Application of Bayesian Hierarchical Prior Modeling to Sparse Channel Estimation

Niels Lovmand Pedersen*, Carles Navarro Manchón*, Dmitriy Shutin† and Bernard Henri Fleury*

*Department of Electronic Systems, Aalborg University

Niels Jernes Vej 12, DK-9220 Aalborg, Denmark, Email: {nlp,cnm,bfl}@es.aau.dk

†Institute of Communications and Navigation, German Aerospace Center

Oberpfaffenhofen, D-82234 Wessling, Germany, Email: dmitriy.shutin@dlr.de

Abstract—Existing methods for sparse channel estimation typically provide an estimate computed as the solution maximizing an objective function defined as the sum of the log-likelihood function and a penalization term proportional to the ℓ_1-norm of the parameter of interest. However, other penalization terms have proven to have strong sparsity-inducing properties. In this work, we design pilot-assisted channel estimators for OFDM wireless receivers within the framework of sparse Bayesian learning by defining hierarchical Bayesian prior models that lead to sparsity-inducing penalization terms. The estimators result as an application of the variational message-passing algorithm on the factor graph representing the signal model extended with the hierarchical prior models. Numerical results demonstrate the superior performance of our channel estimators as compared to traditional and state-of-the-art sparse methods.

I. INTRODUCTION

During the last few years the research on compressive sensing techniques and sparse signal representations [1], [2] applied to channel estimation has received considerable attention, see e.g., [3]–[7]. The reason is that, typically, the impulse response of the wireless channel has a few dominant multipath components. A channel exhibiting this property is said to be sparse [3].

The general goal of sparse signal representations from overcomplete dictionaries is to estimate the sparse vector $\boldsymbol{\alpha}$ in the following system model:

$$\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{\alpha} + \boldsymbol{w}. \tag{1}$$

In this expression $\boldsymbol{y} \in \mathbb{C}^M$ is the vector of measurement samples and $\boldsymbol{w} \in \mathbb{C}^M$ represents the samples of the additive white Gaussian random noise with covariance matrix $\lambda^{-1}\boldsymbol{I}$ and precision parameter $\lambda > 0$. The matrix $\boldsymbol{\Phi} = [\boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_L] \in \mathbb{C}^{M \times L}$ is the overcomplete dictionary with more columns than rows ($L > M$) and $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_L]^{\mathrm{T}} \in \mathbb{C}^L$ is an unknown sparse vector, i.e., $\boldsymbol{\alpha}$ has few nonzero elements at unknown locations.

Often, a sparse channel estimator is constructed by solving the ℓ_1-norm constrained quadratic optimization problem, see among others [4]–[6]:

$$\widehat{\boldsymbol{\alpha}} = \underset{\boldsymbol{\alpha}}{\operatorname{argmin}} \left\{ \|\boldsymbol{y} - \boldsymbol{\Phi}\boldsymbol{\alpha}\|_2^2 + \kappa\|\boldsymbol{\alpha}\|_1 \right\} \tag{2}$$

with $\kappa > 0$ and $\|\cdot\|_p$, $p \geq 1$, denoting the ℓ_p vector norm. This method is also known as Least Absolute Shrinkage and Selection Operator (LASSO) regression [8] or Basis Pursuit Denoising [9]. The popularity of the LASSO regression is mainly attributed to the convexity of the cost function, as well as to its provable sparsity-inducing properties (see [2]). In [4]–[6] the LASSO regression is applied to *orthogonal frequency-division multiplexing* (OFDM) pilot-assisted channel estimation. Various channel estimation algorithms that minimize the LASSO cost function using convex optimization are compared in [6].

Another approach to sparse channel estimation is sparse Bayesian learning (SBL) [7], [10]–[12]. Specifically, SBL aims at finding a sparse *maximum a posteriori* (MAP) estimate of $\boldsymbol{\alpha}$

$$\widehat{\boldsymbol{\alpha}} = \underset{\boldsymbol{\alpha}}{\operatorname{argmin}} \left\{ \|\boldsymbol{y} - \boldsymbol{\Phi}\boldsymbol{\alpha}\|_2^2 + \lambda^{-1}Q(\boldsymbol{\alpha}) \right\} \tag{3}$$

by specifying a prior $p(\boldsymbol{\alpha})$ such that the penalty term $Q(\boldsymbol{\alpha}) \propto^e -\log p(\boldsymbol{\alpha})$ induces a sparse estimate $\widehat{\boldsymbol{\alpha}}$.[1]

Obviously, by comparing (2) and (3) the SBL framework realizes the LASSO cost function by choosing the Laplace prior $p(\boldsymbol{\alpha}) \propto \exp(-a\|\boldsymbol{\alpha}\|_1)$ with $\kappa = \lambda^{-1}a$. However, instead of working directly with the prior $p(\boldsymbol{\alpha})$, SBL models this using a two-layer (2-L) hierarchical structure. This involves specifying a conditional prior $p(\boldsymbol{\alpha}|\boldsymbol{\gamma})$ and a hyperprior $p(\boldsymbol{\gamma})$ such that $p(\boldsymbol{\alpha}) = \int p(\boldsymbol{\alpha}|\boldsymbol{\gamma})p(\boldsymbol{\gamma})\mathrm{d}\boldsymbol{\gamma}$ has a sparsity-inducing nature. The hierarchical approach to the representation of $p(\boldsymbol{\alpha})$ has several important advantages. First of all, one is free to choose simple and analytically tractable probability density functions (pdfs). Second, when carefully chosen, the resulting hierarchical structure allows for the construction of efficient yet computationally tractable iterative inference algorithms with analytical derivation of the inference expressions.

In [13] we propose a 2-L and a three-layer (3-L) prior model for $\boldsymbol{\alpha}$. These hierarchical prior models lead to novel

---

[1]Here $x \propto^e y$ denotes $\exp(x) = \exp(v)\exp(y)$, and thus $x = v + y$, for some arbitrary constant v. We will also make use of $x \propto y$ which denotes $x = vy$ for some positive constant v.

sparsity-inducing priors that include the Laplace prior for complex variables as a special case. This paper adapts the Bayesian probabilistic framework introduced in [13] to OFDM pilot-assisted sparse channel estimation. We then propose a variational message passing (VMP) algorithm that effectively exploits the hierarchical structure of the prior models. This approach leads to novel channel estimators that make use of various priors with strong sparsity-inducing properties. The numerical results reveal the promising potential of our estimators with improved performance as compared to state-of-the-art methods. In particular, the estimators outperform LASSO.

Throughout the paper we shall make use of the following notation: $(\cdot)^{\mathrm{T}}$ and $(\cdot)^{\mathrm{H}}$ denote respectively the transpose and the Hermitian transpose; the expression $\langle f(\boldsymbol{x})\rangle_{q(\boldsymbol{x})}$ denotes the expectation of the function $f(\boldsymbol{x})$ with respect to the density $q(\boldsymbol{x})$; $\mathrm{CN}(\boldsymbol{x}|\boldsymbol{a}, \boldsymbol{B})$ denotes a multivariate complex Gaussian pdf with mean $\boldsymbol{a}$ and covariance matrix $\boldsymbol{B}$; similarly, $\mathrm{Ga}(x|a, b) = \frac{b^a}{\Gamma(a)}x^{a-1}\exp(-bx)$ denotes a Gamma pdf with shape parameter a and rate parameter b.

II. SIGNAL MODEL

We consider a single-input single-output OFDM system with N subcarriers. A cyclic prefix (CP) is added to preserve orthogonality between subcarriers and to eliminate inter-symbol interference between consecutive OFDM symbols. The channel is assumed static during the transmission of each OFDM symbol. The received (baseband) OFDM signal $\boldsymbol{r} \in \mathbb{C}^N$ reads in matrix-vector notation

$$\boldsymbol{r} = \boldsymbol{X}\boldsymbol{h} + \boldsymbol{n}. \tag{4}$$

The diagonal matrix $\boldsymbol{X} = \mathrm{diag}(x_1, x_2, \ldots, x_N)$ contains the transmitted symbols. The components of the vector $\boldsymbol{h} \in \mathbb{C}^N$ are the samples of the channel frequency response at the N subcarriers. Finally, $\boldsymbol{n} \in \mathbb{C}^N$ is a zero-mean complex symmetric Gaussian random vector of independent components with variance λ^{-1}.

To estimate the vector $\boldsymbol{h}$ in (4), a total of M pilot symbols are transmitted at selected subcarriers. The pilot pattern $\mathcal{P} \subseteq \{1, \ldots, N\}$ denotes the set of indices of the pilot subcarriers. The received signals observed at the pilot positions $\boldsymbol{r}_{\mathcal{P}}$ are then divided each by the corresponding pilot symbol $\boldsymbol{X}_{\mathcal{P}} = \mathrm{diag}(x_n : n \in \mathcal{P})$ to produce the vector of observations:

$$\boldsymbol{y} \triangleq (\boldsymbol{X}_{\mathcal{P}})^{-1}\boldsymbol{r}_{\mathcal{P}} = \boldsymbol{h}_{\mathcal{P}} + (\boldsymbol{X}_{\mathcal{P}})^{-1}\boldsymbol{n}_{\mathcal{P}}. \tag{5}$$

We assume that all pilot symbols hold unit power such that the statistics of the noise term $(\boldsymbol{X}_{\mathcal{P}})^{-1}\boldsymbol{n}_{\mathcal{P}}$ remain unchanged, i.e., $\boldsymbol{y} \in \mathbb{C}^M$ yields the samples of the true channel frequency response (at the pilot subcarriers) corrupted by additive complex white Gaussian noise with component variance λ^{-1}.

In this work, we consider a frequency-selective wireless channel that remains constant during the transmission of each OFDM symbol. The maximum relative delay $\tau_{\max}$ is assumed to be large compared to the sampling time T_{s}, i.e., $\tau_{\max}/T_{\mathrm{s}} \gg 1$ [3]. The impulse response of the wireless channel

is modeled as a sum of multipath components:

$$g(\tau) = \sum_{k=1}^{K} \beta_k \delta\left(\tau - \tau_k\right). \tag{6}$$

In this expression, β_k and τ_k are respectively the complex weight and the continuous delay of the kth multipath component, and $\delta(\cdot)$ is the Dirac delta function. The parameter K is the total number of multipath components. The channel parameters K, β_k, and τ_k, $k = 1, \ldots, K$, are random variables. Specifically, the weights β_k, $k = 1, \ldots, K$, are mutually uncorrelated zero-mean with the sum of their variances normalized to one. Additional details regarding the assumptions on the model (6) are provided in Section VI.

III. THE DICTIONARY MATRIX

Our goal is to estimate $\boldsymbol{h}$ in (4) by applying the general optimization problem (3) to the observation model (5). For doing so, we must define a proper dictionary matrix $\boldsymbol{\Phi}$. In this section we give an example of such a matrix. As a starting point, we invoke the parametric model (6) of the channel. Making use of this model, (5) can be written as

$$\boldsymbol{y} = \boldsymbol{T}(\boldsymbol{\tau})\boldsymbol{\beta} + \boldsymbol{w} \tag{7}$$

with $\boldsymbol{h}_{\mathcal{P}} = \boldsymbol{T}(\boldsymbol{\tau})\boldsymbol{\beta}$, $\boldsymbol{w} = (\boldsymbol{X}_{\mathcal{P}})^{-1}\boldsymbol{n}_{\mathcal{P}}$, $\boldsymbol{\beta} = [\beta_1, \ldots, \beta_K]^{\mathrm{T}}$, $\boldsymbol{\tau} = [\tau_1, \ldots, \tau_K]^{\mathrm{T}}$, and $\boldsymbol{T}(\boldsymbol{\tau}) \in \mathbb{C}^{M \times K}$ depending on the pilot pattern $\mathcal{P}$ as well as the unknown delays in $\boldsymbol{\tau}$. Specifically, the (m, k)th entry of $\boldsymbol{T}(\boldsymbol{\tau})$ reads

$$\boldsymbol{T}(\boldsymbol{\tau})_{m,k} \triangleq \exp\left(-j2\pi f_m \tau_k\right), \quad \begin{array}{l} m = 1, 2, \ldots, M \\ k = 1, 2, \ldots, K \end{array} \tag{8}$$

with f_m denoting the frequency of the mth pilot subcarrier. In the general optimization problem (3) the columns of $\boldsymbol{\Phi}$ are known. However, the columns of $\boldsymbol{T}(\boldsymbol{\tau})$ in (7) depend on the unknown delays in $\boldsymbol{\tau}$. To circumvent this discrepancy we follow the same approach as in [5] and consider a grid of uniformly-spaced delay samples in the interval $[0, \tau_{\max}]$:

$$\boldsymbol{\tau}_d = \left[0, \frac{T_{\mathrm{s}}}{\zeta}, \frac{2T_{\mathrm{s}}}{\zeta}, \ldots, \tau_{\max}\right]^{\mathrm{T}} \tag{9}$$

with $\zeta > 0$ such that $\zeta\tau_{\max}/T_{\mathrm{s}}$ is an integer. We now define the dictionary $\boldsymbol{\Phi} \in \mathbb{C}^{M \times L}$ as $\boldsymbol{\Phi} = \boldsymbol{T}(\boldsymbol{\tau}_d)$. Thus, the entries of $\boldsymbol{\Phi}$ are of the form (8) with delay vector $\boldsymbol{\tau}_d$. The number of columns $L = \zeta\tau_{\max}/T_{\mathrm{s}} + 1$ in $\boldsymbol{\Phi}$ is thereby inversely proportional to the selected delay resolution T_{s}/ζ.

It is important to notice that the system model (1) with $\boldsymbol{\Phi}$ defined using discretized delay components is an approximation of the true system model (7). This approximation model is introduced so that (3) can be applied to solve the channel estimation task. The estimate of the channel vector at the pilot subcarriers is then $\widehat{\boldsymbol{h}}_{\mathcal{P}} = \boldsymbol{\Phi}\widehat{\boldsymbol{\alpha}}$. In order to estimate the channel $\boldsymbol{h}$ in (4) the dictionary $\boldsymbol{\Phi}$ is appropriately expanded (row-wise) to include all N subcarrier frequencies.

IV. BAYESIAN PRIOR MODELING

In this section we specify the joint pdf of the system model (1) when it is augmented with the 2-L and the 3-L hierarchical
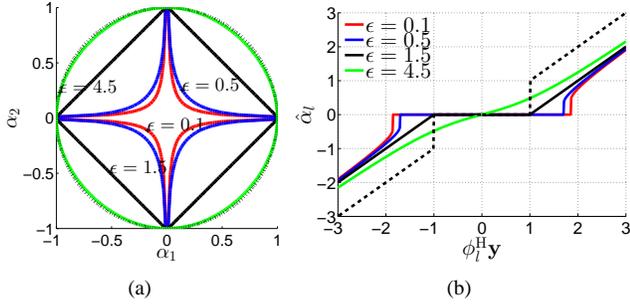
Fig. 1. 2-L hierarchical prior pdf for $\boldsymbol{\alpha} \in \mathbb{C}^2$: (a) Contour plot of the restriction to the $\text{Im}\{\alpha_1\} = \text{Im}\{\alpha_2\} = 0$ – plane of the penalty term $Q(\alpha_1, \alpha_2; \epsilon, \eta) \propto^e -\log(p(\alpha_1; \epsilon, \eta)p(\alpha_2; \epsilon, \eta))$. (b) Restriction to $\text{Im}\{\boldsymbol{\phi}_l^H \boldsymbol{y}\} = 0$ of the resulting MAP estimation rule (3) with ϵ as a parameter in the case when $\boldsymbol{\Phi}$ is orthonormal. The black dashed line indicates the hard-threshold rule and the black solid line the soft-threshold rule (obtained with $\epsilon = 3/2$). The black dashed line indicates the penalty term resulting when the prior pdf is a circular symmetric Gaussian pdf.

Fig. 2. Three-layer hierarchical prior pdf for $\boldsymbol{\alpha} \in \mathbb{C}^2$ with the setting $a = 1$, $b = 0.1$: (a) Restriction to $\text{Im}\{\boldsymbol{\phi}_l^H \boldsymbol{y}\} = 0$ of the resulting MAP estimation rule (3) with ϵ as a parameter in the case when $\boldsymbol{\Phi}$ is orthonormal. The black dashed line indicates the hard-threshold rule and the black solid line the soft-threshold rule. (b) Contour plot of the restriction to the $\text{Im}\{\alpha_1\} = \text{Im}\{\alpha_2\} = 0$ – plane of the penalty term $Q(\alpha_1, \alpha_2; \epsilon, a, b) \propto^e -\log(p(\alpha_1; \epsilon, a, b)p(\alpha_2; \epsilon, a, b))$.

prior model. The joint pdf of (1) augmented with the 2-L hierarchical prior model reads

$$p(\boldsymbol{y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \lambda) = p(\boldsymbol{y}|\boldsymbol{\alpha}, \lambda)p(\lambda)p(\boldsymbol{\alpha}|\boldsymbol{\gamma})p(\boldsymbol{\gamma}; \boldsymbol{\eta}). \quad (10)$$

The 3-L prior model considers the parameter $\boldsymbol{\eta}$ specifying the prior of $\boldsymbol{\gamma}$ in (10) as random. Thus, the joint pdf of (1) augmented with this hierarchical prior model is of the form

$$p(\boldsymbol{y}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\eta}, \lambda) = p(\boldsymbol{y}|\boldsymbol{\alpha}, \lambda)p(\lambda)p(\boldsymbol{\alpha}|\boldsymbol{\gamma})p(\boldsymbol{\gamma}|\boldsymbol{\eta})p(\boldsymbol{\eta}). \quad (11)$$

In (10) and (11) we have $p(\boldsymbol{y}|\boldsymbol{\alpha}, \lambda) = \text{CN}(\boldsymbol{y}|\boldsymbol{\Phi}\boldsymbol{\alpha}, \lambda^{-1}\boldsymbol{I})$ due to (1). Furthermore, we select the conjugate prior $p(\lambda) = p(\lambda; c, d) \triangleq \text{Ga}(\lambda|c, d)$. Finally, we let $p(\boldsymbol{\alpha}|\boldsymbol{\gamma}) = \prod_{l=1}^L p(\alpha_l|\gamma_l)$ with $p(\alpha_l|\gamma_l) \triangleq \text{CN}(\alpha_l|0, \gamma_l)$. In the following we show the main results and properties of these prior models. We refer to [13] for a more detailed analysis.

A. Two-Layer Hierarchical Prior Model

The 2-L prior model assumes that $p(\boldsymbol{\gamma}) = \prod_{l=1}^L p(\gamma_l)$ with $p(\gamma_l) = p(\gamma_l; \epsilon, \eta_l) \triangleq \text{Ga}(\gamma_l|\epsilon, \eta_l)$. We compute the prior of $\boldsymbol{\alpha}$ to be

$$p(\boldsymbol{\alpha}; \epsilon, \boldsymbol{\eta}) = \int_0^\infty p(\boldsymbol{\alpha}|\boldsymbol{\gamma})p(\boldsymbol{\gamma}; \epsilon, \boldsymbol{\eta})\mathrm{d}\boldsymbol{\gamma} = \prod_{l=1}^L p(\alpha_l; \epsilon, \eta_l) \quad (12)$$

with

$$p(\alpha_l; \epsilon, \eta_l) = \frac{2}{\pi\Gamma(\epsilon)}\eta_l^{\frac{(\epsilon+1)}{2}}|\alpha_l|^{\epsilon-1}K_{\epsilon-1}(2\sqrt{\eta_l}|\alpha_l|). \quad (13)$$

In this expression, $K_\nu(\cdot)$ is the modified Bessel function of the second kind with order $\nu \in \mathbb{R}$. The prior (13) leads to the general optimization problem (3) with penalty term

$$Q(\boldsymbol{\alpha}; \epsilon, \boldsymbol{\eta}) = \sum_{l=1}^L \log\left(|\alpha_l|^{\epsilon-1}K_{\epsilon-1}\left(2\sqrt{\eta_l}|\alpha_l|\right)\right). \quad (14)$$

We now show that the 2-L prior model induces the ℓ_1-norm penalty term and thereby the LASSO cost function as a special case. Selecting $\epsilon = 3/2$ and using the identity
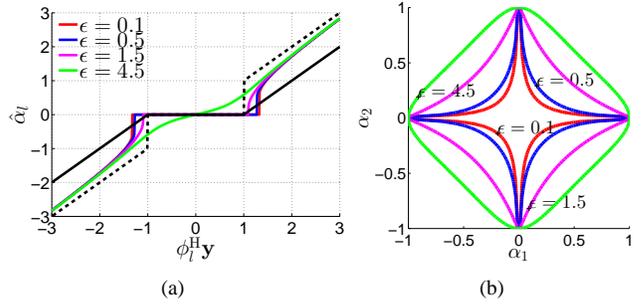
$K_{\frac{1}{2}}(z) = \sqrt{\frac{\pi}{2z}}\exp(-z)$ [14], (13) yields the Laplace prior

$$p(\alpha_l; \epsilon = 3/2, \eta_l) = \frac{2\eta_l}{\pi}\exp(-2\sqrt{\eta_l}|\alpha_l|). \quad (15)$$

With the selection $\eta_l = \eta$, $l = 1, \ldots, L$, we obtain $Q(\boldsymbol{\alpha}; \eta) = 2\sqrt{\eta}\|\boldsymbol{\alpha}\|_1$.

The prior pdf (13) is specified by ϵ and the regularization parameter $\boldsymbol{\eta}$. In order to get insight into the impact of ϵ on the properties of this prior pdf we consider the case $\boldsymbol{\alpha} \in \mathbb{C}^2$. In Fig. 1(a) the contour lines of the restriction to $\mathbb{R}^2$ of $Q(\alpha_1, \alpha_2; \epsilon, \eta) \propto^e -\log(p(\alpha_1; \epsilon, \eta)p(\alpha_2; \epsilon, \eta))$ are visualized;[2] each contour line is computed for a specific choice of ϵ. Notice that as ϵ decreases towards 0 more probability mass accumulates along the $\boldsymbol{\alpha}$-axes; as a consequence, the mode of the resulting posterior is more likely to be located close to the axes, thus promoting a sparse solution. The behavior of the classical ℓ_1 penalty term obtained for $\epsilon = 3/2$ can also be clearly recognized. In Fig. 1(b) we consider the case when $\boldsymbol{\Phi}$ is orthonormal and compute the MAP estimator (3) with penalty term (14) for different values of ϵ. Note the typical soft-threshold-like behavior of the estimators. As $\epsilon \to 0$, more components of $\widehat{\boldsymbol{\alpha}}$ are pulled towards zero since the threshold value increases, thus encouraging a sparser solution.

B. Three-Layer Hierarchical Prior Model

We now turn to the SBL problem with a 3-L prior model for $\boldsymbol{\alpha}$ leading to the joint pdf in (11). Specifically, the goal is to incorporate the regularization parameter $\boldsymbol{\eta}$ into the inference framework. To that end, we define $p(\boldsymbol{\eta}) = \prod_l^L p(\eta_l)$ with $p(\eta_l) = p(\eta_l; a_l, b_l) \triangleq \text{Ga}(\eta_l|a_l, b_l)$ and compute the prior $p(\boldsymbol{\alpha})$. Defining $\boldsymbol{a} \triangleq [a_1, \ldots, a_l]^\mathrm{T}$ and $\boldsymbol{b} \triangleq [b_1, \ldots, b_L]^\mathrm{T}$ we

---

[2]Let f denote a function defined on a set A. The restriction of f to a subset $B \subset A$ is the function defined on B that coincides with f on this subset.
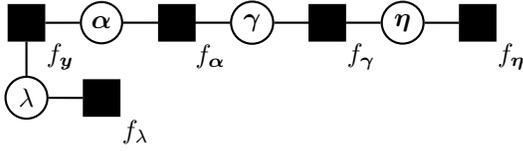
Fig. 3. A factor graph that represents the joint pdf (11). In this figure $f_{\boldsymbol{y}} \equiv p(\boldsymbol{y}|\boldsymbol{\alpha}, \lambda)$, $f_{\boldsymbol{\alpha}} \equiv p(\boldsymbol{\alpha}|\boldsymbol{\gamma})$, $f_{\boldsymbol{\gamma}} \equiv p(\boldsymbol{\gamma}|\boldsymbol{\eta})$, $f_{\boldsymbol{\eta}} \equiv p(\boldsymbol{\eta})$, and $f_{\lambda} \equiv p(\lambda)$.

obtain $p(\boldsymbol{\alpha}; \epsilon, \boldsymbol{a}, \boldsymbol{b}) = \prod_l^L p(\alpha_l; \epsilon, a_l, b_l)$ with

$$
\begin{aligned}
p(\alpha_l; \epsilon, a_l, b_l) &= \int_0^\infty p(\alpha_l|\gamma_l) p(\gamma_l) \mathrm{d}\gamma_l \\
&= \frac{\Gamma(\epsilon + a_l)\Gamma(a_l + 1)}{\pi b_l \Gamma(\epsilon)\Gamma(a_l)} \left(\frac{|\alpha_l|^2}{b_l}\right)^{\epsilon-1} U\left(\epsilon + a_l; \epsilon; \frac{|\alpha_l|^2}{b_l}\right).
\end{aligned}
\tag{16}
$$

In this expression, $U(\cdot; \cdot; \cdot)$ is the confluent hypergeometric function [14]. In Fig. 2(a) we show the estimation rules produced by the MAP solver for different values of ϵ and fixed parameters a_l and b_l when $\boldsymbol{\Phi}$ is orthonormal. It can be seen that the estimation rules obtained with the 3-L prior model approximate the hard-thresholding rule. In Fig. 2(b), we depict the contour lines of the restriction to $\mathbb{R}^2$ of $Q(\alpha_1, \alpha_2; \epsilon, a, b) \propto^e -\log(p(\alpha_1; \epsilon, a, b)p(\alpha_2; \epsilon, a, b))$. Observe that although the contours behave qualitatively similarly to those shown in Fig. 1(a) for the 2-L prior model, the estimation rules in Fig. 2(a) and Fig. 1(b) are different.

Naturally, the 3-L prior model encompasses three free parameters, ϵ, $\boldsymbol{a}$, and $\boldsymbol{b}$. The choice $\epsilon = 0$ and b_l small (practically we let $b_l = 10^{-6}$, $l = 1, \dots, L$) induces a weighted log-sum penalization term. This term is known to strongly promote a sparse estimate [10], [11]. Later in the text we will also adopt this parameter setting.

V. VARIATIONAL MESSAGE PASSING

In this section we present a VMP algorithm for estimating $\boldsymbol{h}$ in (4) given the observation $\boldsymbol{y}$ in (5). Let $\boldsymbol{\Theta} = \{\boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\eta}, \lambda\}$ be the set of unknown parameters and $p(\boldsymbol{y}, \boldsymbol{\Theta})$ be the joint pdf specified in (11). The factor graph [15] that encodes the factorization of $p(\boldsymbol{y}, \boldsymbol{\Theta})$ is shown in Fig. 3. Consider an auxiliary pdf $q(\boldsymbol{\Theta})$ for the unknown parameters that factorizes according to $q(\boldsymbol{\Theta}) = q(\boldsymbol{\alpha})q(\boldsymbol{\gamma})q(\boldsymbol{\eta})q(\lambda)$. The VMP algorithm is an iterative scheme that attempts to compute the auxiliary pdf that minimizes the Kullback-Leibler (KL) divergence $\mathrm{KL}(q(\boldsymbol{\Theta})\|p(\boldsymbol{\Theta}|\boldsymbol{y}))$. In the following we summarize the key steps of the algorithm; the reader is referred to [16] for more information on VMP.

From [16] the auxiliary function $q(\boldsymbol{\theta}_i)$, $\boldsymbol{\theta}_i \in \boldsymbol{\Theta}$, is updated as the product of incoming messages from the neighboring factor nodes f_n to the variable node $\boldsymbol{\theta}_i$:

$$
q(\boldsymbol{\theta}_i) \propto \prod_{f_n \in \mathcal{N}_{\boldsymbol{\theta}_i}} m_{f_n \to \boldsymbol{\theta}_i}.
\tag{17}
$$

In (17) $\mathcal{N}_{\boldsymbol{\theta}_i}$ is the set of factor nodes neighboring the variable node $\boldsymbol{\theta}_i$ and $m_{f_n \to \boldsymbol{\theta}_i}$ denotes the message from factor node f_n to variable node $\boldsymbol{\theta}_i$. This message is computed as

$$
m_{f_n \to \boldsymbol{\theta}_i} = \exp\left(\langle \ln f_n \rangle_{\prod_j q(\boldsymbol{\theta}_j),\, \boldsymbol{\theta}_j \in \mathcal{N}_{f_n} \backslash \{\boldsymbol{\theta}_i\}}\right),
\tag{18}
$$

where $\mathcal{N}_{f_n}$ is the set of variable nodes neighboring the factor node f_n. After an initialization procedure, the individual factors of $q(\boldsymbol{\Theta})$ are then updated iteratively in a round-robin fashion using (17) and (18).

We provide two versions of the VMP algorithm: one applied to the 2-L prior model (referred to as VMP-2L) and another one applied to the 3-L model (VMP-3L). The messages corresponding to VMP-2L are easily obtained as a special case of the messages computed for VMP-3L by assuming $q(\eta_l) = \delta(\eta_l - \hat{\eta}_l)$, where $\hat{\eta}_l$ is some fixed real number.

1) Update of $q(\boldsymbol{\alpha})$: According to (17) and Fig. 3 the computation of the update of $q(\boldsymbol{\alpha})$ requires evaluating the product of messages $m_{f_{\boldsymbol{y}} \to \boldsymbol{\alpha}}$ and $m_{f_{\boldsymbol{\alpha}} \to \boldsymbol{\alpha}}$. Multiplying these two messages yields the Gaussian auxiliary pdf $q(\boldsymbol{\alpha}) = \mathrm{CN}\left(\boldsymbol{\alpha}|\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\Sigma}}_{\boldsymbol{\alpha}}\right)$ with covariance matrix and mean given by

$$
\hat{\boldsymbol{\Sigma}}_{\boldsymbol{\alpha}} = (\langle \lambda \rangle_{q(\lambda)} \boldsymbol{\Phi}^{\mathrm{H}}\boldsymbol{\Phi} + \boldsymbol{V}(\boldsymbol{\gamma}))^{-1},
\tag{19}
$$

$$
\hat{\boldsymbol{\alpha}} = \langle \boldsymbol{\alpha} \rangle_{q(\boldsymbol{\alpha})} = \langle \lambda \rangle_{q(\lambda)} \hat{\boldsymbol{\Sigma}}_{\boldsymbol{\alpha}} \boldsymbol{\Phi}^{\mathrm{H}}\boldsymbol{y}.
\tag{20}
$$

In the above expression we have defined $\boldsymbol{V}(\boldsymbol{\gamma}) = \mathrm{diag}(\langle \gamma_1^{-1} \rangle_{q(\boldsymbol{\gamma})}, \dots, \langle \gamma_L^{-1} \rangle_{q(\boldsymbol{\gamma})})$.

2) Update of $q(\boldsymbol{\gamma})$: The update of $q(\boldsymbol{\gamma})$ is proportional to the product of the messages $m_{f_{\boldsymbol{\alpha}} \to \boldsymbol{\gamma}}$ and $m_{f_{\boldsymbol{\gamma}} \to \boldsymbol{\gamma}}$:

$$
q(\boldsymbol{\gamma}) \propto \prod_{l=1}^L \gamma_l^{\epsilon-2} \exp\left(-\gamma_l^{-1}\langle |\alpha_l|^2 \rangle_{q(\boldsymbol{\alpha})} - \gamma_l \langle \eta_l \rangle_{q(\boldsymbol{\eta})}\right).
\tag{21}
$$

The right-hand side expression in (21) is recognized as the product of Generalized Inverse Gaussian (GIG) pdfs [17] with order $p = \epsilon - 1$. Observe that the computation of $\boldsymbol{V}(\boldsymbol{\gamma})$ in (19) requires evaluating $\langle \gamma_l^{-1} \rangle_{q(\boldsymbol{\gamma})}$ for all $l = 1, \dots, L$. Luckily, the moments of the GIG distribution are given in closed form for any $n \in \mathbb{R}$ [17]:

$$
\langle \gamma_l^n \rangle_{q(\boldsymbol{\gamma})} = \left(\frac{\langle |\alpha_l|^2 \rangle_{q(\boldsymbol{\alpha})}}{\langle \eta_l \rangle_{q(\boldsymbol{\eta})}}\right)^{\frac{n}{2}} \frac{K_{p+n}\left(2\sqrt{\langle \eta_l \rangle_{q(\boldsymbol{\eta})}\langle |\alpha_l|^2 \rangle_{q(\boldsymbol{\alpha})}}\right)}{K_p\left(2\sqrt{\langle \eta_l \rangle_{q(\boldsymbol{\eta})}\langle |\alpha_l|^2 \rangle_{q(\boldsymbol{\alpha})}}\right)}.
\tag{22}
$$

3) Update of $q(\boldsymbol{\eta})$: The update of $q(\boldsymbol{\eta})$ is proportional to the product of messages $m_{f_{\boldsymbol{\eta}} \to \boldsymbol{\eta}}$ and $m_{f_{\boldsymbol{\gamma}} \to \boldsymbol{\eta}}$:

$$
q(\boldsymbol{\eta}) \propto \prod_{l=1}^L \eta_l^{\epsilon+a_l-1} \exp\left(-(\langle \gamma_l \rangle_{q(\boldsymbol{\gamma})} + b_l)\eta_l\right).
\tag{23}
$$

Clearly, $q(\boldsymbol{\eta})$ factorizes as a product of L gamma pdfs, one for each individual entry in $\boldsymbol{\eta}$. The first moment of η_l used in (22) is easily computed as

$$
\langle \eta_l \rangle_{q(\boldsymbol{\eta})} = \frac{\epsilon + a_l}{\langle \gamma_l \rangle_{q(\boldsymbol{\gamma})} + b_l}.
\tag{24}
$$

Naturally, $q(\boldsymbol{\eta})$ is only computed for VMP-3L.

4) Update of $q(\lambda)$: It can be shown that $q(\lambda) = \mathrm{Ga}(\lambda|M + c, \langle \|\boldsymbol{y} - \boldsymbol{\Phi}\boldsymbol{\alpha}\|_2^2 \rangle_{q(\boldsymbol{\alpha})} + d)$. The first moment of λ used in (19)
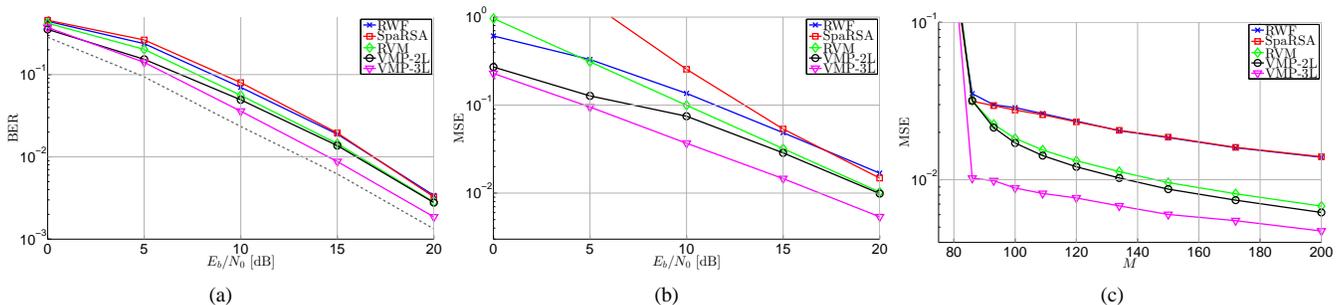
Fig. 4. Comparison of the performance of the VMP-2L, VMP-3L, RWF, RVM, and SparseRSA algorithms: (a) BER versus E_b/N_0, (b) MSE versus E_b/N_0, (c) MSE versus number of available pilots M with fixed $L = 200$ and the ratio between received symbol power and noise variance set to 15 dB. In (a,b) we have $M = 100$ and $L = 200$. In (a) the dashed line shows the BER performance when the true channel vector $\boldsymbol{h}$ in (4) is known.

TABLE I
PARAMETER SETTINGS FOR THE SIMULATIONS. THE CONVOLUTIONAL
CODE AND DECODER HAS BEEN IMPLEMENTED USING [18].

Sampling time, T_s	32.55 ns
CP length	4.69 μs / 144 T_s
Subcarrier spacing	15 kHz
Pilot pattern	Equally spaced, QPSK
Modulation	QPSK
Subcarriers, N	1200
Pilots, M	100
OFDM symbols	1
Information bits	727
Channel interleaver	Random
Convolutional code	$(133, 171, 165)_8$
Decoder	BCJR algorithm [19]

and (20) is therefore

$$\langle\lambda\rangle_{q(\lambda)} = \frac{M + c}{\langle\|\boldsymbol{y} - \boldsymbol{\Phi}\boldsymbol{\alpha}\|_2^2\rangle_{q(\boldsymbol{\alpha})} + d}. \quad (25)$$

VI. NUMERICAL RESULTS

We perform Monte Carlo simulations to evaluate the performance of the two versions of the derived VMP algorithm in Section V. We consider a scenario inspired by the 3GPP LTE standard [20] with the settings specified in Table I. The multipath channel (6) is based on the model used in [21] where, for each realization of the channel, the total number of multipath components K is Poisson distributed with mean of $\langle K\rangle_{p(K)} = 10$ and the delays τ_k, $k = 1, \ldots, K$, are independent and uniformly distributed random variables drawn from the continuous interval [0, 144 T_s] (corresponding to the CP length). The kth nonzero component β_k conditioned on the delay τ_k has a zero-mean complex circular symmetric Gaussian distribution with variance $\sigma^2(\tau_k) = \langle|\beta_k|^2\rangle_{p(\beta_k|\tau_k)} = u\exp(-\tau_k/v)$ and parameters $u, v > 0$.[3]

To initialize the VMP algorithm we set $\langle\lambda\rangle_{q(\lambda)}$ and $\langle\gamma_l^{-1}\rangle_{q(\boldsymbol{\gamma})}$ equal to the inverse of the sample variance of $\boldsymbol{y}$ and

the inverse number of columns L respectively. Furthermore, we let $c = d = 0$ in (25), which corresponds to the Jeffreys noninformative prior for λ. Once the initialization is completed, the algorithm sequentially updates the auxiliary pdfs $q(\boldsymbol{\alpha})$, $q(\boldsymbol{\gamma})$, $q(\boldsymbol{\eta})$, and $q(\lambda)$ until convergence is achieved. Obviously, $q(\boldsymbol{\eta})$ is only updated for VMP-3L, whereas for VMP-2L the entries in $\boldsymbol{\eta}$ are set to M. For both versions we select $\epsilon = 0$ and for VMP-3L we set $a_l = 1$ and $b_l = 10^{-6}$, $l = 1, \ldots, L$. Finally, the dictionary $\boldsymbol{\Phi}$ is specified by M pilot subcarriers and a total of $L = 200$ columns (corresponding to the choice $\tau_{\max} = 144\,T_s$ and $\zeta \approx 1.4$ in (9)).

The VMP is compared to a classical OFDM channel estimator and two state-of-the-art sparse estimation schemes. Specifically, we use as benchmark the robustly-designed Wiener Filter (RWF) [22], the relevance vector machine (RVM) [10], [11],[4] and the *sparse reconstruction by separable approximation* (SpaRSA) algorithm [23].[5] The RVM is an EM algorithm based on the 2-L prior model of the student-t pdf over each α_l, whereas SpaRSA is a proximal gradient method for solving (2). In case of the SpaRSA algorithm the regularization parameter κ needs to be set. In all simulations, we let $\kappa = 2$, which leads to good performance in high signal-to-noise ratio (SNR) regime.

The performance is compared with respect to the resulting bit-error-rate (BER) and mean-squared error (MSE) of the estimate $\widehat{\boldsymbol{h}}$ versus the SNR (E_b/N_0). In addition, in order to quantify the necessary pilot overhead, we evaluate the MSE versus the number of available pilots M. Hence, in this setup M is no longer fixed as in Table I.

In Fig. 4(a) we compare the BER performance of the different schemes. We see that VMP-3L outperforms the other schemes across all the SNR range considered. Specifically, at 1 % BER the gain is approximately 2 dB compared to VMP-2L and RVM and 3 dB compared to SpaRSA and RWF. Also VMP-2L achieves lower BER in the SNR range 0 - 12 dB compared to RVM and across the whole SNR range compared to SpaRSA and RWF.

The superior BER performance of the VMP algorithm is well reflected in the MSE performance shown in Fig. 4(b).

---

[3]The parameter u is computed such that $\langle\sum_{k=1}^{K}|\beta_k(t)|^2\rangle_{p(\boldsymbol{\beta},\boldsymbol{\tau},K)} = 1$, where $p(\boldsymbol{\beta}, \boldsymbol{\tau}, K)$ is the joint pdf of the parameters of the channel model. In the considered simulation scenario, $\langle K\rangle_{p(K)} = 10$, $\tau_{\max} = 144\,T_s$, and $v = 20\,T_s$ (the decay rate).

[4]The software is available on-line at http://dsp.ucsd.edu/~dwipf/.
[5]The software is available on-line at http://www.lx.it.pt/~mtf/SpaRSA/

Again VMP-3L is a clear winner followed by VMP-2L. The bad MSE performance of the SpaRSA for low SNR is due to the difficulty in specifying a suitable regularization parameter κ across a large SNR range.

We next fix the ratio between received symbol power and noise variance to 15 dB[6] and evaluate the MSE versus number of available pilots M. The results are depicted in Fig. 4(c). Observe a noticeable performance gain obtained with VMP-3L. In particular, VMP-3L exhibits the same MSE performance as VMP-2L and RVM using only approximately 85 pilots, roughly half as many as VMP-2L and RVM. Furthermore, VMP-3L, using this number of pilots, significantly outperforms SpaRSA and RWF using 200 pilots.

VII. CONCLUSION

In this paper, we proposed channel estimators based on sparse Bayesian learning. The estimators rely on Bayesian hierarchical prior modeling and variational message passing (VMP). The VMP algorithm effectively exploits the probabilistic structure of the hierarchical prior models and the resulting sparsity-inducing priors. Our numerical results show that the proposed channel estimators yield superior performance in terms of bit-error-rate and mean-squared error as compared to other existing estimators, including the estimator based on the ℓ_1-norm constraint. They also allow for a significant reduction of the amount of pilot subcarriers needed for estimating a given channel.

ACKNOWLEDGMENT

REFERENCES

[1] R. Baraniuk, "Compressive sensing," *IEEE Signal Processing Magazine*, vol. 24, no. 4, pp. 118–121, July 2007.

[2] E. J. Candes and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 21–30, Mar. 2008.

[3] W. Bajwa, J. Haupt, A. Sayeed, and R. Nowak, "Compressed channel sensing: A new approach to estimating sparse multipath channels," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 1058–1076, Jun. 2010.

[4] G. Taubock and F. Hlawatsch, "A compressed sensing technique for OFDM channel estimation in mobile environments: Exploiting channel sparsity for reducing pilots," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing ICASSP 2008*, 2008, pp. 2885–2888.

[5] C. R. Berger, S. Zhou, J. C. Preisig, and P. Willett, "Sparse channel estimation for multicarrier underwater acoustic communication: From subspace methods to compressed sensing," *IEEE Trans. on Sig. Proc.*, vol. 58, no. 3, pp. 1708–1721, 2010.

[6] J. Huang, C. R. Berger, S. Zhou, and J. Huang, "Comparison of basis pursuit algorithms for sparse channel estimation in underwater acoustic OFDM," in *Proc. OCEANS 2010 IEEE - Sydney*, 2010, pp. 1–6.

[7] D. Shutin and B. H. Fleury, "Sparse variational Bayesian SAGE algorithm with application to the estimation of multipath wireless channels," *IEEE Trans. on Sig. Proc.*, vol. 59, pp. 3609–3623, 2011.

[8] R. Tibshirani, "Regression shrinkage and selection via the LASSO," *J. R. Statist. Soc.*, vol. 58, pp. 267–288, 1994.

[9] S. S. Chen, D. L. Donoho, Michael, and A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Journal on Scientific Computing*, vol. 20, pp. 33–61, 1998.

[10] M. Tipping, "Sparse Bayesian learning and the relevance vector machine," *J. of Machine Learning Res.*, vol. 1, pp. 211–244, June 2001.

[11] D. Wipf and B. Rao, "Sparse Bayesian learning for basis selection," *IEEE Trans. on Sig. Proc.*, vol. 52, no. 8, pp. 2153 – 2164, aug. 2004.

[12] D. G. Tzikas, A. C. Likas, and N. P. Galatsanos, "The variational approximation for Bayesian inference," *IEEE Signal Process. Mag.*, vol. 25, no. 6, pp. 131–146, November 2008.

[13] N. L. Pedersen, D. Shutin, C. N. Manchón, and B. H. Fleury, "Sparse estimation using Bayesian hierarchical prior modeling for real and complex models," *submitted to IEEE Trans. on Sig. Proc.*, 2012, arXiv:1108.4324v1.

[14] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, 1972.

[15] F. R. Kschischang, B. J. Frey, and H. A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. on Information Theory*, vol. 47, no. 2, pp. 498–519, Feb 2001.

[16] J. Winn and C. M. Bishop, "Variational message passing," *J. Mach. Learn. Res.*, vol. 6, pp. 661–694, 2005.

[17] B. Jorgensen, *Statistical Properties of the Generalized Inverse Gaussian Distribution (Lecture Notes in Statistics 9)*. Springer-Verlag New York Inc, 1982.

[18] The iterative solutions coded modulation library. [Online]. Available: http://www.iterativesolutions.com

[19] L. Bahl, J. Cocke, F. Jelinek, and J. Raviv, "Optimal decoding of linear codes for minimizing symbol error rate," *IEEE Trans. on Inform. Theory*, vol. 20, no. 2, pp. 284–287, 1974.

[20] 3rd Generation Partnership Project (3GPP) Technical Specification, "Evolved universal terrestrial radio access (e-utra); base station (bs) radio transmission and reception," TS 36.104 V8.4.0, Tech. Rep., 2008.

[21] M. L. Jakobsen, K. Laugesen, C. Navarro Manchón, G. E. Kirkelund, C. Rom, and B. Fleury, "Parametric modeling and pilot-aided estimation of the wireless multipath channel in OFDM systems," in *Proc. IEEE Int Communications (ICC) Conf*, 2010, pp. 1–6.

[22] O. Edfors, M. Sandell, J.-J. van de Beek, S. K. Wilson, and P. O. Börjesson, "OFDM channel estimation by singular value decomposition," *IEEE Trans. on Communications*, vol. 46, no. 7, pp. 931–939, 1998.

[23] S. J. Wright, R. D. Nowak, and M. A. T. Figueiredo, "Sparse reconstruction by separable approximation," *IEEE Trans. on Sig. Proc.*, vol. 57, no. 7, pp. 2479–2493, 2009.

[6]Note that this value does not correspond with E_b/N_0 as represented in Fig. 4(a) and 4(b). The specific E_b/N_0 depends on the number of bits in an OFDM block, which in turn depends on the number of pilot symbols M.