor graph is loopy, even non-approximate SPA is not guaranteed to yield the correct output distributions, because exact inference is NP hard [25]. It is perhaps not surprising that, amidst this uncertainty about exact SPA and its “best” approximation, a number of different factor-graph approaches to JCED over frequency-selective channels have been proposed (e.g., [26]–[29]).

Our approach differs from existing factor-graph JCED designs in that it uses 1) a sparse (i.e., non-Gaussian) channel-tap prior, 2) a clustered (i.e., non-independent) channel-tap prior, and 3) a state-of-the-art SPA approximation known as “generalized approximate message passing” (GAMP), which has been shown to admit rigorous analysis as $N, L \rightarrow \infty$ [13]. In fact, we conjecture that the success of our method is due in large part to the principled approximations used within GAMP. We also note that, although we focus on the case of clustered-sparse channels, our approach could be applied to non-sparse (i.e., Gaussian) or non-clustered (i.e., independent) channel-taps or, e.g., non-sparse channels with unknown length L [26], with minor modifications of our assumed channel prior.

Finally, we mention that this work is an evolution of our earlier work [30], [31] that was limited to an exactly sparse channel, that did not exploit clustering, and that was based on the “relaxed belief propagation” (RBP) algorithm [32], which has higher implementation complexity than GAMP. For example, the JCED scheme from [30], [31] has complexity $\mathcal{O}(NL + N|\mathcal{S}|)$, which grows with the channel length L .

Our paper is organized as follows. In Section II, we detail our assumptions on the OFDM system and the channel prior, and provide an illustrative example of clustered-sparse behavior with the IEEE 802.15.4a channel model. In Section III, we detail our GAMP-based JCED approach, in Section IV we report the results of our simulation study, and in Section V we conclude.

Throughout the paper, we use the following notation. \mathbb{R} denotes the field of reals and \mathbb{C} the complex field. $(\cdot)^*$ denotes conjugate and $\text{Re}(\cdot)$ extracts the real part. Furthermore, $\delta(\tau)$ denotes the Dirac delta waveform while $\{\delta_n\}_{n=-\infty}^{\infty}$ denotes the Kronecker delta sequence. Also, $\langle j \rangle_N$ denotes j -modulo- N , \star convolution, and \propto denotes equality up to a scaling. We use boldface capital letters like \mathbf{B} to denote matrices and boldface small letters like \mathbf{b} to denote vectors. \mathbf{I} denotes the identity matrix, $\mathbf{1}$ denotes the vector of ones, and $\mathcal{D}(\mathbf{b})$ constructs a diagonal matrix from the vector \mathbf{b} . For matrices and vectors, $(\cdot)^T$ denotes transpose and $(\cdot)^H$ denotes conjugate transpose. When x_j is a realization of random variable X_j , we write $x_j \sim X_j$ and use $\mathbb{E}_{X_j}\{x_j\}$ to denote the mean, $\text{var}_{X_j}\{x_j\}$ the variance, $p_{X_j}(x_j)$ the pdf, and $p_{X_j|D_j}(x_j|d_j)$ the pdf conditioned on the event $D_j = d_j$. Sometimes we omit the subscript when there is no danger of confusion, yielding, e.g., $\mathbb{E}\{x_j\}$, $\text{var}\{x_j\}$, $p(x_j)$ and $p(x_j|d_j)$. $\mathcal{CN}(x; \hat{x}, \nu^x) \triangleq (\pi\nu^x)^{-1} \exp(-|x - \hat{x}|^2/\nu^x)$

denotes the circular Gaussian pdf with mean \hat{x} and variance ν^x . In fact, we often use (\hat{v}_j, ν_j^v) when referring to the mean and variance of V_j . For a random vector \mathbf{x} , we use $\text{Cov}(\mathbf{x})$ to denote the covariance matrix.

II. SYSTEM MODEL

A. BICM-OFDM Model

We consider an OFDM system with N subcarriers, each modulated by a quadrature amplitude modulation (QAM) symbol from a 2^M -ary unit-energy constellation \mathcal{S} . Of the N subcarriers, N_p are dedicated as pilots,¹ and the remaining $N_d \triangleq N - N_p$ are used to transmit a total of M_t training bits and $M_d \triangleq N_d M - M_t$ coded/interleaved data bits. The data bits are generated by encoding M_i information bits using a rate- R coder, interleaving them, and partitioning the resulting $M_c \triangleq M_i/R$ bits among an integer number $Q \triangleq M_c/M_d$ of OFDM symbols. We note that the resulting scheme has a spectral efficiency of $\eta \triangleq M_d R/N$ information bits per channel use (bpcu).

In the sequel, we use $s^{(k)} \in \mathcal{S}$ for $k \in \{1, \dots, 2^M\}$ to denote the k th element of the QAM constellation, and $\mathbf{c}^{(k)} \triangleq [c_1^{(k)}, \dots, c_M^{(k)}]^T$ to denote the corresponding bits as defined by the symbol mapping. Likewise, we use $s_i[q] \in \mathcal{S}$ for the QAM symbol transmitted on the i th subcarrier of the q th OFDM symbol and $\mathbf{c}_i[q] \triangleq [c_{i,1}[q], \dots, c_{i,M}[q]]^T$ for the coded/interleaved bits corresponding to that symbol. We use $\mathbf{c}[q] \triangleq [c_0[q], \dots, c_{N-1}[q]]^T$ to denote the coded/interleaved bits in the q th OFDM symbol and $\mathbf{c} \triangleq [\mathbf{c}[1], \dots, \mathbf{c}[Q]]^T$ to denote the entire (interleaved) codeword. The elements of \mathbf{c} that are *a priori* known as pilot or training bits will be referred to as \mathbf{c}_{pt} . The remainder of \mathbf{c} is determined from the information bits $\mathbf{b} \triangleq [b_1, \dots, b_{M_i}]^T$ by coding/interleaving.

To modulate the q th OFDM symbol, an N -point inverse discrete Fourier transform (DFT) Φ^H is applied to the QAM sequence $\mathbf{s}[q] = [s_0[q], \dots, s_{N-1}[q]]^T$, yielding the time-domain sequence $\Phi^H \mathbf{s}[q] = \mathbf{a}[q] = [a_0[q], \dots, a_{N-1}[q]]^T$. The OFDM waveform $a(t)$ is then constructed using L -cyclic-prefixed versions of $\{a_j[q]\}$ and the transmission pulse $g_t(\tau)$:

$$a(t) = \sum_{q=1}^Q \sum_{j=-L}^{N-1} a_{\langle j \rangle_N}[q] g_t(t - jT - q(N+L)T) \quad (1)$$

with T denoting the baud interval (in seconds) and $L < N$.

The waveform $a(t)$ propagates through a noisy channel with an impulse response $h(\tau)$ that is supported on the interval $[\tau_{\min}, \tau_{\max}]$, resulting in the receiver input waveform

$$r(t) = w(t) + \int_{\tau_{\min}}^{\tau_{\max}} h(\tau) a(t - \tau) d\tau \quad (2)$$

where $w(t)$ is a Gaussian noise process with flat power spectral density N_o . We note that a time-invariant channel is assumed for simplicity. The receiver samples $r(t)$ through the reception pulse $g_r(\tau)$, obtaining

$$r_j[q] = \int r(t) g_r(jT + q(N+L)T - t) dt \quad (3)$$

¹For our GAMP decoder, we recommend $N_p = 0$; see Section IV.

and applies an N -DFT Φ to each time-domain sequence $\mathbf{r}[q] = [r_0[q], \dots, r_{N-1}[q]]^T$, yielding the frequency-domain sequences $\Phi \mathbf{r}[q] = \mathbf{y}[q] = [y_0[q], \dots, y_{N-1}[q]]^T$ for $q = 1 \dots Q$.

Defining the pulse-shaped channel response $x(\tau) \triangleq (g_r \star h \star g_t)(\tau)$, it is well known (e.g., [33]) that, when the support of $x(\tau)$ is contained within the interval $[0, LT)$, the frequency domain observation on the i th subcarrier can be written as

$$y_i[q] = s_i[q]z_i[q] + w_i[q] \quad (4)$$

where $z_i[q] \in \mathbb{C}$ is the i th subcarrier's gain and $\{w_i[q]\}$ are Gaussian noise samples. Furthermore, defining the uniformly sampled channel "taps" $x_j[q] \triangleq x(jT + q(N+L)T)$, the subcarrier gains are related to these taps through the DFT

$$z_i[q] = \sum_{j=0}^{L-1} \Phi_{ij} x_j[q]. \quad (5)$$

In addition, when $(g_r \star g_t)(\tau)$ is a Nyquist pulse, $\{w_i[q]\}_{\forall i, q}$ are statistically independent with variance $\nu^w = N_o$.

To simplify the development, we assume that $Q = 1$ in the sequel (but not in the simulations), and drop the index $[q]$ for brevity.

B. Clustered-Sparse Tap Prior

Empirical studies [3]–[6] have suggested that, when the baud rate T^{-1} is sufficiently large, the channel taps $\{x_j\}$ are "sparse" in that the tap distributions tend to be heavy tailed. The same empirical studies suggest that large taps tend to be clustered in the lag domain. Furthermore, both the sparsity and clustering behaviors can be lag-dependent, such as when the receiver's timing-synchronization mechanism aligns the first strong multipath arrivals with a particular reference lag j . A concrete example of these behaviors will be given in Section II-C.

Since our message-passing-based receiver design is inherently Bayesian, we seek a prior on the taps $\{x_j\}$ that is capable of representing this lag-dependent clustered sparsity. For this purpose, we assume a two-state Gaussian mixture (GM2) prior²

$$p(x_j) = (1 - \lambda_j) \mathcal{CN}(x_j; 0, \nu_j^0) + \lambda_j \mathcal{CN}(x_j; 0, \nu_j^1) \quad (6)$$

where $\nu_j^0 \geq 0$ denotes the variance while in the "small" state, $\nu_j^1 > \nu_j^0$ denotes the variance while in the "big" state, and $\lambda_j \triangleq \Pr\{d_j = 1\}$ denotes the prior probability of x_j being in the "big" state. Here, we use $d_j \in \{0, 1\}$ to denote the hidden state, implying the state-conditional pdf $p(x_j | d_j) = \mathcal{CN}(x_j; 0, \nu_j^{d_j})$.

For example, if x_j was presumed to be a "sparse" tap, then we would choose $\lambda_j \ll 1$ and $\nu_j^1 \gg \nu_j^0$ in (6). If, on the other hand, x_j is presumed to be (non-sparse) Rayleigh-fading, we would choose $\lambda_j = 1$ and set ν_j^1 equal to the tap variance, noting that ν_j^0 becomes inconsequential. If x_j is presumed to be Nakagami-fading or similar, we could fit the GM2 parameters $[\lambda_j, \nu_j^0, \nu_j^1]$ appropriately using the EM algorithm, as described in [34, p. 435]. The GM2 prior has been used successfully in many other non-Gaussian inference problems (see, e.g.,

²The message passing algorithm described in Section III-B can also handle non-Gaussian mixtures and/or mixtures with more than two terms.

[35]), and our premise here is that the GM2 model achieves a good balance between fidelity and tractability when modeling channel taps as well.

To capture the big-tap clustering behavior, we employ a hidden Markov model (HMM). For this, we model the tap states $\{d_j\}_{j=0}^{L-1}$ as a Markov chain (MC) with switching probabilities $p_j^{01} \triangleq \Pr\{d_{j+1} = 0 | d_j = 1\}$ and $p_j^{10} \triangleq \Pr\{d_{j+1} = 1 | d_j = 0\}$. Here, $p_j^{01} < 0.5$ implies that the neighbors of a big x_j tend to be big, and $p_j^{10} < 0.5$ implies that the neighbors of a small x_j tend to be small. We note that $\{p_j^{01}, p_j^{10}\}_{j=0}^{L-1}$ must be consistent with $\{\lambda_j\}_{j=0}^{L-1}$ in that the following must hold for all j

$$[\lambda_{j+1} \quad 1 - \lambda_{j+1}] = [\lambda_j \quad 1 - \lambda_j] \begin{bmatrix} 1 - p_j^{01} & p_j^{01} \\ p_j^{10} & 1 - p_j^{10} \end{bmatrix}. \quad (7)$$

Although we allow correlation among the tap states, we assume that the tap *amplitudes* are conditionally independent, i.e., $p(x_{j+1}, x_j | d_{j+1}, d_j) = p(x_j | d_j)p(x_{j+1} | d_{j+1})$. Our experiences with IEEE 802.15.4a channels (see below) suggest that this is a valid assumption.

We emphasize that the model parameters $\{\lambda_j, p_j^{01}, p_j^{10}, \nu_j^1, \nu_j^0\}$ are allowed to vary with lag j , facilitating the exploitation of *a priori* known lag-dependencies in sparsity and/or clustering.

C. Illustrative Example: IEEE 802.15.4a Channels

As an illustrative example of the clustered-sparse tap behavior described above, we generated realizations of the tap vector $\mathbf{x} \triangleq [x_0, \dots, x_{L-1}]^T$ from channel impulse responses $h(\tau)$ generated according to the method specified in the IEEE 802.15.4a "ultra-wideband" standard [16], which uses the Saleh-Valenzuela model [36]

$$h(\tau) = \sum_{c=0}^C \sum_{k=0}^K h_{k,c} e^{j\phi_{k,c}} \delta(\tau - T_c - \tau_{k,c}) \quad (8)$$

where C denotes the number of clusters, T_c the delay of the c th cluster, K the number of components per cluster, $\{\tau_{k,c}\}$ the relative component delays, $\{h_{k,c}\}$ the component amplitudes, and $\{\phi_{k,c}\}$ the component phases. In particular, the 802.15.4a standard specifies the following.

- The cluster arrival times are a Poisson process with rate Λ , i.e., $p(T_c | T_{c-1}) = \Lambda \exp(-\Lambda(T_c - T_{c-1}))$. The initial cluster delay $T_0 \geq \tau_{\min}$, as seen by the receiver, is a function of the timing synchronization algorithm.
- The component arrivals are a mixture of two Poisson processes: $p(\tau_{k,c} | \tau_{k-1,c}) = \beta \lambda_1 \exp(-\lambda_1(\tau_{k,c} - \tau_{k-1,c})) + (1 - \beta) \lambda_2 \exp(-\lambda_2(\tau_{k,c} - \tau_{k-1,c}))$ with $\tau_{0,c} = 0$.
- The component energies obey

$$E\{|h_{k,c}|^2\} = \frac{\exp(-T_l/\Gamma - \tau_{k,l}/\gamma)}{\gamma[(1 - \beta)\lambda_1 + \beta\lambda_2 + 1]} \quad (9)$$

where Γ is the cluster decay time constant and γ is the intra-cluster decay time constant.

- The amplitudes $\{h_{k,c}\}$ are independent and identically distributed (i.i.d.) Nakagami with m -factors randomly generated via i.i.d. $m \sim \mathcal{N}(m_0, \hat{m}_0^2)$.
- The phases $\{\phi_{k,c}\}$ are i.i.d. uniform on $[0, 2\pi)$.




Fig. 1. Histograms of $\text{Re}(x_j)$ for lags $j \in \{5, 23, 128, 230\}$, with “tight” axes. With synchronization delay $L_{\text{pre}} = 20$, note that the histogram appears Gaussian for $j < L_{\text{pre}}$, Laplacian for $j \approx L_{\text{pre}}$, and very sparse for $j \gg L_{\text{pre}}$.

- The number of clusters, C , is Poisson distributed with mean \bar{C} , i.e., $p(C) = (\bar{C})^C \exp(-\bar{C}) / (C!)$.
- The number of components per cluster, K , is set large enough to yield a desired modeling accuracy.

Beyond the above specifications, we assume the following.

- The parameters $\{\Lambda, \lambda_1, \lambda_2, \beta, \Gamma, \gamma, m_0, \hat{m}_0, \bar{C}\}$ are set according to the 802.15.4a “outdoor NLOS” scenario [16].
- $K = 100$ components per cluster are used.
- The pulses $g_t(\tau)$ and $g_r(\tau)$ are square-root raised cosine (SRRC) designs with parameter 0.5.
- The system bandwidth equals $T^{-1} = 256$ MHz.
- The number of taps (and CP length) was set at $L = 256$ (implying a maximal delay spread of 1ν s) in order to capture all significant energy in $h(\tau)$.
- The initial delay was generated via $T_0 = L_{\text{pre}}T + \tilde{T}_0$, where $L_{\text{pre}} = 20$ and where \tilde{T}_0 is exponentially distributed with mean T , i.e., $p(\tilde{T}_0) = \Lambda_0 \exp(-\Lambda_0 \tilde{T}_0)$ for $\Lambda_0 = 1/T$. Here, L_{pre} was chosen so that $\{x_j\}_{j=0}^{L_{\text{pre}}}$ captures the “pre-cursor” energy contributed by the pulse shape, while Λ_0 models a positive synchronization uncertainty.

We now show results from an experiment conducted using $U = 10\,000$ realizations of the tap vector \mathbf{x} . In Fig. 1, we show histograms of $\text{Re}(x_j)$ for lags $j \in \{5, 23, 128, 230\}$. There it can be seen that the empirical distribution of $\text{Re}(x_j)$ changes significantly with lag j : for pre-cursor lags $j < L_{\text{pre}}$, it is approximately Gaussian; for near-cursor lags $j \approx L_{\text{pre}}$, it is approximately Laplacian; and, for post-cursor lags $j \gg L_{\text{pre}}$, it is extremely heavy-tailed. In Fig. 2, we show a typical realization of \mathbf{x} and notice clustering among the big taps. For comparison, we also plot an empirical estimate of the power-delay profile (PDP) $\rho \triangleq [\rho_0, \dots, \rho_{L-1}]^T$ in Fig. 2, where $\rho_j \triangleq \mathbb{E}\{|x_j|^2\}$.

Next, we fit the GM2 parameters $\{\lambda_j, \nu_j^0, \nu_j^1\}_{j=0}^{L-1}$ from the realizations $\{\mathbf{x}_u\}_{u=1}^U$ using the EM algorithm [34, p. 435], which iterates the steps (10)–(13) until convergence:

$$\omega_{j,u} = \frac{\lambda_j \mathcal{CN}(x_{j,u}; 0, \nu_j^1)}{(1-\lambda_j) \mathcal{CN}(x_{j,u}; 0, \nu_j^0) + \lambda_j \mathcal{CN}(x_{j,u}; 0, \nu_j^1)} \forall j, u \quad (10)$$




Fig. 2. Sample realization of channel taps $\{x_j\}$ generated from the IEEE 802.15.4a model with SRRC pulse shaping. Also shown is the empirically estimated PDP, best fits of the GM2 parameters $\{\nu_j^0, \nu_j^1\}$, and the MAP threshold for detecting the hidden state d_j given the tap value x_j .




Fig. 3. Empirically estimated statistics on the tap-states $\{d_j\}$. Top: $\lambda_j \triangleq \Pr\{d_j = 1\}$, middle: $p_j^{01} \triangleq \Pr\{d_{j+1} = 0 \mid d_j = 1\}$, bottom: $p_j^{10} \triangleq \Pr\{d_{j+1} = 1 \mid d_j = 0\}$. The red dashed line shows the synchronization reference, $j = L_{\text{pre}} = 20$.

$$\nu_j^1 = \frac{\sum_{u=1}^U \omega_{j,u} |x_{j,u}|^2}{\sum_{u=1}^U \omega_{j,u} \forall j} \quad (11)$$

$$\nu_j^0 = \frac{\sum_{u=1}^U (1 - \omega_{j,u}) |x_{j,u}|^2}{\sum_{u=1}^U (1 - \omega_{j,u}) \forall j} \quad (12)$$

$$\lambda_j = \frac{1}{U} \sum_{u=1}^U \omega_{j,u} \forall j. \quad (13)$$

Above, $\omega_{j,u}$ is the posterior on the state $d_{j,u}$ of tap $x_{j,u}$, i.e., $\omega_{j,u} = \Pr\{d_{j,u} = 1 \mid x_{j,u}; \lambda_j, \nu_j^0, \nu_j^1\}$. The EM-estimated big-variance profile $\boldsymbol{\nu}^1 \triangleq [\nu_0^1, \dots, \nu_{L-1}^1]^T$ and small-variance profile $\boldsymbol{\nu}^0$ are shown in Fig. 2, while the sparsity profile $\boldsymbol{\lambda} \triangleq [\lambda_0, \dots, \lambda_{L-1}]^T$ is shown in Fig. 3. Not surprisingly, the best-fit GM2 parameters also change significantly with lag j . In particular, as j becomes larger, the variance ratio ν_j^1/ν_j^0

increases while the big-tap-probability λ_j decreases, corresponding to an increase in sparsity. Meanwhile, there exists a peak in λ_j near $j = L_{\text{pre}}$ that results from synchronization.

Next, we empirically estimated the switching probabilities $\mathbf{p}^{01} \triangleq [p_0^{01}, \dots, p_{L-1}^{01}]^T$ and \mathbf{p}^{10} using maximum *a posteriori* (MAP) state estimates, i.e., $\hat{d}_{j,u} = \lfloor \omega_{j,u} + 0.5 \rfloor$. In particular,

$$p_j^{01} = \frac{\sum_{u=1}^U \mathbf{1}_{\{\hat{d}_{j+1,u}=0 \ \& \ \hat{d}_{j,u}=1\}}}{\sum_{u=1}^U \mathbf{1}_{\{\hat{d}_{j,u}=1\}}} \quad (14)$$

$$p_j^{10} = \frac{\sum_{u=1}^U \mathbf{1}_{\{\hat{d}_{j+1,u}=1 \ \& \ \hat{d}_{j,u}=0\}}}{\sum_{u=1}^U \mathbf{1}_{\{\hat{d}_{j,u}=0\}}} \quad (15)$$

where $\mathbf{1}_{\{A\}}$ denotes the indicator function for event A . From the plots in Fig. 3, we see that the estimated switching probabilities are lag-dependent as well.

Finally, using the MAP state estimates $\{\hat{d}_{j,u}\}$, we empirically estimated the normalized conditional correlation

$$\frac{\sum_{u=1}^U \mathbf{1}_{\{\hat{d}_{j+1,u}=1, \hat{d}_{j,u}=1\}} x_{j+1,u} x_{j,u}^*}{\sqrt{\sum_{u=1}^U \mathbf{1}_{\{\hat{d}_{j+1,u}=1, \hat{d}_{j,u}=1\}} |x_{j+1,u}|^2 \sum_{u=1}^U \mathbf{1}_{\{\hat{d}_{j+1,u}=1, \hat{d}_{j,u}=1\}} |x_{j,u}|^2}}$$

and found that the magnitudes were < 0.1 , validating our assumption of conditionally independent tap amplitudes.

In summary, we see that IEEE 802.15.4a channels do indeed yield taps with the lag-dependent clustered sparsity described in Section II-B. Moreover, we have shown how the GM2-HMM parameters can be estimated from realizations of \mathbf{x} . Next, we propose an efficient factor-graph-based approach to joint channel-estimation and decoding (JCED) for BICM-OFDM using the GM2-HMM prior proposed in Section II-B.

III. JOINT CHANNEL ESTIMATION AND DECODING

Our goal is to infer the information bits \mathbf{b} from the OFDM observations \mathbf{y} and the pilot/training bits \mathbf{c}_{pt} , without knowing the channel state \mathbf{x} . In particular, we aim to maximize the posterior pmf $p(\mathbf{b}_m | \mathbf{y}, \mathbf{c}_{\text{pt}})$ of each info bit. To exploit prior knowledge that \mathbf{x} is clustered-sparse, we employ the GM2-HMM prior described in Section II-B. As a result, the info-bit posterior can be decomposed into the following product of factors:

$$\begin{aligned} & p(\mathbf{b}_m | \mathbf{y}, \mathbf{c}_{\text{pt}}) \\ &= \sum_{\mathbf{b}_{-m}} p(\mathbf{b} | \mathbf{y}, \mathbf{c}_{\text{pt}}) \\ &\propto \sum_{\mathbf{b}_{-m}} p(\mathbf{y} | \mathbf{b}, \mathbf{c}_{\text{pt}}) p(\mathbf{b}) \\ &= \int_{\mathbf{x}} \sum_{\mathbf{s}, \mathbf{d}, \mathbf{c}, \mathbf{b}_{-m}} p(\mathbf{y} | \mathbf{s}, \mathbf{x}) p(\mathbf{x} | \mathbf{d}) p(\mathbf{d}) \\ &\quad \times p(\mathbf{s} | \mathbf{c}) p(\mathbf{c} | \mathbf{b}, \mathbf{c}_{\text{pt}}) p(\mathbf{b}) \\ &= \int_{\mathbf{x}} \sum_{\mathbf{d}} \prod_{j=0}^{L-1} p(x_j | d_j) p(d_j | d_{j-1}) \sum_{\mathbf{s}} \prod_{i=0}^{N-1} p(y_i | s_i, \mathbf{x}) \\ &\quad \times \sum_{\mathbf{c}} p(s_i | \mathbf{c}_i) \sum_{\mathbf{b}_{-m}} p(\mathbf{c} | \mathbf{b}, \mathbf{c}_{\text{pt}}) \prod_{m=1}^{M_t} p(\mathbf{b}_m) \end{aligned} \quad (17)$$

where $\mathbf{b}_{-m} \triangleq [b_1, \dots, b_{m-1}, b_{m+1}, \dots, b_{M_t}]^T$. This factorization is illustrated by the *factor graph* in Fig. 4, where the round

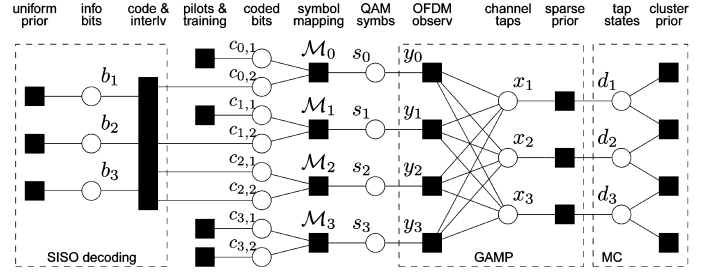


Fig. 4. Factor graph of the JCED problem for a toy example with $M_i = 3$ information bits, $N_p = 1$ pilot subcarrier (at subcarrier index $i = 3$), $M_t = 2$ training bits, $M = 2$ bits per QAM symbol, $N = 4$ OFDM subcarriers, and channel impulse response length $L = 3$.

nodes represent random variables and the square nodes represent the factors of the posterior exposed in (17).

A. Background on Belief Propagation

Although exact evaluation of the posteriors $\{p(\mathbf{b}_m | \mathbf{y}, \mathbf{c}_{\text{pt}})\}$ is computationally impractical for the problem sizes of interest, these posteriors can be approximately evaluated using *belief propagation* (BP) [37] on the factor graph in Fig. 4. In text-book BP, beliefs take the form of pdfs/pmfs that are propagated among nodes of the factor graph via the *sum/product algorithm* (SPA) [18]–[20].

- 1) Say the factor node f is connected to the variable nodes $\{v_a\}_{a=1}^A$. The belief passed from f to v_b is $p_{f \rightarrow v_b}(v_b) \propto \int_{\{v_a\}_{a \neq b}} f(v_1, \dots, v_A) \prod_{a \neq b} p_{v_a \rightarrow f}(v_a)$, given the beliefs $\{p_{v_a \rightarrow f}(\cdot)\}_{a \neq b}$ recently passed to f .
- 2) Say the variable node v is connected to the factor nodes $\{f_1, \dots, f_B\}$. The belief passed from v to f_a is $p_{v \rightarrow f_a}(v) \propto \prod_{b \neq a} p_{f_b \rightarrow v}(v)$, given the beliefs $\{p_{f_b \rightarrow v}(\cdot)\}_{b \neq a}$ recently passed to v .
- 3) Say the variable node v is connected to the factor nodes $\{f_1, \dots, f_B\}$. The posterior on v is the product of all recently arriving beliefs, i.e., $p(v) \propto \prod_{b=1}^B p_{f_b \rightarrow v}(v)$.

When the factor graph contains no loops, SPA-BP yields exact posteriors after two rounds of message passing (i.e., forward and backward), but, in the presence of loops, convergence to the exact posteriors is not guaranteed [25]. That said, there exist many problems to which loopy BP [37] has been successfully applied, including inference on Markov random fields [38], LDPC decoding [14], and compressed sensing [13], [15], [32], [39]–[41]. Our work not only leverages these past successes, but unites them.

B. Background on GAMP

An important sub-problem within our larger bit-inference problem is the estimation of a vector of independent possibly non-Gaussian variables \mathbf{x} that are linearly mixed via $\Phi \in \mathbb{C}^{N \times L}$ to form $\mathbf{z} = \Phi \mathbf{x} = [z_0, \dots, z_{N-1}]^T$, and subsequently observed as noisy measurements \mathbf{y} through the possibly non-Gaussian pdfs $\{p_{Y_i | Z_i}(\cdot | \cdot)\}_{i=0}^{N-1}$. In our case, (6) specifies a GM2 prior on x_j and (4)—given the finite-alphabet uncertainty in s_i —yields the non-Gaussian measurement pdf $p_{Y_i | Z_i}$. This “linear mixing” sub-problem is described by the factor graph shown within the middle dashed box in Fig. 4, where each node “ y_i ” represents the measurement pdf $p_{Y_i | Z_i}$.

TABLE I
GAMP ALGORITHM

definitions:	
$p_{Z_i Y_i}(z y; \hat{z}, \nu^z) = \frac{p_{Y_i Z_i}(y z) \mathcal{CN}(z; \hat{z}, \nu^z)}{\int_{z'} p_{Y_i Z_i}(y z') \mathcal{CN}(z'; \hat{z}, \nu^z)}$	(D1)
$g_{\text{out},i}(y, \hat{z}, \nu^z) = \frac{1}{\nu^z} (\mathbb{E}_{Z_i Y_i} \{z y; \hat{z}, \nu^z\} - \hat{z})$	(D2)
$g'_{\text{out},i}(y, \hat{z}, \nu^z) = \frac{1}{\nu^z} \left(\frac{\text{var}_{Z_i Y_i} \{z y; \hat{z}, \nu^z\}}{\nu^z} - 1 \right)$	(D3)
$p_{X_j}(x; \hat{r}, \nu^r) = \frac{p_{X_j}(x) \mathcal{CN}(x; \hat{r}, \nu^r)}{\int_{x'} p_{X_j}(x') \mathcal{CN}(x'; \hat{r}, \nu^r)}$	(D4)
$g_{\text{in},j}(\hat{r}, \nu^r) = \int_x x p_{X_j}(x; \hat{r}, \nu^r)$	(D5)
$g'_{\text{in},j}(\hat{r}, \nu^r) = \frac{1}{\nu^r} \int_x x - g_{\text{in},j}(\hat{r}, \nu^r) ^2 p_{X_j}(x; \hat{r}, \nu^r)$	(D6)
initialize:	
$\forall j : \hat{x}_j(1) = \int_x x p_{X_j}(x)$	(I1)
$\forall j : \nu_j^x(1) = \int_x x - \hat{x}_j(1) ^2 p_{X_j}(x)$	(I2)
$\forall i : \hat{u}_i(0) = 0$	(I3)
for $n = 1, 2, 3, \dots$	
$\forall i : \hat{z}_i(n) = \sum_{j=0}^{L-1} \Phi_{ij} \hat{x}_j(n)$	(R1)
$\forall i : \nu_i^z(n) = \sum_{j=0}^{L-1} \Phi_{ij} ^2 \nu_j^x(n)$	(R2)
$\forall i : \hat{p}_i(n) = \hat{z}_i(n) - \nu_i^z(n) \hat{u}_i(n-1)$	(R3)
$\forall i : \hat{u}_i(n) = g_{\text{out},i}(y_i, \hat{p}_i(n), \nu_i^z(n))$	(R4)
$\forall i : \nu_i^u(n) = -g'_{\text{out},i}(y_i, \hat{p}_i(n), \nu_i^z(n))$	(R5)
$\forall j : \nu_j^r(n) = \left(\sum_{i=0}^{N-1} \Phi_{ij} ^2 \nu_i^u(n) \right)^{-1}$	(R6)
$\forall j : \hat{r}_j(n) = \hat{x}_j(n) + \nu_j^r(n) \sum_{i=0}^{N-1} \Phi_{ij}^* \hat{u}_i(n)$	(R7)
$\forall j : \nu_j^x(n+1) = \nu_j^r(n) g'_{\text{in},j}(\hat{r}_j(n), \nu_j^r(n))$	(R8)
$\forall j : \hat{x}_j(n+1) = g_{\text{in},j}(\hat{r}_j(n), \nu_j^r(n))$	(R9)
end	

and the node rightward of each node “ x_j ” represents the GM2 prior on x_j .

Building on recent work on multiuser detection by Guo and Wang [42], as well as recent work on message passing algorithms for compressed sensing by Donoho, Maleki, Montanari, and Bayati [40], [41], Rangan proposed a so-called *generalized approximate message passing* (GAMP) scheme that, for the sub-problem described above, admits rigorous analysis³ as $N, L \rightarrow \infty$ [13]. The main ideas behind GAMP are the following. First, although the beliefs flowing leftward from the nodes $\{x_j\}$ are clearly non-Gaussian, the corresponding belief about $z_i = \sum_{j=0}^{L-1} \Phi_{ij} x_j$ can be accurately approximated as Gaussian, when L is large, using the central limit theorem. Moreover, to calculate the parameters of this distribution (i.e., its mean and variance), only the mean and variance of each x_j are needed. Thus, it suffices to pass only means and variances leftward from each x_j node. It is similarly desirable to pass only means and variances rightward from each measurement node. Although the exact rightward flowing beliefs would be non-Gaussian (due to the non-Gaussian assumption on the measurement channels $p_{Y_i|Z_i}$), GAMP approximates them as Gaussian using a second-order Taylor series, and passes only the resulting means and variances. A further simplification employed by GAMP is to approximate the *differences* among the outgoing means/variances of each left node, and the incoming means/variances of each right node, using Taylor series. The GAMP algorithm⁴ is summarized in Table I.

³Since it is difficult to give a concise yet accurate account of GAMP’s technical properties, we refer the interested reader to [13].

⁴To be precise, the GAMP algorithm in Table I is an extension of that proposed in [13]. Table I handles circular *complex-valued* distributions and *non-identically* distributed signals and measurements.

C. Joint Estimation and Decoding Using GAMP

We now detail our application of GAMP to joint channel-estimation and decoding (JCED) under the GM2-HMM tap prior, frequently referring to the factor graph in Fig. 4.

Because our factor graph is loopy, there exists considerable freedom in the message passing schedule. Roughly speaking, we choose to pass messages from the left to the right of Fig. 4 and back again, several times, stopping as soon as the messages converge. Each of these full cycles of message passing will be referred to as a “turbo iteration.” However, during a single turbo iteration, there may be multiple iterations of message passing *between* the GAMP and MC sub-graphs, which will be referred to as “equalizer” iterations. Furthermore, during a single equalizer iteration, there may be multiple iterations of message passing *within* the GAMP sub-graph, while there is at most one forward-backward iteration *within* the MC sub-graph. Finally, the SISO decoding block may itself be implemented using message passing, in which case it may also use several internal iterations. The message passing details are discussed below.

At the start of the first turbo iteration, there is total uncertainty about the information bits, so that $\Pr\{b_m = 1\} = (1)/(2) \forall m$. Thus, the initial bit beliefs flowing rightward out of the coding/interleaving block are uniformly distributed. Meanwhile, the pilot/training bits are known with certainty.

Coded-bit beliefs are then propagated rightward into the symbol mapping nodes. Since the symbol mapping is deterministic, the corresponding pdf factors take the form $p(s^{(k)} | \mathbf{c}^{(l)}) = \delta_{k-l}$. The SPA dictates that the message passed rightward from symbol mapping node “ \mathcal{M}_i ” takes the form

$$p_{\mathcal{M}_i \rightarrow s_i}(s^{(k)}) \propto \sum_{\mathbf{c} \in \{0,1\}^M} p(s^{(k)} | \mathbf{c}) \times \prod_{m=1}^M p_{c_{i,m} \rightarrow \mathcal{M}_i}(c_m) \quad (18)$$

$$= \prod_{m=1}^M p_{c_{i,m} \rightarrow \mathcal{M}_i}(c_m^{(k)}) \quad (19)$$

which is then copied forward as the message passed rightward from node s_i (i.e., $p_{\mathcal{M}_i \rightarrow s_i}(s^{(k)}) = p_{s_i \rightarrow y_i}(s^{(k)})$).

Recall, from Section III-B, that the symbol-belief passed rightward into the measurement node “ y_i ” determines the pdf $p_{Y_i|Z_i}$ used in GAMP. Writing this symbol belief as $\beta_i \triangleq [\beta_i^{(1)}, \dots, \beta_i^{(|S|)}]^\top$ for $\beta_i^{(k)} \triangleq p_{s_i \rightarrow y_i}(s^{(k)})$, (4) implies the measurement pdf

$$p_{Y_i|Z_i}(y|z) = \sum_{k=1}^{|S|} \beta_i^{(k)} \mathcal{CN}(y; s^{(k)} z; \nu^w). \quad (20)$$

From (20), it is shown in Appendix A that the quantities in (D2)–(D3) of Table I become

$$g_{\text{out},i}(y, \hat{z}, \nu^z) = \frac{1}{\nu^z} \hat{e}_i(y, \hat{z}, \nu^z) \quad (21)$$

$$g'_{\text{out},i}(y, \hat{z}, \nu^z) = \frac{1}{\nu^z} \left(\frac{\nu_i^e(y, \hat{z}, \nu^z)}{\nu^z} - 1 \right) \quad (22)$$

for

$$\xi_i^{(k)}(y, \hat{z}, \nu^z) \triangleq \frac{\beta_i^{(k)} \mathcal{CN}(y; s^{(k)} \hat{z}, |s^{(k)}|^2 \nu^z + \nu^w)}{\sum_{k'} \beta_i^{(k')} \mathcal{CN}(y; s^{(k')} \hat{z}, |s^{(k')}|^2 \nu^z + \nu^w)} \quad (23)$$

$$\zeta^{(k)}(\nu^z) \triangleq \frac{|s^{(k)}|^2 \nu^z}{|s^{(k)}|^2 \nu^z + \nu^w} \quad (24)$$

$$\hat{e}^{(k)}(y, \hat{z}, \nu^z) \triangleq \left(\frac{y}{s^{(k)}} - \hat{z} \right) \zeta^{(k)}(\nu^z) \quad (25)$$

$$\hat{e}_i(y, \hat{z}, \nu^z) \triangleq \sum_{k=1}^{|\mathcal{S}|} \xi_i^{(k)}(y, \hat{z}, \nu^z) \hat{e}^{(k)}(y, \hat{z}, \nu^z) \quad (26)$$

$$\nu_i^e(y, \hat{z}, \nu^z) \triangleq \sum_{k=1}^{|\mathcal{S}|} \xi_i^{(k)}(y, \hat{z}, \nu^z) \left(\left| \hat{e}^{(k)}(y, \hat{z}, \nu^z) \right|^2 + \frac{\nu^w \zeta^{(k)}(\nu^z)}{s^{(k)}} \right) - \hat{e}_i(y, \hat{z}, \nu^z) \quad (27)$$

where $\xi_i \triangleq [\xi_i^{(1)}, \dots, \xi_i^{(|\mathcal{S}|)}]^\top$ characterizes the posterior pmf on s_i under the channel model $z_i \sim \mathcal{CN}(\hat{z}, \nu^z)$. Likewise, from (6), it is shown in Appendix B that the quantities (D5)–(D6) take the form

$$g_{\text{in},j}(\hat{r}, \nu^r) = (\alpha_j \gamma_j^1 + (1 - \alpha_j) \gamma_j^0) \hat{r} \quad (28)$$

$$g'_{\text{in},j}(\hat{r}, \nu^r) = \alpha_j (1 - \alpha_j) (\gamma_j^1 - \gamma_j^0)^2 |\hat{r}|^2 / \nu^r + \alpha_j \gamma_j^1 + (1 - \alpha_j) \gamma_j^0 \quad (29)$$

for $\gamma_j^0(\nu^r) \triangleq (1 + \nu^r / \nu_j^0)^{-1}$ and $\gamma_j^1(\nu^r) \triangleq (1 + \nu^r / \nu_j^1)^{-1}$ and

$$\alpha_j(\hat{r}, \nu^r) \triangleq \frac{1}{1 + \underbrace{\left(\frac{\lambda_j}{1 - \lambda_j} \frac{\mathcal{CN}(\hat{r}; 0, \nu_j^1 + \nu^r)}{\mathcal{CN}(\hat{r}; 0, \nu_j^0 + \nu^r)} \right)^{-1}}_{\mathcal{L}_j^{\text{apri}} \quad \mathcal{L}_j^{\text{ext}}(\hat{r}, \nu^r)}} \quad (30)$$

Above, $\mathcal{L}_j^{\text{apri}}$ is the *a priori* likelihood ratio ($\Pr\{d_j = 1\} / \Pr\{d_j = 0\}$) on the hidden state, $\mathcal{L}_j^{\text{ext}}(\hat{r}, \nu^r)$ is GAMP's extrinsic likelihood ratio, and $\alpha_j(\hat{r}, \nu^r)$ is the corresponding posterior probability that $d_j = 1$.

Using (21)–(30), the GAMP algorithm in Table I is iterated until it converges.⁵ In doing so, GAMP generates (a close approximation to) both the conditional means $\hat{\mathbf{x}}$ and variances $\nu^x \triangleq [\nu_0^x, \dots, \nu_{L-1}^x]^\top$ given the observations \mathbf{y} , the soft symbol priors $\beta \triangleq [\beta_0, \dots, \beta_{L-1}]^\top$ and the sparsity prior λ . Conveniently, GAMP also returns (close approximations to) both the conditional means $\hat{\mathbf{z}}$ and variances ν^z of the subchannel gains z , as well as posteriors $\xi \triangleq [\xi_0, \dots, \xi_{L-1}]^\top$ on the symbols \mathbf{s} .

Before continuing, we discuss some GAMP details that are specific to our OFDM-JCED application. First, we notice that, to guarantee that the variance $\nu_i^u(n)$ in (R5) is positive, we must have $\nu_i^e < \nu_z$ in (22). Since this is not necessarily the case during the first few GAMP iterations, we clip ν_i^e at the value $0.99\nu^z$, where 0.99 was chosen heuristically. Second, due to unit-modulus property of the DFT elements Φ_{ij} , step (R2) in Table I simplifies to $\nu_i^z(n) = \sum_j \nu_j^x(n)$ and (R6) simplifies

⁵More precisely, GAMP is iterated until the mean-square tap-estimate difference $\sum_{j=0}^{L-1} |\hat{x}_j(n) - \hat{x}_j(n-1)|^2 / L$ falls below a threshold or a maximum number of GAMP iterations has elapsed.

to $\nu_j^x(n) = (\sum_i \nu_i^u(n))^{-1}$. With these simplifications, the complexity of GAMP is dominated by either the matrix-vector products $\sum_j \Phi_{ij} \hat{x}_j(n)$ in (R1) and $\sum_i \Phi_{ij}^* \hat{u}_i(n)$ in (R7), which can be implemented using a $N \log_2 N$ -multiply FFT when N is a power-of-two, or by the calculation of $\{\hat{e}_i, \nu_i^e\}_{i=0}^{N-1}$ in (26)–(27), which requires $\mathcal{O}(N|\mathcal{S}|)$ multiplies. Thus, GAMP requires only $\mathcal{O}(N \log_2 N + N|\mathcal{S}|)$ multiplies per iteration.

After the messages within the GAMP sub-graph have converged, tap-state beliefs are passed rightward to the MC sub-graph. In particular, the SPA dictates that GAMP passes tap-state likelihoods or, equivalently, the extrinsic likelihood ratios $\mathcal{L}_j^{\text{ext}}$. Since the MC sub-graph is non-loopy, only one iteration of forward–backward message passing is performed,⁶ after which the resulting tap-state likelihoods are passed leftward back to GAMP, where they are treated as tap-state priors λ in the next equalizer iteration. This interaction between the GAMP and MC sub-blocks can be recognized as an incarnation of the structured-sparse reconstruction scheme recently proposed by the authors in [15].

When the tap-state likelihoods passed between GAMP and MC have converged,⁷ the equalizer iterations are terminated and messages are passed leftward from the GAMP block. For this, SPA dictates that a symbol-belief propagates leftward from the y_i node with the form

$$p_{s_i \leftarrow y_i}(s) \propto \int_z \mathcal{CN}(y_i; sz, \nu^w) \mathcal{CN}(z; \hat{z}_i, \nu_i^z) \quad (31)$$

$$= \mathcal{CN}(y_i; s \hat{z}_i, |s|^2 \nu_i^z + \nu^w) \quad (32)$$

where (\hat{z}_i, ν_i^z) play the role of soft channel estimates. The SPA then implies that $p_{\mathcal{M}_i \leftarrow s_i}(s) = p_{s_i \leftarrow y_i}(s)$.

Next, beliefs are passed leftward from each symbol-mapping node \mathcal{M}_i to the corresponding bit nodes $c_{i,m}$. From the SPA, they take the form

$$\begin{aligned} p_{c_{i,m} \leftarrow \mathcal{M}_i}(c) & \propto \sum_{k=1}^{|\mathcal{S}|} \sum_{\mathbf{c}: c_m = c} p(s^{(k)} | \mathbf{c}) p_{\mathcal{M}_i \leftarrow s_i}(s^{(k)}) \\ & \times \prod_{m' \neq m} p_{c_{i,m'} \rightarrow \mathcal{M}_i}(c_{m'}) \\ & = \sum_{k: c_m^{(k)} = c} p_{\mathcal{M}_i \leftarrow s_i}(s^{(k)}) \frac{\prod_{m'=1}^M p_{c_{i,m'} \rightarrow \mathcal{M}_i}(c_{m'}^{(k)})}{p_{c_{i,m} \rightarrow \mathcal{M}_i}(c)} \end{aligned} \quad (33)$$

$$= \frac{1}{p_{c_{i,m} \rightarrow \mathcal{M}_i}(c)} \sum_{k: c_m^{(k)} = c} p_{\mathcal{M}_i \leftarrow s_i}(s^{(k)}) p_{\mathcal{M}_i \rightarrow s_i}(s^{(k)}) \quad (34)$$

for pairs (i, m) that do not correspond to pilot/training bits. (Since the pilot/training bits are known with certainty, there is no need to update their pmfs.)

Finally, messages are passed leftward into the coding/interleaving block. Doing so is equivalent to feeding extrinsic soft

⁶Message passing on the MC factor graph is a standard procedure. For details, we refer the reader to [14], [34].

⁷More precisely, the equalizer iterations are terminated when the mean-square difference in tap-state log-likelihoods falls below a threshold or a maximum number of equalizer iterations has elapsed.

bit estimates to a soft-input/soft-output (SISO) decoder/deinterleaver, which treats them as priors. Since SISO decoding is a well-studied topic [14], [43] and high-performance implementations are readily available (e.g., [44]), we will not elaborate on the details here. It suffices to say that, once the extrinsic outputs of the SISO decoder have been computed, they are re-interleaved and passed rightward from the coding/interleaving block to begin another turbo iteration. These turbo iterations continue until either the decoder detects no bit errors, the soft bit estimates have converged, or a maximum number of iterations has elapsed.

IV. NUMERICAL RESULTS

In this section, we present numerical results⁸ that compare JCED using our GAMP-based scheme to that using soft-input soft-output (SISO) equalizers based on linear MMSE (LMMSE) and LASSO [9], as well as to performance bounds based on perfect channel state information (CSI).

A. Setup

For all results, we used irregular LDPC codes with code-word length ~ 10000 and average column weight 3, generated (and decoded) using the publicly available software [44], with random interleaving. We focus on the case of $N = 1024$ sub-carrier OFDM with 16-QAM (i.e., $M = 4$) operating at a spectral efficiency of $\eta = 2$ bpcu. For bit-to-symbol mapping, we used multilevel Gray-mapping [45], noting recent work [46] that conjectures the optimality of Gray-mapping when BICM is used with a strong code. In some simulations, we used $N_p > 0$ pilot-only subcarriers and $M_t = 0$ interspersed training bits, whereas in others we used $N_p = 0$ and $M_t > 0$. When $N_p > 0$, the pilot subcarriers were placed randomly and modulated with (known) QAM symbols chosen uniformly at random. When $M_t > 0$, the training bits were placed at the most significant bits (MSBs) of uniformly spaced data-subcarriers and modulated with the bit value 1.

Realizations of the tap vector $\mathbf{x}[q]$ were generated from IEEE 802.15.4a outdoor-NLOS impulse responses and SRRC pulses, as described in Section II-C, and *not* from the GM2-HMM model. The tap vectors generated for our simulations are thus as realistic as one can hope to obtain in software. All reported results are averaged over 5000 channel realizations (i.e., 10^7 info bits).

The GM2-HMM parameters $\nu^0, \nu^1, \mathbf{p}^{01}, \mathbf{p}^{10}$ were fit from 10000 realizations of the tap-vector \mathbf{x} using the procedure described in Section II-C. In doing so, we implicitly assumed⁹ that the receiver is designed for the outdoor scenario, and we leverage the prior information made available by the extensive measurement campaign conducted for the IEEE 802.15.4a standard [16]. In all cases, we used a *maximum* of 10 turbo iterations,

⁸Matlab code is available at <http://www.ece.osu.edu/~schniter/turboGAMPdecoding>.

⁹If, instead, we knew that the receiver would be used in a different operating scenario, then we could generate representative realizations of \mathbf{x} for that scenario and fit the GM2-HMM parameters accordingly. Furthermore, one could optimize the receiver for any desired balance between “typical” and “worst case” operating conditions by simply choosing appropriate training realizations \mathbf{x} .

5 equalizer iterations, 15 GAMP iterations, and 25 LDPC decoder iterations, although in most cases the iterations converged early (as described in Section III-C).

B. Comparison With Other Schemes

The proposed GAMP-based equalizer was compared with soft-input soft-output (SISO) equalizers based on LMMSE and LASSO [9], whose constructions are now detailed.

All SISO equalizers are provided with the soft inputs $\hat{\mathbf{s}}[q]$ and $\nu^s[q]$, i.e., the means and variances, respectively, of the symbols $\mathbf{s}[q] \in \mathbb{S}^N$. (Note that, if certain elements in $\mathbf{s}[q]$ are known perfectly as pilots, then the corresponding elements in $\nu^s[q]$ will be zero-valued.) Then, writing $\mathbf{s}[q] = \hat{\mathbf{s}}[q] + \check{\mathbf{s}}[q]$, where $\check{\mathbf{s}}[q]$ an unknown zero-mean deviation, the subcarrier observations $\mathbf{y}[q] = \mathcal{D}(\mathbf{s}[q])\Phi\mathbf{x}[q] + \mathbf{w}[q]$ can be written as

$$\mathbf{y}[q] = \mathcal{D}(\mathbf{s}[q])\Phi\mathbf{x}[q] + \mathbf{v}[q] \quad (35)$$

where $\mathbf{v}[q] \triangleq \mathcal{D}(\check{\mathbf{s}}[q])\Phi\mathbf{x}[q] + \mathbf{w}[q]$ is a zero-mean noise. Treating the elements within $\check{\mathbf{s}}[q]$ as uncorrelated and doing the same with $\mathbf{x}[q]$, and leveraging the fact that Φ is a truncated DFT matrix, it is straightforward to show that $\text{Cov}(\mathbf{v}[q]) = \mathcal{D}(\nu^v[q])$ with $\nu^v[q] = \nu^v\mathbf{1} + (\mathbf{1}^\top \boldsymbol{\rho})\nu^s[q]$, where $\boldsymbol{\rho}$ denotes the channel’s PDP. Without loss of generality, (35) can then be converted to the equivalent white-noise model

$$\mathbf{u}[q] \triangleq \mathcal{D}(\nu^v[q])^{-\frac{1}{2}}\mathbf{y}[q] = \mathbf{A}\mathbf{x}[q] + \mathbf{n}[q] \quad (36)$$

where $\text{Cov}(\mathbf{n}[q]) = \mathbf{I}$ and $\mathbf{A}[q] \triangleq \mathcal{D}(\nu^v[q])^{-\frac{1}{2}}\Phi$ is a known matrix. In summary, (36) provides a mechanism to handle soft inputs for both LASSO and LMMSE.

For LMMSE equalization, we first used (36) to compute

$$\hat{\mathbf{x}}_{\text{lmmse}}[q] = \mathcal{D}(\boldsymbol{\rho})\mathbf{A}^H[q](\mathbf{A}[q]\mathcal{D}(\boldsymbol{\rho})\mathbf{A}^H[q] + \mathbf{I})^{-1}\mathbf{u}[q] \quad (37)$$

from which we obtain the subcarrier gain estimate $\hat{z}_{\text{lmmse}}[q] = \Phi\hat{\mathbf{x}}_{\text{lmmse}}[q]$. The covariance matrix of $\hat{z}_{\text{lmmse}}[q]$ is [2]

$$\Phi(\mathcal{D}(\boldsymbol{\rho}) - \mathcal{D}(\boldsymbol{\rho})\mathbf{A}^H[q](\mathbf{A}[q]\mathcal{D}(\boldsymbol{\rho})\mathbf{A}^H[q] + \mathbf{I})^{-1}\mathbf{A}[q]\mathcal{D}(\boldsymbol{\rho}))\Phi^H$$

whose diagonal elements $\nu_{\text{lmmse}}^z[q]$ are variances on the gain estimates $\hat{z}_{\text{lmmse}}[q]$. Finally, we obtain soft symbol estimates from the soft gain estimates $(\hat{z}_{\text{lmmse}}[q], \nu_{\text{lmmse}}^z[q])$ via (32).

For LASSO,¹⁰ we first computed the tap estimate $\hat{\mathbf{x}}_{\text{lasso}}[q]$ from (36) using the celebrated SPGL1 algorithm [47]. In doing so, we needed to specify the target residual variance, i.e., $\nu_{\text{lasso}}^u \triangleq \|\mathbf{u}[q] - \mathbf{A}[q]\hat{\mathbf{x}}_{\text{lasso}}[q]\|_2^2/N$. Because $\text{Cov}(\mathbf{n}[q]) = \mathbf{I}$, we expect the optimal value of ν_{lasso}^u to be near 1 and, after extensive experimentation, we found that the value $\nu_{\text{lasso}}^u = 0.9$ works well at high SNR and that the value $\nu_{\text{lasso}}^u = 1.5$ works well at low SNR. Thus, for each $\mathbf{u}[q]$, we computed SPGL1 estimates using each of these two¹¹ targets, and kept the one that minimized the squared error $\nu_{\text{lasso}}^x[q] \triangleq \|\mathbf{x}[q] - \hat{\mathbf{x}}_{\text{lasso}}[q]\|_2^2/L$, which we assume a genie is able to provide. For the soft outputs, we set $\hat{z}_{\text{lasso}}[q] = \Phi\hat{\mathbf{x}}_{\text{lasso}}[q]$ and take $\nu_{\text{lasso}}^z[q]$ to be the diagonal elements of $\Phi\text{Cov}(\hat{\mathbf{x}}_{\text{lasso}}[q])\Phi^H$. Assuming

¹⁰The criterion employed by LASSO [9] is equivalent to the one employed in “basis pursuit denoising” [10].

¹¹We also tried running SPGL1 for a dense grid of ν_{lasso}^u values, but often it would get “stuck” at one of them and eventually return an error.




Fig. 5. BER versus number of pilot subcarriers N_p , for $E_b/N_o = 11$ dB, $M_t = 0$ training bits, $\eta = 2$ bpcu, and 16-QAM.

$\text{Cov}(\hat{\mathbf{x}}_{\text{lasso}}[q]) = \nu_{\text{lasso}}^x[q]\mathbf{I}$ and leveraging the fact that Φ is a truncated DFT matrix, we find $\nu_{\text{lasso}}^z[q] = L\nu_{\text{lasso}}^x[q]\mathbf{1}$. Finally, using (32), we obtain soft symbol estimates from the soft gain estimates ($\hat{\mathbf{z}}_{\text{lasso}}[q], \nu_{\text{lasso}}^z[q]$). Due to the genie-aided steps, the performance attained by our LASSO implementation is better than what could be obtained in practice.

These LMMSE- and LASSO-based SISO equalizers were then embedded in the overall factor graph in the same manner as GAMP, with the following exceptions. 1) The LMMSE and LASSO algorithms could not be connected to the MC sub-block, since they are not based on a two-state mixture model. 2) For LASSO, if the genie-aided MSE $\nu_{\text{lasso}}^x[q]$ did not improve during a given turbo iteration, then the corresponding outputs ($\hat{\mathbf{z}}_{\text{lasso}}[q], \nu_{\text{lasso}}^z[q]$) were not updated. This rule was employed to prevent turbo-LASSO from occasionally diverging at low SNR. 3) For LASSO, if $N_p > 0$ and $M_t = 0$, then the LASSO estimates computed during the first turbo iteration use only pilot subcarriers. This makes the performance of SISO-LASSO after the first turbo iteration equal to the performance of the standard pilot-aided LASSO.

C. BER Versus the Number of Pilot Subcarriers N_p

Fig. 5 shows bit error rate (BER) versus the number of pilot subcarriers N_p at $E_b/N_o = 11$ dB and a fixed spectral efficiency of $\eta = 2$ bpcu. In this and other figures, “ALG-#” refers to algorithm ALG with # turbo iterations (and “ALG-fin” after turbo convergence; see Fig. 10) with the MC block disconnected (i.e., there was no attempt to exploit tap clustering). Meanwhile “GAMP-# MC-5” refers to GAMP+MC after # turbo iterations, each containing 5 equalizer iterations. Finally, PCSI refers to MAP equalization under perfect CSI, which yields a bound on the BER performance of any equalizer.

The curves in Fig. 5 exhibit a “U” shape because, as N_p increases, the code rate R must decrease to maintain the fixed spectral efficiency $\eta = 2$ bpcu. While an increase in N_p generally makes channel estimation easier, the reduction in R makes data decoding more difficult. For all schemes under compar-




Fig. 6. BER versus number of interspersed training bits M_t , for $E_b/N_o = 10$ dB, $N_p = 0$ pilots subcarriers, $\eta = 2$ bpcu, and 16-QAM.

ison, Fig. 5 suggests that the choice $N_p \approx 224$ is optimal under the operating conditions. Overall, we see GAMP significantly outperforming both LMMSE and LASSO. Moreover, we see a small but noticeable gain from the MC block.

D. BER Versus the Number of Interspersed Training Bits M_t

Although $N_p > 0$ pilot subcarriers are required for decoupled channel estimation and decoding, JCED can function with $N_p = 0$ as long as a sufficient number M_t of training bits are interspersed among the coded bits used to construct each QAM symbol. To examine this latter case, Fig. 6 shows BER versus M_t at $E_b/N_o = 10$ dB, a fixed spectral efficiency of $\eta = 2$ bpcu, and $N_p = 0$. Again we see the “U” shape, but with GAMP working very well for a relatively wide range of M_t , and again we see a small but noticeable BER improvement when the MC block is used. SISO-LASSO seems to work to some degree with $N_p = 0$, but SISO-LASSO does not.

E. BER Versus E_b/N_o

Fig. 7 shows BER versus E_b/N_o using $N_p = 224$ pilot subcarriers (as suggested by Fig. 5) and $M_t = 0$ training bits. Relative to the perfect-CSI bound, we see SISO-LASSO performing within 5 dB during the first turbo iteration and within 4.5 dB after convergence. Meanwhile, we see SISO-LMMSE performing very poorly during the first turbo iteration, but eventually surpassing SISO-LASSO and coming within 4 dB from the perfect-CSI bound. Remarkably, we see GAMP + MC performing within 0.6 dB of the perfect-CSI bound (and within 1 dB after only 2 turbo iterations). This excellent performance confirms that the proposed GM2-HMM channel model and equalizer design together do an excellent job of capturing and exploiting the lag-dependent clustered-sparse characteristics of the 802.15.4a channel taps. Comparing the GAMP traces to the GAMP + MC traces, we see that the MC block yields a small but noticeable benefit.

Fig. 8 shows BER versus E_b/N_o using $M_t = 448$ interspersed training bits (as suggested by Fig. 6) and $N_p = 0$ pilot subcarriers. There we see that SISO-LASSO does not perform




Fig. 7. BER versus E_b/N_o , for $N_p = 224$ pilot subcarriers, $M_t = 0$ training bits, $\eta = 2$ bpcu, and 16-QAM.




Fig. 8. BER versus E_b/N_o , for $N_p = 0$ pilot subcarriers, $M_t = 448$ training bits, $\eta = 2$ bpcu, and 16-QAM.

well at all. SISO-LMMSE works to some degree after several turbo iterations, although not as well as in the $N_p > 0$ case. Meanwhile, we see GAMP+MC performing within 1 dB of the perfect-CSI case, and GAMP alone performing within 1.5 dB. Comparing Fig. 8 to Fig. 7, we see GAMP with training bits performing about 1 dB better than GAMP with dedicated pilot subcarriers. The perfect-CSI bound likewise improves because, with 16-QAM, $M_t = 448$ training bits constitutes half the overhead of $N_p = 224$ pilot subcarriers, allowing Fig. 8 the use of a stronger code at $\eta = 2$ bpcu.

F. Channel-Tap NMSE Versus E_b/N_o

Fig. 9 shows the channel estimates' normalized mean-squared error (NMSE) $E\{\|\mathbf{x}[q] - \hat{\mathbf{x}}[q]\|_2^2 / \|\mathbf{x}[q]\|_2^2}\}$ versus E_b/N_o , at the point that the turbo iterations were terminated, using $N_p = 224$ pilot subcarriers and $M_t = 0$ training bits. (For comparison, Fig. 7 shows BER for this configuration.) We also show the NMSE attained by the “bit and support genie”




Fig. 9. Channel tap NMSE versus E_b/N_o , for $N_p = 224$ pilot subcarriers, $M_t = 0$ training bits, $\eta = 2$ bpcu, and 16-QAM.




Fig. 10. Average time per turbo iteration (top), average number of turbo iterations (middle), and average total time (bottom), versus E_b/N_o , for $N_p = 224$ pilot subcarriers, $M_t = 0$ training bits, $\eta = 2$ bpcu, and 16-QAM.

(BSG), which calculates MMSE channel estimates using perfect knowledge of both the coded bits and the hidden channel states $\{d_j\}$, and which provides a lower bound for any channel estimator. In the figure, we see that the NMSEs of LMMSE and LASSO channel estimates are within 8–12 dB of the BSG, whereas those of GAMP are within 2–4 dB. Meanwhile, we see that GAMP + MC has a small but noticeable advantage over GAMP alone. We reason that the LMMSE estimates are worse than the GAMP estimates because they do not exploit the non-Gaussianity of the channel taps x_j , and the LASSO estimates are worse than the GAMP estimates because they do not exploit the known priors on the channel taps (i.e., the lag-dependent sparsity λ and PDP ρ).

G. Computational Complexity Versus E_b/N_o

Fig. 10 shows the average time per turbo iteration (in Matlab seconds on a 2.6 GHz CPU), the average number of turbo iterations, and the average total time (to turbo convergence),

as a function of E_b/N_o , using $N_p = 224$ pilot subcarriers and $M_t = 0$ training bits. (For comparison, Fig. 7 shows BER for this configuration and Fig. 9 shows NMSE.) Regarding the average time per turbo iteration, we see GAMP \pm MC taking ≈ 1.5 s at low E_b/N_o and ≈ 0.5 s at high E_b/N_o . GAMP + MC takes only slightly longer than GAMP alone due to the efficiency of the message computations within the MC block, and the fact that both the GAMP iterations and equalizer iterations are terminated as soon as the messages converge. In comparison, SISO-LMMSE takes ≈ 4.5 s per turbo iteration, and SISO-LASSO takes between 1 and 7 s, depending on E_b/N_o . Regarding the number of average number of turbo iterations until convergence, we see that—at low E_b/N_o —GAMP + MC takes about 5 turbo iterations, GAMP alone takes about 7, SISO-LMMSE takes about 5, and SISO-LASSO takes about 3, while—at high E_b/N_o —all algorithms converge after only 1 turbo iteration. Regarding the total time for equalization, GAMP + MC and GAMP are about the same at low E_b/N_o , whereas GAMP alone takes about 30% less time at high E_b/N_o . Meanwhile, SISO-LASSO and SISO-LMMSE are uniformly slower than GAMP and GAMP + MC over the entire E_b/N_o range, in some cases by a factor of 10.

V. CONCLUSION

In this paper, we presented a factor-graph approach to joint channel-estimation and decoding (JCED) for BICM-OFDM that merges recent advances in approximate message passing algorithms [13] with those in structured-sparse signal reconstruction [15] and SISO decoding [14]. Different from existing factor-graph approaches to JCED, ours is able to exploit not only sparse channel taps, but also clustered sparsity patterns that typify large-bandwidth communication channels, such as those that result from pulse-shaped communication over IEEE 802.15.4a modeled channels. For this purpose, we proposed the use of a two-state Gaussian mixture prior with a Markov model on the hidden tap states. The implementation complexity of our JCED scheme is dominated by $\mathcal{O}(N \log_2 N + N|S|)$ multiplies per GAMP iteration, facilitating the application to systems with many subcarriers N and many channel taps $L < N$. Experiments with IEEE 802.15.4a modeled channels showed BER performance within 1 dB of the known-channel bound, and 3–4 dB better than LMMSE- and LASSO-based soft equalizers. These experiments also suggested that, with our proposed approach, the use of interspersed training bits is more efficient than the use of dedicated pilot subcarriers. For very large constellations (e.g., in 1024-QAM, $|S| = 1024$), the algorithm becomes impractical; therefore, a topic of future research is how to further reduce the complexity.

APPENDIX A

DERIVATION OF GAMP FUNCTIONS $g_{out,i}$ AND $g'_{out,i}$

In this Appendix, we derive the GAMP quantities $g_{out,i}(y, \hat{z}, \nu^z)$ and $g'_{out,i}(y, \hat{z}, \nu^z)$ given in (21)–(26).

From (D1), we have that

$$E_{Z_i|Y_i}\{z | y; \hat{z}, \nu^z\} = \frac{1}{p_{Y_i}(y)} \int_z z p_{Y_i|Z_i}(y|z) \mathcal{CN}(z; \hat{z}, \nu^z) \quad (38)$$

where $p_{Y_i}(y) \triangleq \int_z p_{Y_i|Z_i}(y|z) \mathcal{CN}(z; \hat{z}, \nu^z)$. From (20), we rewrite $p_{Y_i|Z_i}(y|z)$ as

$$p_{Y_i|Z_i}(y|z) = \sum_{k=1}^{2^M} \frac{\beta_i^{(k)}}{s^{(k)}} \mathcal{CN}\left(z; \frac{y}{s^{(k)}}, \frac{\nu^w}{|s^{(k)}|^2}\right) \quad (39)$$

so that

$$\begin{aligned} & \int_z z p_{Y_i|Z_i}(y|z) \mathcal{CN}(z; \hat{z}, \nu^z) \\ &= \sum_{k=1}^{2^M} \frac{\beta_i^{(k)}}{s^{(k)}} \int_z z \mathcal{CN}\left(z; \frac{y}{s^{(k)}}, \frac{\nu^w}{|s^{(k)}|^2}\right) \mathcal{CN}(z; \hat{z}, \nu^z) \end{aligned} \quad (40)$$

$$\begin{aligned} p_{Y_i}(y) &= \sum_{k=1}^{2^M} \frac{\beta_i^{(k)}}{s^{(k)}} \int_z \mathcal{CN}\left(z; \frac{y}{s^{(k)}}, \frac{\nu^w}{|s^{(k)}|^2}\right) \mathcal{CN}(z; \hat{z}, \nu^z). \end{aligned} \quad (41)$$

Using the property that

$$\begin{aligned} & \mathcal{CN}(x; \hat{\theta}, \nu^\theta) \mathcal{CN}(x; \hat{\phi}, \nu^\phi) \\ &= \mathcal{CN}\left(x; \frac{\hat{\theta}/\nu^\theta + \hat{\phi}/\nu^\phi}{1/\nu^\theta + 1/\nu^\phi}, \frac{1}{1/\nu^\theta + 1/\nu^\phi}\right) \\ & \quad \times \mathcal{CN}(0; \hat{\theta} - \hat{\phi}, \nu^\theta + \nu^\phi) \end{aligned} \quad (42)$$

we can rewrite

$$\begin{aligned} & \int_z z p_{Y_i|Z_i}(y|z) \mathcal{CN}(z; \hat{z}, \nu^z) \\ &= \sum_{k=1}^{2^M} \frac{\beta_i^{(k)}}{s^{(k)}} \mathcal{CN}\left(0; \frac{y_i}{s} - \hat{z}, \frac{\nu^w}{|s^{(k)}|^2} + \nu^z\right) \\ & \quad \times \int_z z \mathcal{CN}\left(z; \frac{\frac{y}{s^{(k)}} \frac{|s^{(k)}|^2}{\nu^w} + \hat{z}}{\frac{|s^{(k)}|^2}{\nu^w} + \frac{1}{\nu^z}}, \frac{1}{\frac{|s^{(k)}|^2}{\nu^w} + \frac{1}{\nu^z}}\right) \end{aligned} \quad (43)$$

$$= \sum_{k=1}^{2^M} \frac{\beta_i^{(k)}}{s^{(k)}} \mathcal{CN}\left(\frac{y_i}{s}; \hat{z}, \frac{\nu^w}{|s^{(k)}|^2} + \nu^z\right) \frac{\frac{y}{s^{(k)}} \frac{|s^{(k)}|^2}{\nu^w} + \hat{z}}{\frac{|s^{(k)}|^2}{\nu^w} + \frac{1}{\nu^z}} \quad (44)$$

$$\begin{aligned} &= \sum_{k=1}^{2^M} \beta_i^{(k)} \mathcal{CN}\left(y_i; s^{(k)} \hat{z}, |s^{(k)}|^2 \nu^z + \nu^w\right) \\ & \quad \times \underbrace{\left(\left(\frac{y}{s^{(k)}} - \hat{z}\right) \frac{|s^{(k)}|^2 \nu^z}{|s^{(k)}|^2 \nu^z + \nu^w} + \hat{z}\right)}_{\triangleq \hat{e}^{(k)}(y, \hat{z}, \nu^z)} \end{aligned} \quad (45)$$

and, using the same procedure, we get

$$p_{Y_i}(y) = \sum_{k=1}^{2^M} \beta_i^{(k)} \mathcal{CN}\left(y_i; s^{(k)} \hat{z}, |s^{(k)}|^2 \nu^z + \nu^w\right). \quad (46)$$

With $\xi_i^{(k)}(y, \hat{z}, \nu^z)$ defined in (23), (38) and (45) and (46) combine to give

$$E_{Z_i|Y_i}\{z | y; \hat{z}, \nu^z\} = \sum_{k=1}^{2^M} \xi_i^{(k)}(y, \hat{z}, \nu^z) \left(\hat{e}^{(k)}(y, \hat{z}, \nu^z) + \hat{z}\right). \quad (47)$$

Finally, from (47) and the definition of $g_{\text{out},i}(y, \hat{z}, \nu^z)$ in (D2), (21) follows immediately.

From (D1), we have that

$$\text{var}_{Z_i|Y_i}\{z | y; \hat{z}, \nu^z\} = \frac{1}{p_{Y_i}(y)} \int_z |z - \mathbb{E}_{Z_i|Y_i}\{z | y; \hat{z}, \nu^z\}|^2 \times p_{Y_i|Z_i}(y|z) \mathcal{CN}(z; \hat{z}, \nu^z). \quad (48)$$

Similar to (43), we can write

$$\begin{aligned} & \int_z |z - \mathbb{E}_{Z_i|Y_i}\{z | y; \hat{z}, \nu^z\}|^2 p_{Y_i|Z_i}(y|z) \mathcal{CN}(z; \hat{z}, \nu^z) \\ &= \sum_{k=1}^{2^M} \frac{\beta_i^{(k)}}{s^{(k)}} \mathcal{CN}\left(0; \frac{y_i}{s} - \hat{z}, \frac{\nu^w}{|s^{(k)}|^2} + \nu^z\right) \\ & \times \int_z |z - \mathbb{E}_{Z_i|Y_i}\{z | y; \hat{z}, \nu^z\}|^2 \\ & \times \mathcal{CN}\left(z; \frac{\frac{y}{s^{(k)}} \frac{|s^{(k)}|^2}{\nu^w} + \frac{\hat{z}}{\nu^z}}{\frac{|s^{(k)}|^2}{\nu^w} + \frac{1}{\nu^z}}, \frac{1}{\frac{|s^{(k)}|^2}{\nu^w} + \frac{1}{\nu^z}}\right). \quad (49) \end{aligned}$$

Then, using the change-of-variable $\tilde{z} \triangleq z - \mathbb{E}_{Z_i|Y_i}\{z | y; \hat{z}, \nu^z\}$, and absorbing the $s^{(k)}$ terms as done in (45), we get

$$\begin{aligned} & \int_z |z - \mathbb{E}_{Z_i|Y_i}\{z | y; \hat{z}, \nu^z\}|^2 p_{Y_i|Z_i}(y|z) \mathcal{CN}(z; \hat{z}, \nu^z) \\ &= \sum_{k=1}^{2^M} \beta_i^{(k)} \mathcal{CN}\left(y_i; s^{(k)} \hat{z}, |s^{(k)}|^2 \nu^z + \nu^w\right) \\ & \times \int_{\tilde{z}} |\tilde{z}|^2 \mathcal{CN}\left(\tilde{z}; \hat{e}^{(k)} + \underbrace{\hat{z} - \mathbb{E}_{Z_i|Y_i}\{z | y; \hat{z}, \nu^z\}}_{=-\hat{e}_i}, \right. \\ & \left. \times \frac{\nu^w \nu^z}{|s^{(k)}|^2 \nu^z + \nu^w}\right) \quad (50) \end{aligned}$$

$$\begin{aligned} &= \sum_{k=1}^{2^M} \beta_i^{(k)} \mathcal{CN}\left(y_i; s^{(k)} \hat{z}, |s^{(k)}|^2 \nu^z + \nu^w\right) \\ & \times \left(\left| \hat{e}^{(k)} - \hat{e}_i \right|^2 + \frac{\nu^w \nu^z}{|s^{(k)}|^2 \nu^z + \nu^w} \right). \quad (51) \end{aligned}$$

Using $\xi_i^{(k)}(y, \hat{z}, \nu^z)$ defined in (23) and $\zeta^{(k)}(y, \hat{z}, \nu^z)$ defined in (24), (46) and (48) and (51) combine to give

$$\begin{aligned} \text{var}_{Z_i|Y_i}\{z | y; \hat{z}, \nu^z\} &= \sum_{k=1}^{2^M} \xi_i^{(k)}(y, \hat{z}, \nu^z) \\ & \times \left(\frac{\nu^w \zeta^{(k)}(y, \hat{z}, \nu^z)}{|s^{(k)}|^2} + \left| \hat{e}_i(y, \hat{z}, \nu^z) - \hat{e}^{(k)}(y, \hat{z}, \nu^z) \right|^2 \right) \quad (52) \end{aligned}$$

which is rewritten as $\nu_i^e(y, \hat{z}, \nu^z) \triangleq \text{var}_{Z_i|Y_i}\{z | y; \hat{z}, \nu^z\}$ in (27). Finally, plugging $\nu_i^e(y, \hat{z}, \nu^z)$ into the definition of $g'_{\text{out},i}(y, \hat{z}, \nu^z)$ in (D3), we immediately obtain (22).

APPENDIX B

DERIVATION OF GAMP FUNCTIONS $g_{\text{in},j}$ AND $g'_{\text{in},j}$

In this Appendix, we derive the GAMP quantities $g_{\text{in},j}(\hat{r}, \nu^r)$ and $g'_{\text{in},j}(\hat{r}, \nu^r)$ given in (28)–(30).

From (D4)–(D6), we note that $g_{\text{in},j}(\hat{r}, \nu^r)$ and $\nu^r g'_{\text{in},j}(\hat{r}, \nu^r)$ are the mean and variance, respectively, of the pdf

$$\frac{1}{Z_j} p_{X_j}(r) \mathcal{CN}(r; \hat{r}, \nu^r) \quad (53)$$

where $Z_j \triangleq \int_r p_{X_j}(r) \mathcal{CN}(r; \hat{r}, \nu^r)$. Using (42) together with the definition of $p_{X_j}(\cdot)$ from (6), we find

$$\begin{aligned} p_{X_j}(r) \mathcal{CN}(r; \hat{r}, \nu^r) &= \lambda_j \mathcal{CN}(r; 0, \nu_j^1) \mathcal{CN}(r; \hat{r}, \nu^r) \\ &+ (1 - \lambda_j) \mathcal{CN}(r; 0, \nu_j^0) \mathcal{CN}(r; \hat{r}, \nu^r) \quad (54) \end{aligned}$$

$$\begin{aligned} &= \lambda_j \mathcal{CN}(\hat{r}; 0, \nu_j^1 + \nu^r) \mathcal{CN}(r; \hat{r} \gamma_j^1(\nu^r), \nu^r \gamma_j^1(\nu^r)) \\ &+ (1 - \lambda_j) \mathcal{CN}(\hat{r}; 0, \nu_j^0 + \nu^r) \\ &\times \mathcal{CN}(r; \hat{r} \gamma_j^0(\nu^r), \nu^r \gamma_j^0(\nu^r)) \quad (55) \end{aligned}$$

for $\gamma_j^0(\nu^r) \triangleq (1 + \nu^r / \nu_j^0)^{-1}$ and $\gamma_j^1(\nu^r) \triangleq (1 + \nu^r / \nu_j^1)^{-1}$. This implies that

$$Z_j = \lambda_j \mathcal{CN}(\hat{r}; 0, \nu_j^1 + \nu^r) + (1 - \lambda_j) \mathcal{CN}(\hat{r}; 0, \nu_j^0 + \nu^r). \quad (56)$$

Thus, the mean obeys

$$\begin{aligned} g_{\text{in},j}(\hat{r}, \nu^r) &= \frac{1}{Z_j} \int_r r p_{X_j}(r) \mathcal{CN}(r; \hat{r}, \nu^r) \quad (57) \\ &= \frac{\lambda_j \mathcal{CN}(\hat{r}; 0, \nu_j^1 + \nu^r)}{Z_j} \gamma_j^1(\nu^r) \hat{r} \\ &= \underbrace{\alpha_j(\hat{r}, \nu^r)}_{= \alpha_j(\hat{r}, \nu^r)} \\ &+ \frac{(1 - \lambda_j) \mathcal{CN}(\hat{r}; 0, \nu_j^0 + \nu^r)}{Z_j} \gamma_j^0(\nu^r) \hat{r} \\ &= \underbrace{1 - \alpha_j(\hat{r}, \nu^r)}_{= 1 - \alpha_j(\hat{r}, \nu^r)} \quad (58) \end{aligned}$$

yielding (28), where a straightforward manipulation relates the expression for $\alpha_j(\hat{r}, \nu^r)$ above with its definition in (30).

Since, for the pdf in (53), $g_{\text{in},j}$ is the mean and $\nu^r g'_{\text{in},j}$ is the variance, we can write

$$\begin{aligned} \nu^r g'_{\text{in},j}(\hat{r}, \nu^r) &= \frac{1}{Z_j} \int_r |r|^2 p_{X_j}(r) \mathcal{CN}(r; \hat{r}, \nu^r) - |g_{\text{in},j}|^2 \quad (59) \\ &= \alpha_j \left(|\hat{r} \gamma_j^1|^2 + \nu^r \gamma_j^1 \right) + (1 - \alpha_j) \left(|\hat{r} \gamma_j^0|^2 + \nu^r \gamma_j^0 \right) \\ &\quad - |\alpha_j \gamma_j^1 \hat{r} + (1 - \alpha_j) \gamma_j^0 \hat{r}|^2 \quad (60) \end{aligned}$$

which can be simplified to yield (29).

REFERENCES

- [1] A. F. Molisch, *Wireless Communications*. New York: Wiley-IEEE Press, 2005.
- [2] H. V. Poor, *An Introduction to Signal Detection and Estimation*, 2nd ed. New York: Springer, 1994.
- [3] R. J.-M. Cramer, R. A. Scholtz, and M. Z. Win, "Evaluation of an ultra-wide-band propagation channel," *IEEE Trans. Antennas Propag.*, vol. 50, no. 5, pp. 561–570, May 2002.

- [4] J. C. Preisig and G. Deane, "Surface wave focusing and acoustic communications in the surf zone," *J. Acoust. Soc. Amer.*, vol. 116, pp. 2067–2080, Oct. 2004.
- [5] A. F. Molisch, "Ultrawideband propagation channels—Theory, measurement, and modeling," *IEEE Trans. Veh. Tech.*, vol. 54, no. 5, pp. 1528–1545, Sep. 2005.
- [6] N. Czink, X. Yin, H. Ozelik, M. Herdin, E. Bonek, and B. H. Fleury, "Cluster characteristics in a MIMO indoor propagation environment," *IEEE Trans. Wireless Commun.*, vol. 6, no. 4, pp. 1465–1475, Apr. 2007.
- [7] M. Medard, "The effect upon channel capacity in wireless communication of perfect and imperfect knowledge of the channel," *IEEE Trans. Inf. Theory*, vol. 46, no. 3, pp. 933–946, May 2000.
- [8] "Special Issue on Compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, Mar. 2008.
- [9] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. R. Statist. Soc. B*, vol. 58, no. 1, pp. 267–288, 1996.
- [10] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, 1998.
- [11] W. U. Bajwa, J. Haupt, A. M. Sayeed, and R. Nowak, "Compressed channel sensing: A new approach to estimating sparse multipath channels," *Proc. IEEE*, vol. 98, pp. 1058–1076, June 2010.
- [12] A. P. Kannu and P. Schniter, "On communication over unknown sparse frequency-selective block-fading channels," *IEEE Trans. Inf. Theory*, vol. 57, no. 10, pp. 6619–6632, Oct. 2011.
- [13] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," in *Proc. Int. Symp. Inf. Theory*, Oct. 2010, pp. 2168–2172.
- [14] D. J. C. MacKay, *Infatation Theory, Inference, and Learning Algorithms*. New York: Cambridge Univ. Press, 2003.
- [15] P. Schniter, "Turbo reconstruction of structured sparse signals," in *Proc. Conf. Inf. Sci. Syst.*, Princeton, NJ, Mar. 2010.
- [16] A. F. Molisch, K. Balakrishnan, C.-C. Chong, S. Emami, A. Fort, J. Karedal, J. Kunisch, H. Schantz, U. Schuster, and K. Siwiak, "IEEE 802.15.4a Channel Model—Final Report," Document IEEE 802.1504-0062-02-004a, 2005, Tech. Rep..
- [17] A. P. Worthen and W. E. Stark, "Unified design of iterative receivers using factor graphs," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 843–849, Feb. 2001.
- [18] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan Kaufman, 1988.
- [19] H.-A. Loeliger, J. Dauwels, J. Hu, S. Korl, L. Ping, and F. Kschischang, "The factor graph approach to model-based signal processing," *Proc. IEEE*, vol. 95, no. 6, pp. 1295–1322, Jun. 2007.
- [20] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 498–519, Feb. 2001.
- [21] J. Dauwels, S. Korl, and H.-A. Loeliger, "Expectation maximization as message passing," in *Proc. IEEE Int. Symp. Inf. Theory*, Adelaide, Australia, Sep. 2005, pp. 425–429.
- [22] J. Dauwels, S. Korl, and H.-A. Loeliger, "Particle methods as message passing," in *Proc. IEEE Int. Symp. Inf. Theory*, Seattle, WA, Jul. 2006, pp. 2052–2056.
- [23] J. Dauwels, "On variational message passing on factor graphs," in *Proc. IEEE Int. Symp. Inf. Theory*, Nice, France, Jun. 2007, pp. 2546–2550.
- [24] J. Dauwels, S. Korl, and H.-A. Loeliger, "Steepest descent as message passing," in *Proc. Inf. Theory Workshop*, Awaji, Japan, Oct. 2005, pp. 42–46.
- [25] G. F. Cooper, "The computational complexity of probabilistic inference using Bayesian belief networks," *Artif. Intell.*, vol. 42, pp. 393–405, 1990.
- [26] C. Novak, G. Matz, and F. Hlawatsch, "Factor graph based design of an OFDM-IDMA receiver performing joint data detection, channel estimation, and channel length selection," in *Proc. IEEE Int. Conf. Acoust. Speech, Signal Process.*, Taipei, Taiwan, Apr. 2009, pp. 2433–2436.
- [27] Y. Liu, L. Brunel, and J. Boutros, "Joint channel estimation and decoding using Gaussian approximation in a factor graph over multipath channel," in *Proc. IEEE Int. Symp. Pers. Indoor Mobile Radio Commun.*, Tokyo, Japan, Sep. 2009, pp. 3164–3168.
- [28] C. Kneivel, Z. Shi, P. A. Hoeher, and G. Auer, "2D graph-based soft channel estimation for MIMO-OFDM," in *Proc. IEEE Int. Conf. Commun.*, Cape Town, South Africa, Jul. 2010, pp. 1–5.
- [29] G. E. Kirkelund, C. N. Manchón, L. P. B. Christensen, E. Riegler, and B. H. Fleury, "Variational message-passing for joint channel estimation and decoding in MIMO-OFDM," in *Proc. IEEE Global Telecomm. Conf.*, Miami, FL, Dec. 2010, pp. 1–6.
- [30] P. Schniter, "Joint estimation and decoding for sparse channels via relaxed belief propagation," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Pacific Grove, CA, Nov. 2010.
- [31] P. Schniter, "Belief-propagation-based joint channel estimation and decoding for spectrally efficient communication over unknown sparse channels," *Phys. Commun.: Special Iss. in Compress. Sens. Commun.*, Oct. 2011, doi:10.1016/j.phycom.2011.09.004.
- [32] S. Rangan, "Estimation with random linear mixing, belief propagation and compressed sensing," in *Proc. 44th Annu. Conf. Inf. Sci. Syst.*, May 2010, pp. 1–6 [Online]. Available: arXiv:1001.2228v2
- [33] L. J. Cimini, Jr., "Analysis and simulation of a digital mobile radio channel using orthogonal frequency division multiplexing," *IEEE Trans. Commun.*, vol. 33, no. 7, pp. 665–765, Jul. 1985.
- [34] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York: Springer, 2007.
- [35] H. Ishwaran and J. S. Rao, "Spike and slab variable selection: Frequentist and Bayesian strategies," *Ann. Statist.*, vol. 33, pp. 730–773, Apr. 2005.
- [36] A. Saleh and R. A. Valenzuela, "A statistical model for indoor multipath propagation," *IEEE J. Sel. Areas Commun.*, vol. 5, no. 2, pp. 128–137, Feb. 1987.
- [37] B. J. Frey and D. J. C. MacKay, M. Jordan, M. S. Kearns, and S. A. Solla, Eds., "A revolution: Belief propagation in graphs with cycles," in *Advances in Neural Information Processing Systems*. Cambridge, MA: MIT Press, 1998.
- [38] W. T. Freeman, E. C. Pasztor, and O. T. Carmichael, "Learning low-level vision," *Intl. J. Computer Vision*, vol. 40, pp. 25–47, Oct. 2000.
- [39] D. Baron, S. Sarvotham, and R. G. Baraniuk, "Bayesian compressive sensing via belief propagation," *IEEE Trans. Signal Process.*, vol. 58, no. 1, pp. 2269–280, Jan. 2010.
- [40] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing," *Proc. Nat. Acad. Sci.*, vol. 106, pp. 18914–18919, Nov. 2009.
- [41] M. Bayati and A. Montanari, "The dynamics of message passing on dense graphs, with applications to compressed sensing," *IEEE Inf. Theory*, vol. 57, no. 2, pp. 764–785, Feb. 2011.
- [42] D. Guo and C.-C. Wang, "Random sparse linear systems observed via arbitrary channels: A decoupling principle," in *Proc. IEEE Int. Symp. Inf. Theory*, Nice, France, Jun. 2007, pp. 946–950.
- [43] T. J. Richardson and R. L. Urbanke, *Modern Coding Theory*. New York: Cambridge Univ. Press, 2009.
- [44] I. Kozintsev, Matlab Programs for Encoding and Decoding of LDPC Codes over GF(2^m), [Online]. Available: <http://www.kozintsev.net/soft.html>
- [45] Y. L. C. de Jong and T. J. Willink, "Iterative tree search detection for MIMO wireless systems," *IEEE Trans. Commun.*, vol. 53, no. 6, pp. 930–935, Jun. 2005.
- [46] M. Samuel, M. Barsoum, and M. P. Fitz, "On the suitability of Gray bit mappings to outer channel codes in iteratively decoded BICM," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Pacific Grove, CA, Nov. 2009, pp. 982–985.
- [47] E. van den Berg and M. P. Friedlander, "Probing the Pareto frontier for basis pursuit solutions," *SIAM J. Sci. Comput.*, vol. 31, no. 2, pp. 890–912, 2008.



Philip Schniter received the B.S. and M.S. degrees in electrical and computer engineering from the University of Illinois at Urbana-Champaign in 1992 and 1993, respectively, and the Ph.D. degree in electrical engineering from Cornell University in Ithaca, NY, in 2000.

From 1993 to 1996, he was with Tektronix, Inc., Beaverton, OR, as a Systems Engineer, and in 2000. Subsequently, he joined the Department of Electrical and Computer Engineering, The Ohio State University, Columbus, where he is now an

Associate Professor and a member of the Inflation Processing Systems (IPS) Lab. In 2008–2009, he was a Visiting Professor at Eurecom (Sophia Antipolis, France) and Supélec (Gif-sur-Yvette, France). His areas of interest include statistical signal processing, wireless communications and networks, and machine learning.

Prof. Schniter received the National Science Foundation CAREER Award in 2003.