

Efficient High-Dimensional Inference in the Multiple Measurement Vector Problem

Justin Ziniel *Student Member, IEEE* and Philip Schniter *Senior Member, IEEE*

Abstract

In this work, a Bayesian approximate message passing algorithm is proposed for solving the multiple measurement vector (MMV) problem in compressive sensing, in which a collection of sparse signal vectors that share a common support are recovered from undersampled noisy measurements. The algorithm, AMP-MMV, is capable of exploiting temporal correlations in the amplitudes of non-zero coefficients, and provides soft estimates of the signal vectors as well as the underlying support. Central to the proposed approach is an extension of recently developed approximate message passing techniques to the amplitude-correlated MMV setting. Aided by these techniques, AMP-MMV offers a computational complexity that is linear in all problem dimensions. In order to allow for automatic parameter tuning, an expectation-maximization algorithm that complements AMP-MMV is described. Finally, a detailed numerical study demonstrates the power of the proposed approach and its particular suitability for application to high-dimensional problems.

I. INTRODUCTION

As the field of compressive sensing (CS) [1]–[3] matures, researchers have begun to explore numerous extensions of the classical sparse signal recovery problem, in which a signal with few non-zero coefficients is reconstructed from a handful of incoherent linear measurements. One such extension, known as the *multiple measurement vector* (MMV) problem, generalizes the sparse signal recovery, or *single*

Copyright © 2012 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

The authors are with the Department of Electrical and Computer Engineering, The Ohio State University, Columbus, Ohio. E-mail: {zinielj, schniter}@ece.osu.edu.

Work supported in part by NSF grant CCF-1018368, DARPA/ONR grant N66001-10-1-4090, and an allocation of computing time from the Ohio Supercomputer Center. Portions of this work were previously presented at the 2011 Asilomar Conference on Signals, Systems, and Computers.

measurement vector (SMV), problem to the case where a group of measurement vectors has been obtained from a group of signal vectors that are assumed to be jointly sparse, sharing a common support. Such a problem has many applications, including magnetoencephalography [4], [5], direction-of-arrival estimation [6] and parallel magnetic resonance imaging (pMRI) [7].

Mathematically, given T length- M measurement vectors, the traditional MMV objective is to recover a collection of length- N sparse vectors $\{\mathbf{x}^{(t)}\}_{t=1}^T$, when $M < N$. Each measurement vector, $\mathbf{y}^{(t)}$, is obtained as

$$\mathbf{y}^{(t)} = \mathbf{A}\mathbf{x}^{(t)} + \mathbf{e}^{(t)}, \quad t = 1, \dots, T, \quad (1)$$

where \mathbf{A} is a known measurement matrix and $\mathbf{e}^{(t)}$ is corrupting additive noise. The unique feature of the MMV problem is the assumption of joint sparsity: the support of each sparse signal vector $\mathbf{x}^{(t)}$ is identical. Oftentimes, the collection of measurement vectors forms a time-series, thus we adopt a temporal viewpoint of the MMV problem, without loss of generality.

A straightforward approach to solving the MMV problem is to break it apart into independent SMV problems and apply one of the many SMV algorithms. While simple, this approach ignores valuable temporal structure in the signal that can be exploited to provide improved recovery performance. Indeed, under mild conditions, the probability of recovery failure can be made to decay exponentially as the number of timesteps T grows, when taking into account the joint sparsity [8].

Another approach (e.g., [9]) to the joint-sparse MMV problem is to restate (1) as the block-sparse SMV model

$$\bar{\mathbf{y}} = \mathcal{D}(\mathbf{A})\bar{\mathbf{x}} + \bar{\mathbf{e}}, \quad (2)$$

where $\bar{\mathbf{y}} \triangleq [\mathbf{y}^{(1)\top}, \dots, \mathbf{y}^{(T)\top}]^\top$, $\bar{\mathbf{x}} \triangleq [\mathbf{x}^{(1)\top}, \dots, \mathbf{x}^{(T)\top}]^\top$, $\bar{\mathbf{e}} \triangleq [\mathbf{e}^{(1)\top}, \dots, \mathbf{e}^{(T)\top}]^\top$, and $\mathcal{D}(\mathbf{A})$ denotes a block diagonal matrix consisting of T replicates of \mathbf{A} . In this case, \mathbf{x} is block-sparse, where the n^{th} block (for $n = 1, \dots, N$) consists of the coefficients $\{x_n, x_{n+N}, \dots, x_{n+(T-1)N}\}$. Equivalently, one could express (1) using the matrix model

$$\mathbf{Y} = \mathbf{A}\mathbf{X} + \mathbf{E}, \quad (3)$$

where $\mathbf{Y} \triangleq [\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(T)}]$, $\mathbf{X} \triangleq [\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(T)}]$, and $\mathbf{E} \triangleq [\mathbf{e}^{(1)}, \dots, \mathbf{e}^{(T)}]$. Under the matrix model, joint sparsity in (1) manifests as row-sparsity in \mathbf{X} . Algorithms developed for the matrix MMV problem are oftentimes intuitive extensions of SMV algorithms, and therefore share a similar taxonomy. Among the different techniques that have been proposed are mixed-norm minimization methods [5], [10]–[12], greedy pursuit methods [5], [13], [14], and Bayesian methods [6], [15]–[18]. Existing literature suggests

that greedy pursuit techniques are outperformed by mixed-norm minimization approaches, which in turn are surpassed by Bayesian methods [5], [15], [18].

In addition to work on the MMV problem, related work has been performed on a similar problem sometimes referred to as the “*dynamic CS*” problem [19]–[23]. The dynamic CS problem also shares the trait of working with multiple measurement vectors, but instead of joint sparsity, considers a situation in which the support of the signal changes slowly over time.

Given the plethora of available techniques for solving the MMV problem, it is natural to wonder what, if any, improvements can be made. In this work, we will primarily address two deficiencies evident in the available MMV literature. The first deficiency is the inability of many algorithms to account for amplitude correlations in the non-zero rows of \mathbf{X} .¹ Incorporating this temporal correlation structure is crucial, not only because many real-world signals possess such structure, but because the performance of MMV algorithms is particularly sensitive to this structure [8], [14], [15], [18], [24]. The second deficiency is that of computational complexity: while Bayesian MMV algorithms appear to offer the strongest recovery performance, it comes at the cost of increased complexity relative to simpler schemes, such as those based on greedy pursuit. For high-dimensional datasets, the complexity of Bayesian techniques may prohibit their application.

Our goal is to develop an MMV algorithm that offers the best of both worlds, combining the recovery performance of Bayesian techniques, even in the presence of substantial amplitude correlation and apriori unknown signal statistics, with the linear complexity scaling of greedy pursuit methods. Aiding us in meeting our goal is a powerful algorithmic framework known as *approximate message passing* (AMP), first proposed by Donoho et al. for the SMV CS problem [25]. In its early SMV formulations, AMP was shown to perform rapid and highly accurate probabilistic inference on models with known i.i.d. signal and noise priors, and i.i.d. random \mathbf{A} matrices [25], [26]. More recently, AMP was extended to the block-sparse SMV problem under similar conditions [27]. While this block-sparse SMV AMP does solve a simple version of the MMV problem via the formulation (2), it does not account for intra-block amplitude correlation (i.e., temporal correlation in the MMV model). Recently, Kim et al. proposed an AMP-based MMV algorithm that does exploit temporal amplitude correlation [16]. However, their approach requires knowledge of the signal and noise statistics (e.g., sparsity, power, correlation) and uses matrix inversions at each iteration, implying a complexity that grows superlinearly in the problem dimensions.

In this work, we propose an AMP-based MMV algorithm (henceforth referred to as AMP-MMV) that

¹Notable exceptions include [16], [12], and [18], which explicitly model amplitude correlations.

exploits temporal amplitude correlation and learns the signal and noise statistics directly from the data, all while maintaining a computational complexity that grows linearly in the problem dimensions. In addition, AMP-MMV can easily accommodate time-varying measurement matrices $\mathbf{A}^{(t)}$, implicit measurement operators (e.g., FFT-based \mathbf{A}), and complex-valued quantities. (These latter scenarios occur in, e.g., digital communication [28] and pMRI [29].) The key to our approach lies in combining the “turbo AMP” framework of [30], where the usual AMP factor graph is augmented with additional hidden variable nodes and inference is performed on the augmented factor graph, with an EM-based approach to hyperparameter learning. Details are provided in Sections II, IV, and V.

In Section VI, we present a detailed numerical study of AMP-MMV that includes a comparison against three state-of-the-art MMV algorithms. In order to establish an absolute performance benchmark, in Section III we describe a tight, oracle-aided performance lower bound that is realized through a support-aware Kalman smoother (SKS). To the best of our knowledge, our numerical study is the first in the MMV literature to use the SKS as a benchmark. Our numerical study demonstrates that AMP-MMV performs near this oracle performance bound under a wide range of problem settings, and that AMP-MMV is especially suitable for application to high-dimensional problems. In what represents a less-explored direction for the MMV problem, we also explore the effects of measurement matrix time-variation (cf. [6]). Our results show that measurement matrix time-variation can significantly improve reconstruction performance and thus we advocate the use of time-varying measurement operators whenever possible.

Notation: Boldfaced lower-case letters, e.g., \mathbf{a} , denote vectors, while boldfaced upper-case letters, e.g., \mathbf{A} , denote matrices. The letter t is strictly used to index a timestep, $t = 1, 2, \dots, T$, the letter n is strictly used to index the coefficients of a signal, $n = 1, \dots, N$, and the letter m is strictly used to index the measurements, $m = 1, \dots, M$. The superscript (t) indicates a timestep-dependent quantity, while a superscript without parentheses, e.g., k , indicates a quantity whose value changes according to some algorithmic iteration index. Subscripted variables such as $x_n^{(t)}$ are used to denote the n^{th} element of the vector $\mathbf{x}^{(t)}$. The m^{th} row of the matrix \mathbf{A} is denoted by \mathbf{a}_m^\top , and the transpose (conjugate transpose) by \mathbf{A}^\top (\mathbf{A}^H). An M -by- M identity matrix is denoted by \mathbf{I}_M , a length- N vector of ones is given by $\mathbf{1}_N$ and $\mathcal{D}(\mathbf{a})$ designates a diagonal matrix whose diagonal entries are given by the elements of the vector \mathbf{a} . Finally, $\mathcal{CN}(\mathbf{a}; \mathbf{b}, \mathbf{C})$ refers to the complex normal distribution that is a function of the vector \mathbf{a} , with mean \mathbf{b} and covariance matrix \mathbf{C} .

II. SIGNAL MODEL

In this section, we elaborate on the signal model outlined in Section I, and make precise our modeling assumptions. Our signal model, as well as our algorithm, will be presented in the context of complex-valued signals, but can be easily modified to accommodate real-valued signals.

As noted in Section I, we consider the linear measurement model (1), in which the signal $\mathbf{x}^{(t)} \in \mathbb{C}^N$ at timestep t is observed as $\mathbf{y}^{(t)} \in \mathbb{C}^M$ through the linear operator $\mathbf{A} \in \mathbb{C}^{M \times N}$. We assume $\mathbf{e}^{(t)} \sim \mathcal{CN}(\mathbf{0}, \sigma_e^2 \mathbf{I}_M)$ is circularly symmetric complex white Gaussian noise. We use $\mathcal{S} \triangleq \{n | x_n^{(t)} \neq 0\}$ to denote the indices of the time-invariant support of the signal, which is assumed to be suitably sparse, i.e., $|\mathcal{S}| \leq M$.²

Our approach to specifying a prior distribution for the signal, $p(\{\mathbf{x}^{(t)}\}_{t=1}^T)$, is motivated by a desire to separate the support, \mathcal{S} , from the amplitudes of the non-zero, or ‘‘active,’’ coefficients. To accomplish this, we decompose each coefficient $x_n^{(t)}$ as the product of two hidden variables:

$$x_n^{(t)} = s_n \cdot \theta_n^{(t)} \quad \Leftrightarrow \quad p(x_n^{(t)} | s_n, \theta_n^{(t)}) = \begin{cases} \delta(x_n^{(t)} - \theta_n^{(t)}), & s_n = 1, \\ \delta(x_n^{(t)}), & s_n = 0, \end{cases} \quad (4)$$

where $s_n \in \{0, 1\}$ is a binary variable that indicates support set membership, $\theta_n^{(t)} \in \mathbb{C}$ is a variable that provides the amplitude of coefficient $x_n^{(t)}$, and $\delta(\cdot)$ is the Dirac delta function. When $s_n = 0$, $x_n^{(t)} = 0$ and $n \notin \mathcal{S}$, and when $s_n = 1$, $x_n^{(t)} = \theta_n^{(t)}$ and $n \in \mathcal{S}$. To model the sparsity of the signal, we treat each s_n as a Bernoulli random variable with $\Pr\{s_n = 1\} \triangleq \lambda_n < 1$.

In order to model the temporal correlation of signal amplitudes, we treat the evolution of amplitudes over time as stationary first-order Gauss-Markov random processes. Specifically, we assume that $\theta_n^{(t)}$ evolves according to the following linear dynamical system model:

$$\theta_n^{(t)} = (1 - \alpha)(\theta_n^{(t-1)} - \zeta) + \alpha w_n^{(t)} + \zeta, \quad (5)$$

where $\zeta \in \mathbb{C}$ is the mean of the amplitude process, $w_n^{(t)} \sim \mathcal{CN}(0, \rho)$ is a circularly symmetric white Gaussian perturbation process, and $\alpha \in [0, 1]$ is a scalar that controls the correlation of $\theta_n^{(t)}$ across time. At one extreme, $\alpha = 0$, the random process is perfectly correlated ($\theta_n^{(t)} = \theta_n^{(t-1)}$), while at the other extreme, $\alpha = 1$, the amplitudes evolve independently over time. Note that the binary support vector, \mathbf{s} , is independent of the amplitude random process, $\{\boldsymbol{\theta}^{(t)}\}_{t=1}^T$, which implies that there are hidden amplitude

²If the signal being recovered is not itself sparse, it is assumed that there exists a known basis, incoherent with the measurement matrix, in which the signal possesses a sparse representation. Without loss of generality, we will assume the underlying signal is sparse in the canonical basis.

“trajectories”, $\{\theta_n^{(t)}\}_{t=1}^T$, associated with inactive coefficients. Consequently, $\theta_n^{(t)}$ should be thought of as the conditional amplitude of $x_n^{(t)}$, conditioned on $s_n = 1$.

Under our model, the prior distribution of any signal coefficient, $x_n^{(t)}$, is a Bernoulli-Gaussian or “spike-and-slab” distribution:

$$p(x_n^{(t)}) = (1 - \lambda_n)\delta(x_n^{(t)}) + \lambda_n\mathcal{CN}(x_n^{(t)}; \zeta, \sigma^2), \quad (6)$$

where $\sigma^2 \triangleq \frac{\alpha\rho}{2-\alpha}$ is the steady-state variance of $\theta_n^{(t)}$. We note that when $\lambda_n < 1$, (6) is an effective sparsity-promoting prior due to the point mass at $x_n^{(t)} = 0$.

III. THE SUPPORT-AWARE KALMAN SMOOTHER

Prior to describing AMP-MMV in detail, we first motivate the type of inference we wish to perform. Suppose for a moment that we are interested in obtaining a minimum mean square error (MMSE) estimate of $\{\mathbf{x}^{(t)}\}_{t=1}^T$, and that we have access to an oracle who can provide us with the support, \mathcal{S} . With this knowledge, we can concentrate solely on estimating $\{\boldsymbol{\theta}^{(t)}\}_{t=1}^T$, since, conditioned on \mathcal{S} , an MMSE estimate of $\{\boldsymbol{\theta}^{(t)}\}_{t=1}^T$ can provide an MMSE estimate of $\{\mathbf{x}^{(t)}\}_{t=1}^T$. For the linear dynamical system of (5), the support-aware Kalman smoother (SKS) provides the appropriate oracle-aided MMSE estimator of $\{\boldsymbol{\theta}^{(t)}\}_{t=1}^T$ [31]. The state-space model used by the SKS is:

$$\boldsymbol{\theta}^{(t)} = (1 - \alpha)\boldsymbol{\theta}^{(t-1)} + \alpha\zeta\mathbf{1}_N + \alpha\mathbf{w}^{(t)}, \quad (7)$$

$$\mathbf{y}^{(t)} = \mathbf{A}\mathcal{D}(\mathbf{s})\boldsymbol{\theta}^{(t)} + \mathbf{e}^{(t)}, \quad (8)$$

where \mathbf{s} is the binary support vector associated with \mathcal{S} . If $\hat{\boldsymbol{\theta}}^{(t)}$ is the MMSE estimate returned by the SKS, then an MMSE estimate of $\mathbf{x}^{(t)}$ is given by $\hat{\mathbf{x}}^{(t)} = \mathcal{D}(\mathbf{s})\hat{\boldsymbol{\theta}}^{(t)}$.

The state-space model (7), (8) provides a useful interpretation of our signal model. In the context of Kalman smoothing, the state vector $\boldsymbol{\theta}^{(t)}$ is only partially observable (due to the action of $\mathcal{D}(\mathbf{s})$ in (8)). Since $\mathcal{D}(\mathbf{s})\boldsymbol{\theta}^{(t)} = \mathbf{x}^{(t)}$, noisy linear measurements of $\mathbf{x}^{(t)}$ are used to infer the state $\boldsymbol{\theta}^{(t)}$. However, since only those $\theta_n^{(t)}$ for which $n \in \mathcal{S}$ are observable, and thus identifiable, they are the only ones whose posterior distributions will be meaningful.

Since the SKS performs optimal MMSE estimation, given knowledge of the true signal support, it provides a useful lower bound on the achievable performance of any support-agnostic Bayesian algorithm that aims to perform MMSE estimation of $\{\mathbf{x}^{(t)}\}_{t=1}^T$.




Fig. 1: Factor graph representation of the $p(\bar{\mathbf{x}}, \bar{\boldsymbol{\theta}}, \mathbf{s}|\bar{\mathbf{y}})$ decomposition in (9).

IV. THE AMP-MMV ALGORITHM

In Section II, we decomposed each signal coefficient, $x_n^{(t)}$, as the product of a binary support variable, s_n , and an amplitude variable, $\theta_n^{(t)}$. We now develop an algorithm that infers a marginal posterior distribution on each variable, enabling both soft estimation and soft support detection.

The statistical structure of the signal model from Section II becomes apparent from a factorization of the posterior joint pdf of all random variables. Recalling from (2) the definitions of $\bar{\mathbf{y}}$ and $\bar{\mathbf{x}}$, and defining $\bar{\boldsymbol{\theta}}$ similarly, the posterior joint distribution factors as follows:

$$p(\bar{\mathbf{x}}, \bar{\boldsymbol{\theta}}, \mathbf{s}|\bar{\mathbf{y}}) \propto \prod_{t=1}^T \left(\prod_{m=1}^M p(y_m^{(t)}|\mathbf{x}^{(t)}) \prod_{n=1}^N p(x_n^{(t)}|\theta_n^{(t)}, s_n) p(\theta_n^{(t)}|\theta_n^{(t-1)}) \right) \prod_{n=1}^N p(s_n), \quad (9)$$

where \propto indicates equality up to a normalizing constant, and $p(\theta_n^{(1)}|\theta_n^{(0)}) \triangleq p(\theta_n^{(1)})$. A convenient graphical representation of this decomposition is given by a *factor graph* [32], which is an undirected bipartite graph that connects the pdf “factors” of (9) with the variables that make up their arguments. The factor graph for the decomposition of (9) is shown in Fig. 1. The *factor nodes* are denoted by filled squares, while the *variable nodes* are denoted by circles. In the figure, the signal variable nodes at timestep t , $\{x_n^{(t)}\}_{n=1}^N$, are depicted as lying in a plane, or “frame”, with successive frames stacked one after another. Since during inference the measurements $\{y_m^{(t)}\}$ are known observations and not random variables, they do not appear explicitly in the factor graph. The connection between the frames occurs through the amplitude and support indicator variables, providing a graphical representation of the temporal correlation in the signal. For visual clarity, these $\{\theta_n^{(t)}\}_{t=1}^T$ and s_n variable nodes have been removed from the graph for the intermediate index n , but should in fact be present at every index $n = 1, \dots, N$.

Factor	Distribution	Functional Form
$g_m^{(t)}(\mathbf{x}^{(t)})$	$p(y_m^{(t)} \mathbf{x}^{(t)})$	$\mathcal{CN}(y_m^{(t)}; \mathbf{a}_m^T \mathbf{x}^{(t)}, \sigma_e^2)$
$f_n^{(t)}(x_n^{(t)}, s_n, \theta_n^{(t)})$	$p(x_n^{(t)} s_n, \theta_n^{(t)})$	$\delta(x_n^{(t)} - s_n \theta_n^{(t)})$
$h_n(s_n)$	$p(s_n)$	$(1 - \lambda_n)^{(1-s_n)} (\lambda_n)^{s_n}$
$d_n^{(1)}(\theta_n^{(1)})$	$p(\theta_n^{(1)})$	$\mathcal{CN}(\theta_n^{(1)}; \zeta, \sigma^2)$
$d_n^{(t)}(\theta_n^{(t)}, \theta_n^{(t-1)})$	$p(\theta_n^{(t)} \theta_n^{(t-1)})$	$\mathcal{CN}(\theta_n^{(t)}; (1 - \alpha)\theta_n^{(t-1)} + \alpha\zeta, \alpha^2\rho)$

TABLE I: The factors, underlying distributions, and functional forms associated with the signal model of Section II.

The factor nodes in Fig. 1 have all been assigned alphabetic labels; the correspondence between these labels and the distributions they represent, as well as the functional form of each distribution, is presented in Table I.

A natural approach to performing statistical inference on a signal model that possesses a convenient factor graph representation is through a message passing algorithm known as belief propagation [33]. In belief propagation, the messages exchanged between connected nodes of the graph represent probability distributions. In cycle-free graphs, belief propagation can be viewed as an instance of the sum-product algorithm [32], allowing one to obtain an exact posterior marginal distribution for each unobserved variable, given a collection of observed variables. When the factor graph contains cycles, the same rules that define the sum-product algorithm can still be applied, however convergence is no longer guaranteed [32]. Despite this, there exist many problems to which loopy belief propagation [34] has been successfully applied, including inference on Markov random fields [35], LDPC decoding [36], and compressed sensing [25], [30], [37]–[40].

We now proceed with a high-level description of AMP-MMV, an algorithm that follows the sum-product methodology while leveraging recent advances in message approximation [25]. In what follows, we use $\nu_{a \rightarrow b}(\cdot)$ to denote a message that is passed from node a to a connected node b .

A. Message Scheduling

Since the factor graph of Fig. 1 contains many cycles there are a number of valid ways to schedule, or sequence, the messages that are exchanged in the graph. We will describe two message passing schedules that empirically provide good convergence behavior and straightforward implementation. We refer to these two schedules as the *parallel message schedule* and the *serial message schedule*. In both cases, messages are first initialized to agnostic values, and then iteratively exchanged throughout the graph according to the chosen schedule until either convergence occurs, or a maximum number of allowable iterations is reached.

Conceptually, both message schedules can be decomposed into four distinct phases, differing only in which messages are initialized and the order in which the phases are sequenced. We label each phase using the mnemonics **(into)**, **(within)**, **(out)**, and **(across)**. In phase **(into)**, messages are passed from the s_n and $\theta_n^{(t)}$ variable nodes *into* frame t . Loosely speaking, these messages convey current beliefs about the values of s and $\theta^{(t)}$. In phase **(within)**, messages are exchanged *within* frame t , producing an estimate of $x^{(t)}$ using the current beliefs about s and $\theta^{(t)}$ together with the available measurements $\mathbf{y}^{(t)}$. In phase **(out)**, the estimate of $x^{(t)}$ is used to refine the beliefs about s and $\theta^{(t)}$ by passing messages *out* of frame t . Finally, in phase **(across)**, messages are sent from $\theta_n^{(t)}$ to either $\theta_n^{(t+1)}$ or $\theta_n^{(t-1)}$, thus conveying information *across* time about temporal correlation in the signal amplitudes.

The parallel message schedule begins by performing phase **(into)** in parallel for each frame $t = 1, \dots, T$ simultaneously. Then, phase **(within)** is performed simultaneously for each frame, followed by phase **(out)**. Next, information about the amplitudes is exchanged between the different timesteps by performing phase **(across)** in the forward direction, i.e., messages are passed from $\theta_n^{(1)}$ to $\theta_n^{(2)}$, and then from $\theta_n^{(2)}$ to $\theta_n^{(3)}$, proceeding until $\theta_n^{(T)}$ is reached. Finally, phase **(across)** is performed in the backward direction, where messages are passed consecutively from $\theta_n^{(T)}$ down to $\theta_n^{(1)}$. At this point, a single iteration of AMP-MMV has been completed, and a new iteration can commence starting with phase **(into)**. In this way, all of the available measurements, $\{\mathbf{y}^{(t)}\}_{t=1}^T$, are used to influence the recovery of the signal at each timestep.

The serial message schedule is similar to the parallel schedule except that it operates on frames in a sequential fashion, enabling causal processing of MMV signals. Beginning at the initial timestep, $t = 1$, the serial schedule first performs phase **(into)**, followed by phases **(within)** and **(out)**. Outgoing messages from the initial frame are then used in phase **(across)** to pass messages from $\theta_n^{(1)}$ to $\theta_n^{(2)}$. The messages arriving at $\theta_n^{(2)}$, along with updated beliefs about the value of s , are used to initiate phase **(into)** at timestep $t = 2$. Phases **(within)** and **(out)** are performed for frame 2, followed by another round of phase **(across)**, with messages being passed forward to $\theta_n^{(3)}$. This procedure continues until phase **(out)** is completed at frame T . Until now, only causal information has been used in producing estimates of the signal. If the application permits smoothing, then message passing continues in a similar fashion, but with messages now propagating backward in time, i.e., messages are passed from $\theta_n^{(T)}$ to $\theta_n^{(T-1)}$, phases **(into)**, **(within)**, and **(out)** are performed at frame $T - 1$, and then messages move from $\theta_n^{(T-1)}$ to $\theta_n^{(T-2)}$. The process continues until messages arrive at $\theta_n^{(1)}$, at which point a single *forward/backward pass* has been completed. We complete multiple such passes, resulting in a smoothed estimate of the signal.

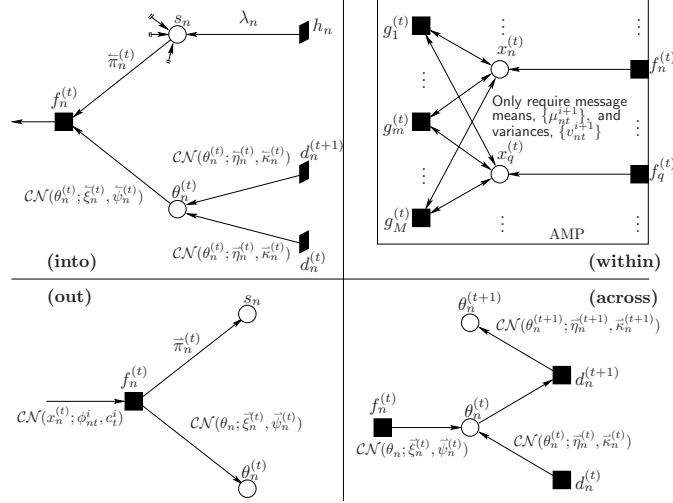


Fig. 2: A summary of the four message passing phases, including message notation and form.

B. Implementing the Message Passes

Space constraints prohibit us from providing a full derivation of all the messages that are exchanged through the factor graph of Fig. 1. Most messages can be derived by straightforward application of the rules of the sum-product algorithm. Therefore, in this sub-section we will restrict our attention to a handful of messages in the **(within)** and **(out)** phases whose implementation requires a departure from the sum-product rules for one reason or another.

To aid our discussion, in Fig. 2 we summarize each of the four phases, focusing primarily on a single coefficient index n at some intermediate frame t . Arrows indicate the direction that messages are moving, and only those nodes and edges participating in a particular phase are shown in that phase. For the **(across)** phase we show messages being passed forward in time, and omit a graphic for the corresponding backwards pass. The figure also introduces the notation that we adopt for the different variables that serve to parameterize the messages. Certain variables, e.g., $\bar{\eta}_n^{(t)}$ and $\bar{\eta}_n^{(t)}$, are accented with directional arrows. This is to distinguish variables associated with messages moving in one direction from those associated with messages moving in another. For Bernoulli message pdfs, we show only the nonzero probability, e.g., $\lambda_n = \nu_{h_n \rightarrow s_n}(s_n = 1)$.

Phase **(within)** entails using the messages transmitted from s_n and $\theta_n^{(t)}$ to $f_n^{(t)}$ to compute the messages that pass between $x_n^{(t)}$ and the $\{g_m^{(t)}\}$ nodes. Inspection of Fig. 2 reveals a dense interconnection between the $\{x_n^{(t)}\}$ and $\{g_m^{(t)}\}$ nodes. As a consequence, applying the standard sum-product rules to compute the $\nu_{g_m^{(t)} \rightarrow x_n^{(t)}}(\cdot)$ messages would result in an algorithm that required the evaluation of multi-dimensional integrals that grew exponentially in number in both N and M . Since we are strongly motivated to apply

AMP-MMV to high-dimensional problems, this approach is clearly infeasible. Instead, we turn to a recently developed algorithm known as *approximate message passing* (AMP).

AMP was originally proposed by Donoho et al. [25] as a message passing algorithm designed to solve the noiseless SMV CS problem known as Basis Pursuit ($\min \|\mathbf{x}\|_1$ s.t. $\mathbf{y} = \mathbf{A}\mathbf{x}$), and was subsequently extended [26] to support MMSE estimation under white-Gaussian-noise-corrupted observations and generic signal priors of the form $p(\mathbf{x}) = \prod p(x_n)$ through an approximation of the sum-product algorithm. In both cases, the associated factor graph looks identical to that of the **(within)** segment of Fig. 2. Conventional wisdom holds that loopy belief propagation only works well when the factor graph is locally tree-like. For general, non-sparse \mathbf{A} matrices, the **(within)** graph will clearly not possess this property, due to the many short cycles between the $x_n^{(t)}$ and $g_m^{(t)}$ nodes. Reasoning differently, Donoho et al. showed that the density of connections could prove beneficial, if properly exploited.

In particular, central limit theorem arguments suggest that the messages propagated from the g_m nodes to the x_n nodes under the sum-product algorithm can be well-approximated as Gaussian when the problem dimensionality is sufficiently high. Moreover, the computation of these Gaussian-approximated messages only requires knowledge of the mean and variance of the sum-product messages from the x_n to the g_m nodes. Finally, when $|A_{mn}|^2$ scales as $\mathcal{O}(1/M)$ for all (m, n) , the differences between the variances of the messages emitted by the x_n nodes vanish as M grows large, as do those of the g_m nodes when N grows large, allowing each to be approximated by a single, common variance. Together, these sum-product approximations yield an iterative thresholding algorithm with a particular first-order correction term that ensures both Gaussianity and independence in the residual error vector over the iterations. The complexity of this iterative thresholding algorithm is dominated by a single multiplication by \mathbf{A} and \mathbf{A}^H per iteration, implying a per-iteration computational cost of $\mathcal{O}(MN)$ flops. Furthermore, the state-evolution equation that governs the transient behavior of AMP shows that the number of required iterations does not scale with either M or N , implying that the total complexity is itself $\mathcal{O}(MN)$ flops.

AMP's suitability for the MMV problem stems from several considerations. First, AMP's probabilistic construction, coupled with its message passing implementation, makes it well-suited for incorporation as a subroutine within a larger message passing algorithm. In the MMV problem it is clear that $p(\bar{\mathbf{x}}) \neq \prod p(x_n^{(t)})$ due to the joint sparsity and amplitude correlation structure, and therefore AMP does not appear to be directly applicable. Fortunately, by modeling this structure through the hidden variables \mathbf{s} and $\bar{\boldsymbol{\theta}}$, we can exploit the conditional independence of the signal coefficients: $p(\bar{\mathbf{x}}|\mathbf{s}, \bar{\boldsymbol{\theta}}) = \prod p(x_n^{(t)}|s_n, \theta_n^{(t)})$. In particular, we replace the $p(x_n^{(t)})$ that AMP traditionally expects with $\nu_{f_n^{(t)} \rightarrow x_n^{(t)}}(\cdot)$, the most recent message moving into the **(within)** segment of Fig. 2. This message represents a "local prior" on $x_n^{(t)}$

given the current belief about the hidden variables s_n and $\theta_n^{(t)}$, and assumes the Bernoulli-Gaussian form

$$\nu_{f_n^{(t)} \rightarrow x_n^{(t)}}(x_n^{(t)}) = (1 - \bar{\pi}_n^{(t)})\delta(x_n^{(t)}) + \bar{\pi}_n^{(t)}\mathcal{CN}(x_n^{(t)}; \bar{\xi}_n^{(t)}, \bar{\psi}_n^{(t)}). \quad (10)$$

This ‘‘local prior’’ determines the AMP soft-thresholding functions defined in (D1) - (D4) of Table II. The derivation of these thresholding functions closely follows those outlined in [30], which considered the special case of a zero-mean Bernoulli-Gaussian prior.

Beyond the ease with which AMP is included into the larger message passing algorithm, a second factor that favors using AMP is the tremendous computational efficiency it imparts on high-dimensional problems. Using AMP to perform the most computationally intensive message passes enables AMP-MMV to attain a linear complexity scaling in all problem dimensions. To see why this is the case, note that the **(into)**, **(out)**, and **(across)** steps can be executed in $\mathcal{O}(N)$ flops/timestep, while AMP allows the **(within)** step to be executed in $\mathcal{O}(MN)$ flops/timestep (see (A4) - (A8) of Table II). Since these four steps are executed $\mathcal{O}(T)$ times per AMP-MMV iteration for both the serial and parallel message schedules, it follows that AMP-MMV’s overall complexity is $\mathcal{O}(TMN)$.³

A third appealing feature of AMP is that it is theoretically well-grounded; a recent analysis [40] shows that, for Gaussian \mathbf{A} in the large-system limit (i.e., $M, N \rightarrow \infty$ with M/N fixed), the behavior of AMP is governed by a state evolution whose fixed points, when unique, correspond to MMSE-optimal signal estimates.

After using AMP to implement phase **(within)**, we must pass messages out of frame t in order to update our beliefs about the values of \mathbf{s} and $\boldsymbol{\theta}^{(t)}$ in the **(out)** phase. Applying the sum-product algorithm rules to compute the message $\nu_{f_n^{(t)} \rightarrow \theta_n^{(t)}}(\cdot)$ results in the expression

$$\nu_{f_n^{(t)} \rightarrow \theta_n^{(t)}}^{\text{exact}}(\theta_n^{(t)}) \triangleq (1 - \bar{\pi}_n^{(t)})\mathcal{CN}(0; \phi_{nt}, c_t) + \bar{\pi}_n^{(t)}\mathcal{CN}(\theta_n^{(t)}; \phi_{nt}, c_t), \quad (11)$$

which is an improper distribution due to the constant (w.r.t. $\theta_n^{(t)}$) term $\mathcal{CN}(0; \phi_{nt}, c_t)$. This behavior is a consequence of the conditional signal model (4). In particular, when $s_n = 0$, $x_n^{(t)}$ provides no information about the value of $\theta_n^{(t)}$. Roughly speaking, the term $\mathcal{CN}(0; \phi_{nt}, c_t)$ corresponds to the distribution of $\theta_n^{(t)}$ conditioned on the case $s_n = 0$.

³The primary computational burden of executing AMP-MMV involves performing matrix-vector products with \mathbf{A} and \mathbf{A}^H , allowing it to be easily applied in problems where the measurement matrix is never stored explicitly, but rather is implemented implicitly through subroutines. Fast implicit \mathbf{A} operators can provide significant computational savings in high-dimensional problems; implementing a Fourier transform as a fast Fourier transform (FFT) subroutine, for example, would drop AMP-MMV’s complexity from $\mathcal{O}(TMN)$ to $\mathcal{O}(TN \log_2 N)$.

% Define soft-thresholding functions:	
$F_{nt}(\phi; c) \triangleq (1 + \gamma_{nt}(\phi; c))^{-1} \left(\frac{\overleftarrow{\psi}_n^{(t)} \phi + \overleftarrow{\xi}_n^{(t)} c}{\overleftarrow{\psi}_n^{(t)} + c} \right)$	(D1)
$G_{nt}(\phi; c) \triangleq (1 + \gamma_{nt}(\phi; c))^{-1} \left(\frac{\overleftarrow{\psi}_n^{(t)} c}{\overleftarrow{\psi}_n^{(t)} + c} \right) + \gamma_{nt}(\phi; c) F_{nt}(\phi; c) ^2$	(D2)
$F'_{nt}(\phi; c) \triangleq \frac{\partial}{\partial \phi} F_{nt}(\phi; c) = \frac{1}{c} G_{nt}(\phi; c)$	(D3)
$\gamma_{nt}(\phi; c) \triangleq \left(\frac{1 - \overleftarrow{\pi}_n^{(t)}}{\overleftarrow{\pi}_n^{(t)}} \right) \left(\frac{\overleftarrow{\psi}_n^{(t)} + c}{c} \right) \times \exp \left(- \left[\frac{\overleftarrow{\psi}_n^{(t)} \phi ^2 + \overleftarrow{\xi}_n^{(t)} * c \phi + \overleftarrow{\xi}_n^{(t)} c \phi^* - c \overleftarrow{\xi}_n^{(t)} ^2}{c(\overleftarrow{\psi}_n^{(t)} + c)} \right] \right)$	(D4)
% Begin passing messages . . .	
for $t = 1, \dots, T, \forall n :$	
% Execute the (into) phase . . .	
$\overleftarrow{\pi}_n^{(t)} = \frac{\lambda_n \cdot \prod_{t' \neq t} \overleftarrow{\pi}_n^{(t')}}{(1 - \lambda_n) \cdot \prod_{t' \neq t} (1 - \overleftarrow{\pi}_n^{(t')}) + \lambda_n \cdot \prod_{t' \neq t} \overleftarrow{\pi}_n^{(t')}}$	(A1)
$\overleftarrow{\psi}_n^{(t)} = \frac{\overleftarrow{\kappa}_n^{(t)} \cdot \overleftarrow{\kappa}_n^{(t)}}{\overleftarrow{\kappa}_n^{(t)} + \overleftarrow{\kappa}_n^{(t)}}$	(A2)
$\overleftarrow{\xi}_n^{(t)} = \overleftarrow{\psi}_n^{(t)} \cdot \left(\frac{\overleftarrow{\eta}_n^{(t)}}{\overleftarrow{\kappa}_n^{(t)}} + \frac{\overleftarrow{\zeta}_n^{(t)}}{\overleftarrow{\kappa}_n^{(t)}} \right)$	(A3)
% Initialize AMP-related variables . . .	
$\forall m : z_{mt}^1 = y_m^{(t)}, \forall n : \mu_{nt}^1 = 0, \text{ and } c_t^1 = 100 \cdot \sum_{n=1}^N \psi_n^{(t)}$	
% Execute the (within) phase using AMP . . .	
for $i = 1, \dots, I, \forall n, m :$	
$\phi_{nt}^i = \sum_{m=1}^M A_{mn}^* z_{mt}^i + \mu_{nt}^i$	(A4)
$\mu_{nt}^{i+1} = F_{nt}(\phi_{nt}^i; c_t^i)$	(A5)
$v_{nt}^{i+1} = G_{nt}(\phi_{nt}^i; c_t^i)$	(A6)
$c_t^{i+1} = \sigma_e^2 + \frac{1}{M} \sum_{n=1}^N v_{nt}^{i+1}$	(A7)
$z_{mt}^{i+1} = y_m^{(t)} - \mathbf{a}_m^T \boldsymbol{\mu}_t^{i+1} + \frac{z_{mt}^i}{M} \sum_{n=1}^N F'_{nt}(\phi_{nt}^i; c_t^i)$	(A8)
end	
$\hat{x}_n^{(t)} = \mu_{nt}^{I+1}$ % Store current estimate of $x_n^{(t)}$	(A9)
% Execute the (out) phase . . .	
$\overleftarrow{\pi}_n^{(t)} = \left(1 + \left(\frac{\overleftarrow{\pi}_n^{(t)}}{1 - \overleftarrow{\pi}_n^{(t)}} \right) \gamma_{nt}(\phi_{nt}^I; c_t^{I+1}) \right)^{-1}$	(A10)
$(\overleftarrow{\xi}_n^{(t)}, \overleftarrow{\psi}_n^{(t)}) = \text{taylor_approx}(\overleftarrow{\pi}_n^{(t)}, \phi_{nt}^I, c_t^I)$	(A11)
% Execute the (across) phase from $\theta_n^{(t)}$ to $\theta_n^{(t+1)}$. . .	
$\overleftarrow{\eta}_n^{(t+1)} = (1 - \alpha) \left(\frac{\overleftarrow{\kappa}_n^{(t)} \overleftarrow{\psi}_n^{(t)}}{\overleftarrow{\kappa}_n^{(t)} + \overleftarrow{\psi}_n^{(t)}} \right) \left(\frac{\overleftarrow{\eta}_n^{(t)}}{\overleftarrow{\kappa}_n^{(t)}} + \frac{\overleftarrow{\xi}_n^{(t)}}{\overleftarrow{\psi}_n^{(t)}} \right) + \alpha \zeta$	(A12)
$\overleftarrow{\kappa}_n^{(t+1)} = (1 - \alpha)^2 \left(\frac{\overleftarrow{\kappa}_n^{(t)} \overleftarrow{\psi}_n^{(t)}}{\overleftarrow{\kappa}_n^{(t)} + \overleftarrow{\psi}_n^{(t)}} \right) + \alpha^2 \rho$	(A13)
end	

TABLE II: Message update equations for executing a single forward pass using the serial message schedule.

As a means of circumventing the improper message pdf above, we will regard our original signal model, in which $s_n \in \{0, 1\}$, as the limiting case of a signal model in which $s_n \in \{\varepsilon, 1\}$ with $\varepsilon \rightarrow 0$. For any fixed, positive ε , $\nu_{f_n^{(t)} \rightarrow \theta_n^{(t)}}(\cdot)$ is given by the proper pdf

$$\nu_{f_n^{(t)} \rightarrow \theta_n^{(t)}}^{\text{mod}}(\theta_n^{(t)}) = (1 - \Omega(\overleftarrow{\pi}_n^{(t)})) \mathcal{CN}(\theta_n^{(t)}; \frac{1}{\varepsilon} \phi_{nt}, \frac{1}{\varepsilon^2} c_t) + \Omega(\overleftarrow{\pi}_n^{(t)}) \mathcal{CN}(\theta_n^{(t)}; \phi_{nt}, c_t), \quad (12)$$

where

$$\Omega(\pi) \triangleq \frac{\varepsilon^2 \pi}{(1 - \pi) + \varepsilon^2 \pi}. \quad (13)$$

Equation (12) is a binary Gaussian mixture density. When $\varepsilon \ll 1$, the first Gaussian component is extremely broad, and conveys little information about the possible value of $\theta_n^{(t)}$. The second component is a more informative Gaussian whose mean, ϕ_{nt} , and variance, c_t , are determined by the product of the messages $\{\nu_{g_m^{(t)} \rightarrow x_n^{(t)}}(\cdot)\}_{m=1}^M$. The relative mass assigned to each Gaussian component is a function of the incoming activity probability $\bar{\pi}_n^{(t)}$ (see (10)). Note that the limiting case of $\Omega(\cdot)$ is a simple indicator function:

$$\lim_{\varepsilon \rightarrow 0} \Omega(\pi) = \begin{cases} 0 & \text{if } 0 \leq \pi < 1, \\ 1 & \text{if } \pi = 1. \end{cases} \quad (14)$$

When implementing AMP-MMV, we therefore fix ε at a small positive value, e.g., $\varepsilon = 1 \times 10^{-7}$. If desired, (12) could then be used as the outgoing message, however this would present a further difficulty. Propagating a Gaussian mixture along a given edge would result in an exponential growth in the number of mixture components that would need to be propagated along subsequent edges. To avoid this outcome, we collapse our binary Gaussian mixture to a single Gaussian component, an approach sometimes referred to as *Gaussian sum approximation* [41], [42]. Since, for $\varepsilon \ll 1$, $\Omega(\cdot)$ behaves nearly like the indicator function in (14), one of the two Gaussian components will typically have negligible mass. For this reason, collapsing the mixture to a single Gaussian appears justifiable.

To carry out the collapsing, we perform a second-order Taylor series approximation of $-\log \nu_{f_n^{(t)} \rightarrow \theta_n^{(t)}}^{\text{mod}}(\theta_n^{(t)})$ with respect to $\theta_n^{(t)}$ about the point ϕ_{nt} .⁴ This provides the mean, $\bar{\xi}_n^{(t)}$, and variance, $\bar{\psi}_n^{(t)}$, of the single Gaussian that serves as $\nu_{f_n^{(t)} \rightarrow \theta_n^{(t)}}(\cdot)$. (See Fig. 2.) In Appendix A we summarize the Taylor approximation procedure, and in Table III provide the pseudocode function, `taylor_approx`, for computing $\bar{\xi}_n^{(t)}$ and $\bar{\psi}_n^{(t)}$.

With the exception of the messages discussed above, all the remaining messages can be derived using the standard sum-product algorithm rules [32]. For convenience, we summarize the results in Table II, where we provide a pseudocode implementation of a single forward pass of AMP-MMV using the serial message schedule.

V. ESTIMATING THE MODEL PARAMETERS

The signal model of Section II depends on the sparsity parameters $\{\lambda_n\}_{n=1}^N$, amplitude parameters ζ , α , and ρ , and noise variance σ_e^2 . While some of these parameters may be known accurately from prior

⁴For technical reasons, the Taylor series approximation is performed in \mathbb{R}^2 instead of \mathbb{C} .

function $(\vec{\xi}, \vec{\psi}) = \text{taylor_approx}(\bar{\pi}, \phi, c)$	
% Define useful variables:	
$a \triangleq \varepsilon^2(1 - \Omega(\bar{\pi}))$	(T1)
$\bar{a} \triangleq \Omega(\bar{\pi})$	(T2)
$b \triangleq \frac{\varepsilon^2}{c} (1 - \frac{1}{\varepsilon})\phi ^2$	(T3)
$\vartheta_r \triangleq -\frac{2\varepsilon^2}{c} (1 - \frac{1}{\varepsilon})\Re\{\phi\}$	(T4)
$\vartheta_i \triangleq -\frac{2\varepsilon^2}{c} (1 - \frac{1}{\varepsilon})\Im\{\phi\}$	(T5)
% Compute outputs:	
$\vec{\psi} = \frac{(a^2e^{-b} + a\bar{a} + \bar{a}^2e^b)c}{\varepsilon^2a^2e^{-b} + a\bar{a}(\varepsilon^2 + 1 - \frac{1}{2}c\vartheta_r^2) + \bar{a}^2e^b}$	(T6)
$\vec{\xi}_r = \phi_r - \frac{1}{2}\vec{\psi} \frac{-ae^{-b}\vartheta_r}{ae^{-b} + \bar{a}}$	(T7)
$\vec{\xi}_i = \phi_i - \frac{1}{2}\vec{\psi} \frac{-ae^{-b}\vartheta_i}{ae^{-b} + \bar{a}}$	(T8)
$\vec{\xi} = \vec{\xi}_r + j\vec{\xi}_i$	(T9)
return $(\vec{\xi}, \vec{\psi})$	

TABLE III: Pseudocode function for computing a single-Gaussian approximation of (12).

information, it is likely that many will require tuning. To this end, we develop an expectation-maximization (EM) algorithm that couples with the message passing procedure described in Section IV-A to provide a means of learning all of the model parameters while simultaneously estimating the signal \bar{x} and its support s .

The EM algorithm [43] is an appealing choice for performing parameter estimation for two primary reasons. First and foremost, the EM algorithm is a well-studied and principled means of parameter estimation. At every EM iteration, the data likelihood function is guaranteed to increase until convergence to a local maximum of the likelihood function occurs [43]. For multimodal likelihood functions, local maxima will, in general, not coincide with the global maximum likelihood (ML) estimator, however a judicious initialization can help in ensuring the EM algorithm reaches the global maximum [44]. Second, the expectation step of the EM algorithm relies on quantities that have already been computed in the process of executing AMP-MMV. Ordinarily, this step constitutes the major computational burden of any EM algorithm, thus the fact that we can perform it essentially for free makes our EM procedure highly efficient.

We let $\Gamma \triangleq \{\lambda, \zeta, \alpha, \rho, \sigma_e^2\}$ denote the set of all model parameters, and let Γ^k denote the set of parameter estimates at the k^{th} EM iteration. Here we have assumed that the binary support indicator variables share a common activity probability, λ , i.e., $\Pr\{s_n = 1\} = \lambda \forall n$. For all parameters except σ_e^2 we use s and $\bar{\theta}$ as the so-called ‘‘missing’’ data of the EM algorithm, while for σ_e^2 we use \bar{x} .

For the first iteration of AMP-MMV, the model parameters are initialized based on either prior signal

knowledge, or according to some heuristic criteria. Using these parameter values, AMP-MMV performs either a single iteration of the parallel message schedule, or a single forward/backward pass of the serial message schedule, as described in Section IV-A. Upon completing this first iteration, approximate marginal posterior distributions are available for each of the underlying random variables, e.g., $p(x_n^{(t)}|\bar{\mathbf{y}})$, $p(s_n|\bar{\mathbf{y}})$, and $p(\theta_n^{(t)}|\bar{\mathbf{y}})$. Additionally, belief propagation can provide pairwise joint posterior distributions, e.g., $p(\theta_n^{(t)}, \theta_n^{(t-1)}|\bar{\mathbf{y}})$, for any variable nodes connected by a common factor node [45]. With these marginal, and pairwise joint, posterior distributions, it is possible to perform the iterative expectation and maximization steps required to maximize $p(\bar{\mathbf{y}}|\Gamma)$ in closed-form. We adopt a Gauss-Seidel scheme, performing coordinate-wise maximization, e.g.,

$$\lambda^{k+1} = \operatorname{argmax}_{\lambda} \mathbb{E}_{\mathbf{s}, \bar{\boldsymbol{\theta}}|\bar{\mathbf{y}}} \left[\log p(\bar{\mathbf{y}}, \mathbf{s}, \bar{\boldsymbol{\theta}}; \lambda, \Gamma^k \setminus \{\lambda^k\}) \middle| \bar{\mathbf{y}}, \Gamma^k \right],$$

where k is the iteration index common to both AMP-MMV and the EM algorithm.

In Table IV we provide the EM parameter update equations for our signal model. In practice, we found that the robustness and convergence behavior of our EM procedure were improved if we were selective about which parameters we updated on a given iteration. For example, the parameters α and ρ are tightly coupled to one another, since $\operatorname{var}\{\theta_n^{(t)}|\theta_n^{(t-1)}\} = \alpha^2\rho$. Consequently, if the initial choices of α and ρ are too small, it is possible that the EM procedure will overcompensate on the first iteration by producing revised estimates of both parameters that are too large. This leads to an oscillatory behavior in the EM updates that can be effectively combated by avoiding updating both α and ρ on the same iteration.

VI. NUMERICAL STUDY

In this section we describe the results of an extensive numerical study that was conducted to explore the performance characteristics and tradeoffs of AMP-MMV. MATLAB code⁵ was written to implement both the parallel and serial message schedules of Section IV-A, along with the EM parameter estimation procedure of Section V.

For comparison to AMP-MMV, we tested two other Bayesian algorithms for the MMV problem, MSBL [15] and T-MSBL⁶ [18], which have been shown to offer “best in class” performance on the MMV problem. We also included a recently proposed greedy algorithm designed specifically for highly correlated signals, subspace-augmented MUSIC⁷ (SA-MUSIC), which has been shown to outperform

⁵Code available at ece.osu.edu/~schniter/turboAMPmmv.

⁶Code available at dsp.ucsd.edu/~zhilin/Software.html.

⁷Code obtained through personal correspondence with authors.

% Define key quantities obtained from AMP-MMV at iteration k :	
$E[s_n \bar{\mathbf{y}}] = \frac{\lambda_n \prod_{t=1}^T \bar{\pi}_n^{(t)}}{\lambda_n \prod_{t=1}^T \bar{\pi}_n^{(t)} + (1-\lambda_n) \prod_{t=1}^T (1-\bar{\pi}_n^{(t)})}$	(Q1)
$\tilde{v}_n^{(t)} \triangleq \text{var}\{\theta_n^{(t)} \bar{\mathbf{y}}\} = \left(\frac{1}{\kappa_n^{(t)}} + \frac{1}{\psi_n^{(t)}} + \frac{1}{\kappa_n^{(t)}} \right)^{-1}$	(Q2)
$\tilde{\mu}_n^{(t)} \triangleq E[\theta_n^{(t)} \bar{\mathbf{y}}] = \tilde{v}_n^{(t)} \cdot \left(\frac{\eta_n^{(t)}}{\kappa_n^{(t)}} + \frac{\xi_n^{(t)}}{\psi_n^{(t)}} + \frac{\zeta_n^{(t)}}{\kappa_n^{(t)}} \right)$	(Q3)
$v_n^{(t)} \triangleq \text{var}\{x_n^{(t)} \bar{\mathbf{y}}\}$	% See (A6) of Table II
$\mu_n^{(t)} \triangleq E[x_n^{(t)} \bar{\mathbf{y}}]$	% See (A5) of Table II
% EM update equations:	
$\lambda^{k+1} = \frac{1}{N} \sum_{n=1}^N E[s_n \bar{\mathbf{y}}]$	(E1)
$\zeta^{k+1} = \left(\frac{N(T-1)}{\rho^k} + \frac{N}{(\sigma^2)^k} \right)^{-1} \left(\frac{1}{(\sigma^2)^k} \sum_{n=1}^N \tilde{\mu}_n^{(1)} + \sum_{t=2}^T \sum_{n=1}^N \frac{1}{\alpha^k \rho^k} (\tilde{\mu}_n^{(t)} - (1-\alpha^k) \tilde{\mu}_n^{(t-1)}) \right)$	(E2)
$\alpha^{k+1} = \frac{1}{4N(T-1)} \left(\mathfrak{b} - \sqrt{\mathfrak{b}^2 + 8N(T-1)\mathfrak{c}} \right)$	(E3)
where:	
$\mathfrak{b} \triangleq \frac{2}{\rho^k} \sum_{t=2}^T \sum_{n=1}^N \Re\{E[\theta_n^{(t)*} \theta_n^{(t-1)} \bar{\mathbf{y}}]\} - \Re\{(\tilde{\mu}_n^{(t)} - \tilde{\mu}_n^{(t-1)})^* \zeta^k\} - \tilde{v}_n^{(t-1)} - \tilde{\mu}_n^{(t-1)} ^2$	
$\mathfrak{c} \triangleq \frac{2}{\rho^k} \sum_{t=2}^T \sum_{n=1}^N \tilde{v}_n^{(t)} + \tilde{\mu}_n^{(t)} ^2 + \tilde{v}_n^{(t-1)} + \tilde{\mu}_n^{(t-1)} ^2 - 2\Re\{E[\theta_n^{(t)*} \theta_n^{(t-1)} \bar{\mathbf{y}}]\}$	
$\rho^{k+1} = \frac{1}{(\alpha^k)^2 N(T-1)} \sum_{t=2}^T \sum_{n=1}^N \tilde{v}_n^{(t)} + \tilde{\mu}_n^{(t)} ^2 + (\alpha^k)^2 \zeta^k ^2 - 2(1-\alpha^k) \Re\{E[\theta_n^{(t)*} \theta_n^{(t-1)} \bar{\mathbf{y}}]\} - 2\alpha^k \Re\{(\tilde{\mu}_n^{(t)*} \zeta^k\} + 2\alpha^k (1-\alpha^k) \Re\{(\tilde{\mu}_n^{(t-1)*} \zeta^k\} + (1-\alpha^k)(\tilde{v}_n^{(t-1)} + \tilde{\mu}_n^{(t-1)} ^2)$	(E4)
$\sigma_e^{2k+1} = \frac{1}{TM} \left(\sum_{t=1}^T \ \mathbf{y}^{(t)} - \mathbf{A}\boldsymbol{\mu}^{(t)}\ ^2 + \mathbf{1}_N^T \mathbf{v}^{(t)} \right)$	(E5)

TABLE IV: EM algorithm update equations for the signal model parameters of Section II.

MMV basis pursuit and several correlation-agnostic greedy methods [14]. Finally, we implemented the support-aware Kalman smoother (SKS), which, as noted in Section III, provides a lower bound on the achievable MSE of any algorithm. To implement the SKS, we took advantage of the fact that $\bar{\mathbf{y}}$, $\bar{\mathbf{x}}$, and $\bar{\boldsymbol{\theta}}$ are jointly Gaussian when conditioned on the support, \mathbf{s} , and thus Fig. 1 becomes a Gaussian graphical model. Consequently, the sum-product algorithm yields closed-form expressions (i.e., no approximations are required) for each of the messages traversing the graph. Therefore, it is possible to obtain the desired posterior means (i.e., MMSE estimates of $\bar{\mathbf{x}}$) despite the fact that the graph is loopy [46, Claim 5].

In all of our experiments, performance was analyzed on synthetically generated datasets, and averaged over 250 independent trials. Since MSBL and T-MSBL were derived for real-valued signals, we used a real-valued equivalent of the signal model described in Section II, and ran a real-valued version of AMP-MMV. Our data generation procedure closely mirrors the one used to characterize T-MSBL in [18]. Unless otherwise stated, the measurement matrices were i.i.d. Gaussian random matrices with unit-norm columns, $T = 4$ measurement vectors were generated, the stationary variance of the amplitude process

was set at $\sigma^2 \triangleq \frac{\alpha\rho}{2-\alpha} = 1$, and the noise variance σ_e^2 was set to yield an SNR of 25 dB.

Three performance metrics were considered throughout our tests. The first metric, which we refer to as the time-averaged normalized MSE (TNMSE), is defined as

$$\text{TNMSE}(\bar{\mathbf{x}}, \hat{\mathbf{x}}) \triangleq \frac{1}{T} \sum_{t=1}^T \frac{\|\mathbf{x}^{(t)} - \hat{\mathbf{x}}^{(t)}\|_2^2}{\|\mathbf{x}^{(t)}\|_2^2},$$

where $\hat{\mathbf{x}}^{(t)}$ is an estimate of $\mathbf{x}^{(t)}$. The second metric, intended to gauge the accuracy of the recovered support, is the normalized support error rate (NSER), which is defined as the number of indices in which the true and estimated support differ, normalized by the cardinality of the true support \mathcal{S} . The third and final metric is runtime, which is an important metric given the prevalence of high-dimensional datasets.

The algorithms were configured and executed as follows: to obtain support estimates for MSBL, T-MSBL, and AMP-MMV, we adopted the technique utilized in [18] of identifying the K amplitude trajectories with the largest ℓ_2 norms as the support set, where $K \triangleq |\mathcal{S}|$. Note that this is an optimistic means of identifying the support, as it assumes that an oracle provides the true value of K . For this reason, we implemented an additional *non-oracle-aided* support estimate for AMP-MMV that consisted of those indices n for which $\hat{p}(s_n|\bar{\mathbf{y}}) > \frac{1}{2}$. In all simulations, AMP-MMV was given imperfect knowledge of the signal model parameters, and refined the initial parameter choices according to the EM update procedure given in Table IV. In particular, the noise variance was initialized at $\sigma_e^2 = 1 \times 10^{-3}$. The remaining parameters were initialized agnostically using simple heuristics that made use of sample statistics derived from the available measurements, $\bar{\mathbf{y}}$. Equation (A9) of Table II was used to produce $\hat{\mathbf{x}}^{(t)}$, which corresponds to an MMSE estimate of $\mathbf{x}^{(t)}$ under AMP-MMV's estimated posteriors $\hat{p}(x_n^{(t)}|\bar{\mathbf{y}})$. In the course of running simulations, we monitored the residual energy, $\sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{A}\hat{\mathbf{x}}^{(t)}\|_2^2$, and would automatically switch the schedule, e.g., from parallel to serial, and/or change the maximum number of iterations whenever the residual energy exceeded a noise variance-dependent threshold. The SKS was given perfect parameter and support knowledge and was run until convergence. Both MSBL and T-MSBL were tuned in a manner recommended by the codes' authors. SA-MUSIC was given the true value of K , and upon generating an estimate of the support, $\hat{\mathcal{S}}$, a conditional MMSE signal estimate was produced, e.g., $\hat{\mathbf{x}}^{(t)} = \text{E}[\mathbf{x}^{(t)}|\hat{\mathcal{S}}, \mathbf{y}^{(t)}]$.

A. Performance Versus Sparsity, M/K

As a first experiment, we studied how performance changes as a function of the measurements-to-active-coefficients ratio, M/K . For this experiment, $N = 5000$, $M = 1563$, and $T = 4$. The activity




Fig. 3: A plot of the TNMSE (in dB), NSER, and runtime of T-MSBL, MSBL, SA-MUSIC, AMP-MMV, and the SKS versus M/K . Correlation coefficient $1 - \alpha = 0.90$.

probability, λ , was swept over the range $[0.096, 0.22]$, implying that the ratio of measurements-to-active-coefficients, M/K , ranged from 1.42 to 3.26.

In Fig. 3, we plot the performance when the temporal correlation of the amplitudes is $1 - \alpha = 0.90$. For AMP-MMV, two traces appear on the NSER plot, with the \circ marker corresponding to the K -largest-trajectory-norm method of support estimation, and the \triangle marker corresponding to the support estimate obtained from the posteriors $\hat{p}(s_n|\bar{\mathbf{y}})$. We see that, when $M/K \geq 2$, the TNMSE performance of both AMP-MMV and T-MSBL is almost identical to that of the oracle-aided SKS. However, when $M/K < 2$, every algorithm's support estimation performance (NSER) degrades, and the TNMSE consequently grows. Indeed, when $M/K < 1.50$, all of the algorithms perform poorly compared to the SKS, although T-MSBL performs the best of the four. We also note the superior NSER performance of AMP-MMV over much of the range, even when using $p(s_n|\bar{\mathbf{y}})$ to estimate \mathcal{S} (and thus not requiring apriori knowledge of K). From the runtime plot we see the tremendous efficiency of AMP-MMV. Over the region in which AMP-MMV is performing well (and thus not cycling through multiple configurations in vain), we see that AMP-MMV's runtime is more than one order-of-magnitude faster than SA-MUSIC, and two orders-of-magnitude faster than either T-MSBL or MSBL.

In Fig. 4 we repeat the same experiment, but with increased amplitude correlation $1 - \alpha = 0.99$. In this case we see that AMP-MMV and T-MSBL still offer a TNMSE performance that is comparable to the SKS when $M/K \geq 2.50$, whereas the performance of both MSBL and SA-MUSIC has degraded across-the-board. When $M/K < 2.5$, the NSER and TNMSE performance of AMP-MMV and T-MSBL decay sharply, and all the methods considered perform poorly compared to the SKS. Our finding that performance is adversely affected by increased temporal correlation is consistent with the theoretical and empirical findings of [8], [14], [15], [18]. Interestingly, the performance of the SKS shows a modest




Fig. 4: A plot of the TNMSE (in dB), NSER, and runtime of T-MSBL, MSBL, SA-MUSIC, AMP-MMV, and the SKS versus M/K . Correlation coefficient $1 - \alpha = 0.99$.




Fig. 5: A plot of the TNMSE (in dB), NSER, and runtime of T-MSBL, MSBL, SA-MUSIC, AMP-MMV, and the SKS versus T . Correlation coefficient $1 - \alpha = 0.90$.

improvement compared to Fig. 3, reflecting the fact that the slower temporal variations of the amplitudes are easier to track when the support is known.

B. Performance Versus T

In a second experiment, we studied how performance is affected by the number of measurement vectors, T , used in the reconstruction. For this experiment, we used $N = 5000$, $M = N/5$, and $\lambda = 0.10$ ($M/K = 2$). Figure 5 shows the performance with a correlation of $1 - \alpha = 0.90$. Comparing to Fig. 3, we see that MSBL's performance is strongly impacted by the reduced value of M . AMP-MMV and T-MSBL perform more-or-less equivalently across the range of T , although AMP-MMV does so with an order-of-magnitude reduction in complexity. It is interesting to observe that, in this problem regime, the SKS TNMSE bound is insensitive to the number of measurement vectors acquired.




Fig. 6: A plot of the TNMSE (in dB), NSER, and runtime of T-MSBL, MSBL, SA-MUSIC, AMP-MMV, and the SKS versus SNR. Correlation coefficient $1 - \alpha = 0.95$.

C. Performance Versus SNR

To understand how AMP-MMV performs in low SNR environments, we conducted a test in which SNR was swept from 5 dB to 25 dB.⁸ The problem dimensions were fixed at $N = 5000$, $M = N/5$, and $T = 4$. The sparsity rate, λ , was chosen to yield $M/K = 3$ measurements-per-active-coefficient, and the correlation was set at $1 - \alpha = 0.95$.

Our findings are presented in Fig. 6. Both T-MSBL and MSBL operate within 5 - 10 dB of the SKS in TNMSE across the range of SNRs, while AMP-MMV operates ≈ 5 dB from the SKS when the SNR is at or below 10 dB, and approaches the SKS in performance as the SNR elevates. We also note that using AMP-MMV's posteriors on s_n to estimate the support does not appear to perform much worse than the K -largest-trajectory-norm method for high SNRs, and shows a slight advantage at low SNRs. The increase in runtime exhibited by AMP-MMV in this experiment is a consequence of our decision to configure AMP-MMV identically for all experiments; our initialization of the noise variance, σ_e^2 , was more than an order-of-magnitude off over the majority of the SNR range, and thus AMP-MMV cycled through many different schedules in an effort to obtain an (unrealistic) residual energy. Runtime could be drastically improved in this experiment by using a more appropriate initialization of σ_e^2 .

D. Performance Versus Undersampling Rate, N/M

As mentioned in Section I, one of the principal aims of CS is to reduce the number of measurements that must be acquired while still obtaining a good solution. In the MMV problem, dramatic reductions in the sampling rate are possible. To illustrate this, in Fig. 7 we present the results of an experiment in

⁸In lower SNR regimes, learning rules for the noise variance are known to become less reliable [15], [18]. Still, for high-dimensional problems, a sub-optimal learning rule may be preferable to a computationally costly cross-validation procedure. For this reason, we ran all three Bayesian algorithms with a learning rule for the noise variance enabled.

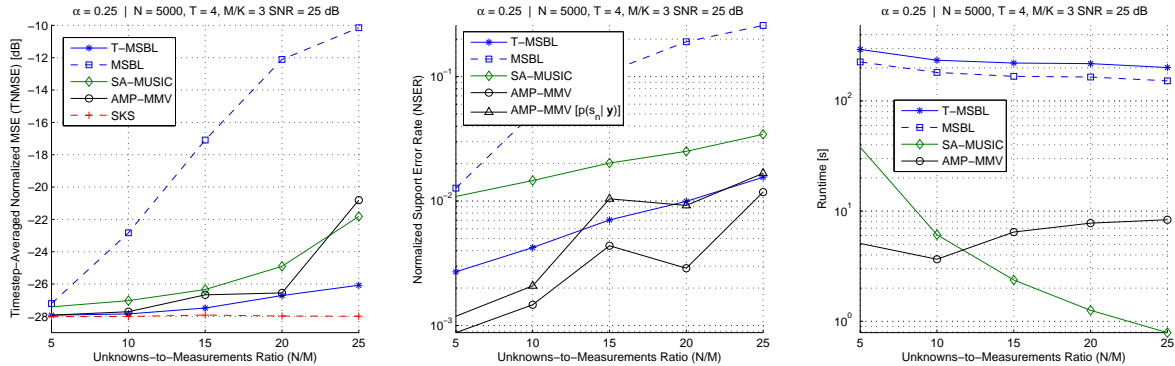


Fig. 7: A plot of the TNMSE (in dB), NSER, and runtime of T-MSBL, MSBL, SA-MUSIC, AMP-MMV, and the SKS versus undersampling rate, N/M . Correlation coefficient $1 - \alpha = 0.75$, which the undersampling factor, N/M , was varied from 5 to 25 unknowns-per-measurement. Specifically, N was fixed at 5000, while M was varied. λ was likewise adjusted in order to keep M/K fixed at 3 measurements-per-active-coefficient. In Fig. 7, we see that MSBL quickly departs from the SKS performance bound, whereas AMP-MMV, T-MSBL, and SA-MUSIC are able to remain close to the bound when $N/M \leq 20$. At $N/M = 25$, both AMP-MMV and SA-MUSIC have diverged from the bound, and, while still offering an impressive TNMSE, they are outperformed by T-MSBL. In conducting this test, we observed that AMP-MMV’s performance is strongly tied to the number of smoothing iterations performed. Whereas for other tests, 5 smoothing iterations were often sufficient, in scenarios with a high degree of undersampling, (e.g., $N/M \geq 15$), 50 – 100 smoothing iterations were often required to obtain good signal estimates. This suggests that messages must be exchanged between neighboring timesteps over many iterations in order to arrive at consensus in severely underdetermined problems.

E. Performance Versus Signal Dimension, N

As we have indicated throughout this paper, a key consideration of our method was ensuring that it would be suitable for high-dimensional problems. Our complexity analysis indicated that a single iteration of AMP-MMV could be completed in $\mathcal{O}(TNM)$ flops. This linear scaling of the complexity with respect to problem dimensions gives encouragement that our algorithm should efficiently handle large problems, but if the number of iterations required to obtain a solution grows too rapidly with problem size, our technique would be of limited practical utility. To ensure that this was not the case, we performed an experiment in which the signal dimension, N , was swept logarithmically over the range $[100, 10000]$. M was scaled proportionally such that $N/M = 3$. The sparsity rate was fixed at $\lambda = 0.15$ so that $M/K \approx 2$, and the correlation was set at $1 - \alpha = 0.95$.

The results of this experiment are provided in Fig. 8. Several features of these plots are of interest.

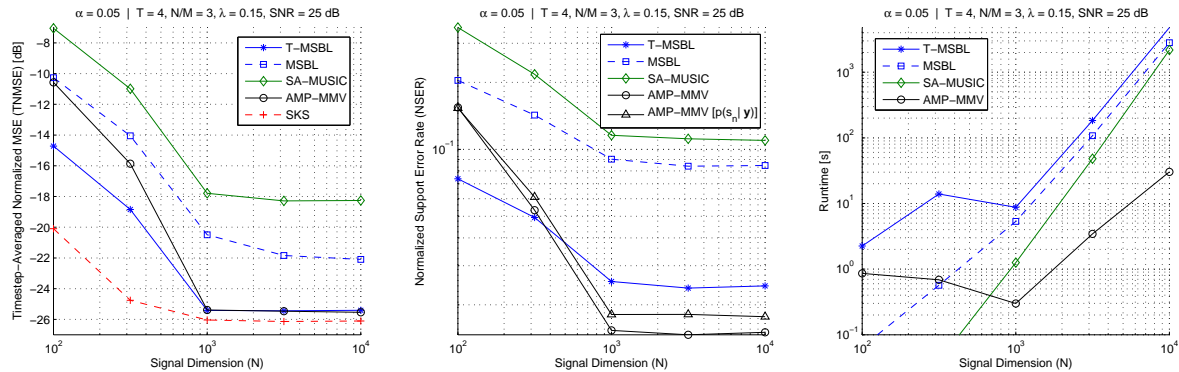


Fig. 8: A plot of the TNMSE (in dB), NSER, and runtime of T-MSBL, MSBL, SA-MUSIC, AMP-MMV, and the SKS versus signal dimension, N . Correlation coefficient $1 - \alpha = 0.95$.

First, we observe that the performance of every algorithm improves noticeably as problem dimensions grow from $N = 100$ to $N = 1000$, with AMP-MMV and T-MSBL converging in TNMSE performance to the SKS bound. The second observation that we point out is that AMP-MMV works extremely quickly. Indeed, a problem with $NT = 40000$ unknowns can be solved accurately in just under 30 seconds. Finally, we note that at small problem dimensions, AMP-MMV is not as quick as either MSBL or SA-MUSIC, however AMP-MMV scales with increasing problem dimensions more favorably than the other methods; at $N = 10000$ we note that AMP-MMV runs at least two orders-of-magnitude faster than the other techniques.

F. Performance With Time-Varying Measurement Matrices

In all of the previous experiments, we considered the standard MMV problem (1), in which all of the measurement vectors were acquired using a single, common measurement matrix. While this setup is appropriate for many tasks, there are a number of practical applications in which a joint-sparse signal is measured through distinct measurement matrices.

To better understand what, if any, gains can be obtained from diversity in the measurement matrices, we designed an experiment that explored how performance is affected by the rate-of-change of the measurement matrix over time. For simplicity, we considered a first-order Gauss-Markov random process to describe how a given measurement matrix changed over time. Specifically, we started with a matrix whose columns were drawn i.i.d. Gaussian as in previous experiments, which was then used as the measurement matrix to collect the measurements at timestep $t = 1$. At subsequent timesteps, the matrix evolved according to

$$\mathbf{A}^{(t)} = (1 - \beta)\mathbf{A}^{(t-1)} + \beta\mathbf{U}^{(t)}, \quad (15)$$




Fig. 9: A plot of the TNMSE (in dB), NSER, and runtime of AMP-MMV and the SKS versus rate-of-change of the measurement matrix, β . Correlation coefficient $1 - \alpha = 0.99$, where $\mathbf{U}^{(t)}$ was a matrix whose elements were drawn i.i.d. Gaussian, with a variance chosen such that the column norm of $\mathbf{A}^{(t)}$ would (in expectation) equal one.

In the test, β was swept over a range, providing a quantitative measure of the rate-of-change of the measurement matrix over time. Clearly, $\beta = 0$ would correspond to the standard MMV problem, while $\beta = 1$ would represent a collection of statistically independent measurement matrices.

In Fig. 9 we show the performance when $N = 5000$, $N/M = 30$, $M/K = 2$, and the correlation is $1 - \alpha = 0.99$. For the standard MMV problem, this configuration is effectively impossible. Indeed, for $\beta < 0.03$, we see that AMP-MMV is entirely failing at recovering the signal. However, once $\beta \approx 0.08$, we see that the NSER has dropped dramatically, as has the TNMSE. Once $\beta \geq 0.10$, AMP-MMV is performing almost to the level of the noise. As this experiment should hopefully convince the reader, even modest amounts of diversity in the measurement process can enable accurate reconstruction in operating environments that are otherwise impossible.

VII. CONCLUSION

In this work we introduced AMP-MMV, a Bayesian message passing algorithm for solving the MMV problem (1) when temporal correlation is present in the amplitudes of the non-zero signal coefficients. Our algorithm, which leverages Donoho, Maleki, and Montanari's AMP framework [25], performs rapid inference on high-dimensional MMV datasets. In order to establish a reference point for the quality of solutions obtained by AMP-MMV, we described and implemented the oracle-aided support-aware Kalman smoother (SKS). In numerical experiments, we found a range of problems over which AMP-MMV performed nearly as well as the SKS, despite the fact that AMP-MMV was given crude hyperparameter initializations that were refined from the data using an expectation-maximization algorithm. In comparing against two alternative Bayesian techniques, and one greedy technique, we found that AMP-MMV

offers an unrivaled performance-complexity tradeoff, particular in high-dimensional settings. We also demonstrated that substantial gains can be obtained in the MMV problem by incorporating diversity into the measurement process. Such diversity is particularly important in settings where the temporal correlation between coefficient amplitudes is substantial.

APPENDIX

In this appendix we summarize the procedure used to collapse the binary Gaussian mixture of (12) to a single Gaussian. For simplicity, we drop the n and (t) sub- and superscripts.

Let $\theta_r \triangleq \Re\{\theta\}$, let $\theta_i \triangleq \Im\{\theta\}$, and let ϕ_r and ϕ_i be defined similarly. Define

$$\begin{aligned}\tilde{g}(\theta_r, \theta_i) &\triangleq \nu_{f \rightarrow \theta}^{\text{mod}}(\theta_r + j\theta_i), \\ &= (1 - \Omega(\bar{\pi})) \mathcal{CN}(\theta_r + j\theta_i; \frac{1}{\varepsilon}\phi, \frac{1}{\varepsilon^2}c) + \Omega(\bar{\pi}) \mathcal{CN}(\theta_r + j\theta_i; \phi, c) \\ \tilde{f}(\theta_r, \theta_i) &\triangleq -\log \tilde{g}(\theta_r, \theta_i).\end{aligned}$$

Our objective is to approximate $\tilde{f}(\theta_r, \theta_i)$ using a two-dimensional second-order Taylor series expansion, $\check{f}(\theta_r, \theta_i)$, about the point ϕ :

$$\begin{aligned}\check{f}(\theta_r, \theta_i) &= \tilde{f}(\phi_r, \phi_i) + (\theta_r - \phi_r) \frac{\partial \tilde{f}}{\partial \theta_r} + (\theta_i - \phi_i) \frac{\partial \tilde{f}}{\partial \theta_i} \\ &\quad + \frac{1}{2} \left[(\theta_r - \phi_r)^2 \frac{\partial^2 \tilde{f}}{\partial \theta_r^2} + (\theta_r - \phi_r)(\theta_i - \phi_i) \frac{\partial^2 \tilde{f}}{\partial \theta_r \partial \theta_i} + (\theta_i - \phi_i)^2 \frac{\partial^2 \tilde{f}}{\partial \theta_i^2} \right],\end{aligned}$$

with all partial derivatives evaluated at ϕ . It can be shown that, for Taylor series expansions about the point ϕ , $\frac{\partial^2 \tilde{f}}{\partial \theta_r \partial \theta_i} = \mathcal{O}(\varepsilon^2)$ and $\left| \frac{\partial^2 \tilde{f}}{\partial \theta_r^2} - \frac{\partial^2 \tilde{f}}{\partial \theta_i^2} \right| = \mathcal{O}(\varepsilon^2)$. Since $\varepsilon \ll 1$, it is reasonable to therefore adopt a further approximation and assume $\frac{\partial^2 \tilde{f}}{\partial \theta_r \partial \theta_i} = 0$ and $\frac{\partial^2 \tilde{f}}{\partial \theta_r^2} = \frac{\partial^2 \tilde{f}}{\partial \theta_i^2}$. With this approximation, note that

$$\exp(-\check{f}(\theta_r, \theta_i)) \propto \mathcal{CN}(\theta_r + j\theta_i; \vec{\xi}, \vec{\psi}),$$

with

$$\vec{\psi} \triangleq 2 \frac{\partial^2 \tilde{f}}{\partial \theta_r^2}^{-1}, \tag{16}$$

$$\vec{\xi} \triangleq \phi_r + j\phi_i - \frac{\vec{\psi}}{2} \left(\frac{\partial \tilde{f}}{\partial \theta_r} + j \frac{\partial \tilde{f}}{\partial \theta_i} \right). \tag{17}$$

The pseudocode function, `taylor_approx`, that computes (16), (17) given the parameters of $\nu_{f \rightarrow \theta}^{\text{mod}}(\cdot)$ is provided in Table III.

REFERENCES

- [1] E. J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inform. Theory*, vol. 52, pp. 489 – 509, Feb. 2006.
- [2] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inform. Theory*, vol. 52, pp. 1289–1306, Apr. 2006.
- [3] E. J. Candes and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, pp. 21–30, Mar. 2008.
- [4] B. D. Rao and K. Kreutz-Delgado, "Sparse solutions to linear inverse problems with multiple measurement vectors," in *IEEE Digital Signal Process. Workshop*, (Bryce Canyon, UT), pp. 1–4, June 1998.
- [5] S. F. Cotter, B. D. Rao, K. Engan, and K. Kreutz-Delgado, "Sparse solutions to linear inverse problems with multiple measurement vectors," *IEEE Trans. Signal Process.*, vol. 53, pp. 2477–2488, Jul. 2005.
- [6] G. Tzagkarakis, D. Miliotis, and P. Tsakalides, "Multiple-measurement Bayesian compressed sensing using GSM priors for DOA estimation," in *IEEE Int'l Conf. Acoust., Speech & Signal Process. (ICASSP)*, (Dallas, TX), pp. 2610–2613, Mar. 2010.
- [7] D. Liang, L. Ying, and F. Liang, "Parallel MRI acceleration using M-FOCUSS," in *Int'l Conf. Bioinformatics & Biomed. Eng. (ICBBE)*, (Beijing, China), pp. 1–4, June 2009.
- [8] Y. Eldar and H. Rauhut, "Average case analysis of multichannel sparse recovery using convex relaxation," *IEEE Trans. Inform. Theory*, vol. 56, pp. 505–519, Jan. 2010.
- [9] Y. Eldar and M. Mishali, "Robust recovery of signals from a structured union of subspaces," *IEEE Trans. Inform. Theory*, vol. 55, pp. 5302–5316, Nov. 2009.
- [10] J. A. Tropp, A. C. Gilbert, and M. J. Strauss, "Algorithms for simultaneous sparse approximation. Part II: Convex relaxation," *Signal Processing*, vol. 86, pp. 589–602, Apr. 2006.
- [11] M. M. Hyder and K. Mahata, "A robust algorithm for joint sparse recovery," *IEEE Signal Process. Lett.*, vol. 16, pp. 1091–1094, Dec. 2009.
- [12] Z. Zhang and B. D. Rao, "Iterative reweighted algorithms for sparse signal recovery with temporally correlated source vectors," in *IEEE Int'l Conf. Acoust., Speech & Signal Process. (ICASSP)*, (Prague, Czech Republic), pp. 3932–3935, May 2011.
- [13] J. A. Tropp, A. C. Gilbert, and M. J. Strauss, "Algorithms for simultaneous sparse approximation. Part I: Greedy pursuit," *Signal Processing*, vol. 86, pp. 572–588, Apr. 2006.
- [14] K. Lee, Y. Bresler, and M. Junge, "Subspace methods for joint sparse recovery." arXiv:1004.3071, Apr. 2010.
- [15] D. P. Wipf and B. D. Rao, "An empirical Bayesian strategy for solving the simultaneous sparse approximation problem," *IEEE Trans. Signal Process.*, vol. 55, pp. 3704–3716, Jul. 2007.
- [16] J. Kim, W. Chang, B. Jung, D. Baron, and J. C. Ye, "Belief propagation for joint sparse recovery." arXiv:1102.3289v1 [cs.IT], Feb. 2011.
- [17] S. Shedthikere and A. Chockalingam, "Bayesian framework and message passing for joint support and signal recovery of approximately sparse signals," in *IEEE Int'l Conf. Acoust., Speech & Signal Process. (ICASSP)*, (Prague, Czech Republic), pp. 4032–4035, May 2011.
- [18] Z. Zhang and B. D. Rao, "Sparse signal recovery with temporally correlated source vectors using Sparse Bayesian Learning," *IEEE J. Selected Topics Signal Process.*, vol. 5, pp. 912–926, Sept. 2011.
- [19] N. Vaswani, "Kalman filtered compressed sensing," in *IEEE Int'l Conf. Image Process. (ICIP)*, pp. 893 –896, 12-15 2008.

- [20] D. Angelosante, G. B. Giannakis, and E. Grossi, "Compressed sensing of time-varying signals," in *Int'l Conf. Digital Signal Process.*, pp. 1–8, July 2009.
- [21] D. Angelosante, S. Roumeliotis, and G. Giannakis, "Lasso-Kalman smoother for tracking sparse signals," in *Asilomar Conf. Signals, Syst., & Comput.*, (Pacific Grove, CA), pp. 181 – 185, Nov. 2009.
- [22] N. Vaswani and W. Lu, "Modified-CS: Modifying compressive sensing for problems with partially known support," *IEEE Trans. Signal Process.*, vol. 58, pp. 4595–4607, Sept. 2010.
- [23] J. Ziniel, L. C. Potter, and P. Schniter, "Tracking and smoothing of time-varying sparse signals via approximate belief propagation," in *Asilomar Conf. Signals, Syst., & Comput.*, (Pacific Grove, CA), pp. 808–812, Nov. 2010.
- [24] E. van den Berg and M. P. Friedlander, "Theoretical and empirical results for recovery from multiple measurements," *IEEE Trans. Inform. Theory*, vol. 56, pp. 2516–2527, May 2010.
- [25] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing," in *Proc. Nat'l Acad. Sciences*, vol. 106, pp. 18914–18919, Nov. 2009.
- [26] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing: I. Motivation and construction," in *Proc. Inform. Theory Workshop*, pp. 1–5, Jan. 2010.
- [27] D. L. Donoho, I. Johnstone, and A. Montanari, "Accurate prediction of phase transitions in compressed sensing via a connection to minimax denoising." arXiv:1111.1041v1 [cs.IT], Nov. 2011.
- [28] 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, Jun. 2010.
- [29] M. A. Ohliger and D. K. Sodickson, "An introduction to coil array design for parallel MRI," *NMR in Biomed.*, vol. 19, pp. 300–315, May 2006.
- [30] P. Schniter, "Turbo reconstruction of structured sparse signals," in *Conf. on Inform. Sciences & Syst. (CISS)*, Mar. 2010.
- [31] R. L. Eubank, *A Kalman Filter Primer*. Boca Raton, FL: Chapman & Hall/CRC, 2006.
- [32] F. R. Kschischang, B. J. Frey, and H. A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inform. Theory*, vol. 47, pp. 498–519, Feb. 2001.
- [33] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan Kaufman, 1988.
- [34] B. J. Frey and D. J. C. MacKay, "A revolution: Belief propagation in graphs with cycles," in *Adv. in Neural Inform. Process. Sys.* (M. Jordan, M. S. Kearns, and S. A. Solla, eds.), MIT Press, 1998.
- [35] W. T. Freeman, E. C. Pasztor, and O. T. Carmichael, "Learning low-level vision," *Int'l J. Comput. Vision*, vol. 40, pp. 25–47, Oct. 2000.
- [36] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms*. New York: Cambridge University Press, 2003.
- [37] D. Baron, S. Sarvotham, and R. G. Baraniuk, "Bayesian compressive sensing via belief propagation," *IEEE Trans. Signal Process.*, vol. 58, pp. 269 –280, Jan. 2010.
- [38] S. Rangan, "Estimation with random linear mixing, belief propagation, and compressed sensing." arXiv:1001.2228v2 [cs.IT], 2010.
- [39] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing." arXiv:1010.5141v1 [cs.IT], 2010.
- [40] M. Bayati and A. Montanari, "The dynamics of message passing on dense graphs, with applications to compressed sensing," *IEEE Trans. Inform. Theory*, vol. 57, pp. 764–785, Feb. 2011.
- [41] D. L. Alspach and H. W. Sorenson, "Nonlinear Bayesian estimation using Gaussian sum approximations," *IEEE Trans. Autom. Control*, vol. 17, pp. 439–448, Aug. 1972.

- [42] D. Barber and A. T. Cemgil, “Graphical models for time series,” *IEEE Signal Process. Mag.*, vol. 27, pp. 18–28, Nov. 2010.
- [43] A. P. Dempster, N. M. Laird, and D. B. Rubin, “Maximum likelihood from incomplete data via the EM algorithm,” *J. Roy. Statist. Soc., B*, vol. 39, pp. 1–38, 1977.
- [44] T. K. Moon, “The expectation-maximization algorithm,” *IEEE Signal Process. Mag.*, vol. 13, pp. 47–60, Nov. 1996.
- [45] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York: Springer-Verlag, 2006.
- [46] Y. Weiss and W. T. Freeman, “Correctness of belief propagation in Gaussian graphical models of arbitrary topology,” *Neural Computation*, vol. 13, pp. 2173–2200, Oct. 2001.



Justin Ziniel is a Ph.D. student in the Information Processing Systems Lab of the Department of Electrical and Computer Engineering at The Ohio State University (OSU). He received the B.S. and M.S. degrees in electrical and computer engineering from OSU in 2007 and 2012, respectively. His research interests include statistical signal processing, machine learning, and interdisciplinary “big data” inference problems.



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, Ithaca, NY, in 2000. From 1993 to 1996, he was a Systems Engineer with Tektronix Inc., Beaverton, OR. After receiving the Ph.D. degree, in 2000, he joined the Department of Electrical and Computer Engineering at The Ohio State University, Columbus, where he is currently an Associate Professor and a Member of the Information Processing Systems (IPS) Lab. In 2008–2009, he was a Visiting Professor at Eurecom, Sophia Antipolis, France, and Supelec, Gif-sur-Yvette, France. His areas of interest include statistical signal processing, wireless communications and networks, and machine learning. Dr. Schniter received the National Science Foundation CAREER Award, in 2003.