

Bilinear Generalized Approximate Message Passing—Part I: Derivation

Jason T. Parker, *Member, IEEE*, Philip Schniter, *Fellow, IEEE*, and Volkan Cevher, *Senior Member, IEEE*

Abstract—In this paper, we extend the generalized approximate message passing (G-AMP) approach, originally proposed for high-dimensional generalized-linear regression in the context of compressive sensing, to the generalized-bilinear case, which enables its application to matrix completion, robust PCA, dictionary learning, and related matrix-factorization problems. Here, in Part I of a two-part paper, we derive our Bilinear G-AMP (BiG-AMP) algorithm as an approximation of the sum-product belief propagation algorithm in the high-dimensional limit, where central-limit theorem arguments and Taylor-series approximations apply, and under the assumption of statistically independent matrix entries with known priors. In addition, we propose an adaptive damping mechanism that aids convergence under finite problem sizes, an expectation-maximization (EM)-based method to automatically tune the parameters of the assumed priors, and two rank-selection strategies. In Part II of the paper, we will discuss the specializations of EM-BiG-AMP to the problems of matrix completion, robust PCA, and dictionary learning, and we will present the results of an extensive empirical study comparing EM-BiG-AMP to state-of-the-art algorithms on each problem.

Index Terms—Approximate message passing, belief propagation, bilinear estimation, matrix completion, dictionary learning, robust principal components analysis, matrix factorization.

I. INTRODUCTION

IN this work, we present a new algorithmic framework for the following *generalized bilinear* inference problem: estimate the matrices $\mathbf{A} = [a_{mn}] \in \mathbb{R}^{M \times N}$ and $\mathbf{X} = [x_{nl}] \in \mathbb{R}^{N \times L}$ from a matrix observation $\mathbf{Y} \in \mathbb{R}^{M \times L}$ that is statistically coupled to their product, $\mathbf{Z} \triangleq \mathbf{AX}$. In doing so, we treat \mathbf{A} and \mathbf{X} as realizations of independent random matrices \mathbf{A} and

\mathbf{X} with known separable pdfs (or pmfs in the case of discrete models)

$$p_{\mathbf{A}}(\mathbf{A}) = \prod_m \prod_n p_{a_{mn}}(a_{mn}) \quad (1)$$

$$p_{\mathbf{X}}(\mathbf{X}) = \prod_n \prod_l p_{x_{nl}}(x_{nl}), \quad (2)$$

and we likewise assume that the likelihood function of \mathbf{Z} is known and separable, i.e.,

$$p_{\mathbf{Y}|\mathbf{Z}}(\mathbf{Y}|\mathbf{Z}) = \prod_m \prod_l p_{y_{ml}|z_{ml}}(y_{ml}|z_{ml}). \quad (3)$$

Recently, various special cases of this problem have gained the intense interest of the research community, e.g.,

- 1) *Matrix Completion*: In this problem, one observes a few (possibly noise-corrupted) entries of a low-rank matrix and the goal is to infer the missing entries. In our framework, $\mathbf{Z} = \mathbf{AX}$ would represent the complete low-rank matrix (with tall \mathbf{A} and wide \mathbf{X}) and $p_{y_{ml}|z_{ml}}$ the observation mechanism, which would be (partially) informative about z_{ml} at the observed entries $(m, l) \in \Omega$ and non-informative at the missing entries $(m, l) \notin \Omega$.
- 2) *Robust PCA*: Here, the objective is to recover a low-rank matrix (or its principal components) observed in the presence of noise and sparse outliers. In our framework, $\mathbf{Z} = \mathbf{AX}$ could again represent the low-rank matrix, and $p_{y_{ml}|z_{ml}}$ the noise-and-outlier-corrupted observation mechanism. Alternatively, \mathbf{X} could also capture the outliers, as described in the sequel.
- 3) *Dictionary Learning*: Here, the objective is to learn a dictionary \mathbf{A} for which there exists a sparse data representation \mathbf{X} such that \mathbf{AX} closely matches the observed data \mathbf{Y} . In our framework, $\{p_{x_{nl}}\}$ would be chosen to induce sparsity, $\mathbf{Z} = \mathbf{AX}$ would represent the noiseless observations, and $\{p_{y_{ml}|z_{ml}}\}$ would model the (possibly noisy) observation mechanism.

While a plethora of approaches to these problems have been proposed based on optimization techniques (e.g., [5]–[15]), greedy methods (e.g., [16]–[20]), Bayesian sampling methods (e.g., [21], [22]), variational methods (e.g., [23]–[27]), and discrete message passing (e.g., [28]), ours is based on the *Approximate Message Passing* (AMP) framework, an instance of loopy belief propagation (LBP) [29] that was recently developed to tackle *linear* [30]–[32] and *generalized linear* [33] inference problems encountered in the context of compressive sensing (CS). In the generalized-linear CS problem, one estimates $\mathbf{x} \in \mathbb{R}^N$ from observations $\mathbf{y} \in \mathbb{R}^M$ that are statistically coupled to the transform outputs $\mathbf{z} = \mathbf{Ax}$ through a separable

Manuscript received December 09, 2013; revised June 05, 2014; accepted August 19, 2014. Date of publication September 12, 2014; date of current version October 22, 2014. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Ivan W. Selesnick. The work of J. T. Parker on this project has been supported by AFOSR Lab Task 11RY02COR. The work of P. Schniter on this project has been supported by NSF grants IIP-0968910, CCF-1018368, CCF-1218754, and by DARPA/ONR grant N66001-10-1-4090. The work of V. Cevher was supported in part by the European Commission under Grant MIRG-268398, ERC Future Proof, SNF 200021-132548, SNF 200021-146750 and SNF CRSII2-147633. Portions of this work were presented at the Workshop on Signal Processing with Adaptive Sparse Structured Representations, Edinburgh, Scotland, 2011 [2], the Information Theory and Applications Workshop, La Jolla, CA, USA, 2012 [3], and at the Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, USA, 2012 [4].

J. T. Parker is with the Sensors Directorate, Air Force Research Laboratory, Dayton, OH 45433 USA (e-mail: jason.parker.13@us.af.mil).

P. Schniter is with The Ohio State University, Columbus OH 43210 USA (e-mail: schniter@ece.osu.edu).

V. Cevher is with the École Polytechnique Fédérale de Lausanne, Lausanne 1015, Switzerland (e-mail: volkan.cevher@epfl.ch).

Digital Object Identifier 10.1109/TSP.2014.2357776

likelihood function $p_{\mathbf{y}|\mathbf{z}}(\mathbf{y}|\mathbf{z})$, where in this case the transform \mathbf{A} is *fixed and known*.

In the context of CS, the AMP framework yields algorithms with remarkable properties: i) solution trajectories that, in the large-system limit (i.e., as $M, N \rightarrow \infty$ with M/N fixed, under iid sub-Gaussian \mathbf{A}) are governed by a state-evolution whose fixed points—when unique—yield the true posterior means [34], [35] and ii) a low implementation complexity (i.e., dominated by one multiplication with \mathbf{A} and \mathbf{A}^\top per iteration, and relatively few iterations) [32]. Thus, a natural question is whether the AMP framework can be successfully applied to the generalized *bilinear* problem described earlier.

In this manuscript, which is Part I of a two-part work, we propose an AMP-based approach to generalized bilinear inference that we henceforth refer to as *Bilinear Generalized AMP* (BiG-AMP), and we uncover special cases under which the general approach can be simplified. In addition, we propose an adaptive damping [36] mechanism, an expectation-maximization (EM)-based [37] method of tuning the parameters of $p_{a_{mn}}$, $p_{x_{nl}}$, and $p_{y_{ml}|z_{ml}}$ (in case they are unknown), and methods to select the rank N (in case it is unknown). In the case that $p_{a_{mn}}$, $p_{x_{nl}}$, and/or $p_{y_{ml}|z_{ml}}$ are completely unknown, they can be modeled as Gaussian-mixtures with mean/variance/weight parameters learned via EM [38]. In Part II [1], we detail the application of BiG-AMP to matrix completion, robust PCA, and dictionary learning, and present the results of an extensive numerical investigation into the performance of BiG-AMP in each application. Those empirical results demonstrate that BiG-AMP yields an excellent combination of estimation accuracy and runtime when compared to existing state-of-the-art algorithms for each application.

Although the AMP methodology is itself restricted to separable known pdfs (1)–(3), the results of Part II suggest that this limitation is not an issue for many practical problems of interest. However, in problems where the separability assumption is too constraining, it can be relaxed through the use of hidden (coupling) variables, as originally proposed in the context of “turbo-AMP” [39] and applied to BiG-AMP in [40]. Due to space limitations, however, this approach will not be discussed here. Finally, although we focus on real-valued random variables, all of the methodology described in this work can be easily extended to circularly symmetric complex-valued random variables.

We now discuss related work. One possibility of applying AMP methods to matrix completion was suggested by Montanari in [32, Sec. 9.7.3] but the approach described there differs from BiG-AMP in that it was i) constructed from a factor graph with *vector-valued* variables and ii) restricted to the (incomplete) additive white Gaussian noise (AWGN) observation model. Moreover, no concrete algorithm nor performance evaluation was reported. Since we first reported on BiG-AMP in [2], [3], Rangan and Fletcher [41] proposed an AMP-based approach for the estimation of *rank-one* matrices from AWGN-corrupted observations, and showed that it can be characterized by a state evolution in the large-system limit. More recently, Krzakala, Mézard, and Zdeborová [42] proposed an AMP-based approach to blind calibration and dictionary learning in AWGN that bears similarity to a special case of BiG-AMP, and derived a state-evolution using the cavity method. Their method,

however, was not numerically successful in solving dictionary learning problems [42]. The BiG-AMP algorithm that we derive here is a generalization of those in [41], [42] in that it handles *generalized* bilinear observations rather than AWGN-corrupted ones. Moreover, our work is the first to detail adaptive damping, parameter tuning, and rank-selection mechanisms for AMP based bilinear inference, and it is the first to present an in-depth numerical investigation involving both synthetic and real-world datasets. An application/extension of the BiG-AMP algorithm described here to hyperspectral unmixing (an instance of non-negative matrix factorization) was recently proposed in [40].

The remainder of the document is organized as follows. Section II derives the BiG-AMP algorithm, and Section III presents several special-case simplifications of BiG-AMP. Section IV describes the adaptive damping mechanism, and Section V the EM-based tuning of prior parameters and selection of rank N . Finally, Section VI concludes. In Part II [1], we detail the application of BiG-AMP to matrix completion, robust PCA, and dictionary learning, and present the results of an extensive numerical investigation into the performance of BiG-AMP in each application.

Notation: Throughout, we use san-serif font (e.g., \mathbf{x}) for random variables and serif font (e.g., x) otherwise. We use boldface capital letters (e.g., \mathbf{X} and \mathbf{X}) for matrices, boldface small letters (e.g., \mathbf{x} and \mathbf{x}) for vectors, and non-bold small letters (e.g., x and x) for scalars. We then use $p_x(x)$ to denote the pdf of random quantity x , and $\mathcal{N}(x; \hat{x}, \nu^x)$ to denote the Gaussian pdf for a scalar random variable with mean \hat{x} and variance ν^x . Also, we use $E\{x\}$ and $\text{var}\{x\}$ to denote mean and variance of x , respectively, and $D(p_1||p_2)$ for the Kullback-Leibler (KL) divergence between pdfs p_1 and p_2 . For a matrix \mathbf{X} , we use $x_{nl} = [\mathbf{X}]_{nl}$ to denote the entry in the n^{th} row and l^{th} column, $\|\mathbf{X}\|_F$ to denote the Frobenius norm, and \mathbf{X}^\top to denote transpose. Similarly, we use x_n to denote the n^{th} entry in vector \mathbf{x} and $\|\mathbf{x}\|_p = (\sum_n |x_n|^p)^{1/p}$ to denote the ℓ_p norm.

II. BILINEAR GENERALIZED AMP

A. Problem Formulation

For the statistical model (1)–(3), the posterior distribution is

$$p_{\mathbf{X}, \mathbf{A} | \mathbf{Y}}(\mathbf{X}, \mathbf{A} | \mathbf{Y}) = p_{\mathbf{Y} | \mathbf{X}, \mathbf{A}}(\mathbf{Y} | \mathbf{X}, \mathbf{A}) p_{\mathbf{X}}(\mathbf{X}) p_{\mathbf{A}}(\mathbf{A}) / p_{\mathbf{Y}}(\mathbf{Y}) \quad (4)$$

$$\propto p_{\mathbf{Y} | \mathbf{Z}}(\mathbf{Y} | \mathbf{A}\mathbf{X}) p_{\mathbf{X}}(\mathbf{X}) p_{\mathbf{A}}(\mathbf{A}) \quad (5)$$

$$= \left[\prod_m \prod_l p_{y_{ml}|z_{ml}} \left(y_{ml} \left| \sum_k a_{mk} x_{kl} \right. \right) \right] \times \left[\prod_n \prod_l p_{x_{nl}}(x_{nl}) \right] \left[\prod_m \prod_n p_{a_{mn}}(a_{mn}) \right], \quad (6)$$

where (4) employs Bayes’ rule and \propto denotes equality up to a constant scale factor.

The posterior distribution can be represented with a factor graph, as depicted in Fig. 1. There, the factors of $p_{\mathbf{X}, \mathbf{A} | \mathbf{Y}}$ from (6) are represented by “factor nodes” that appear as black boxes, and the random variables are represented by “variable nodes”




Fig. 1. The factor graph for generalized bilinear inference for (toy-sized) problem dimensions $M = 4$, $L = 3$, and $N = 2$.

that appear as white circles. Each variable node is connected to every factor node in which that variable appears. The observed data $\{y_{ml}\}$ are treated as parameters of the $p_{y_{ml}|z_{ml}}$ factor nodes in the middle of the graph, and not as random variables. The structure of Fig. 1 becomes intuitive when recalling that $\mathbf{Z} = \mathbf{AX}$ implies $z_{ml} = \sum_{n=1}^N a_{mn}x_{nl}$.

B. Loopy Belief Propagation

In this work, we aim to compute minimum mean-squared error (MMSE) estimates of \mathbf{X} and \mathbf{A} , i.e., the means¹ of the marginal posteriors $p_{x_{nl}|\mathbf{Y}}(\cdot|\mathbf{Y})$ and $p_{a_{mn}|\mathbf{Y}}(\cdot|\mathbf{Y})$, for all pairs (n, l) and (m, n) . Although exact computation of these quantities is generally prohibitive, they can be efficiently approximated using loopy belief propagation (LBP) [29].

In LBP, beliefs about the random variables (in the form of pdfs or log pdfs) are propagated among the nodes of the factor graph until they converge. The standard way to compute these beliefs, known as the *sum-product algorithm* (SPA) [43], [44], stipulates that the belief emitted by a variable node along a given edge of the graph is computed as the product of the incoming beliefs from all other edges, whereas the belief emitted by a factor node along a given edge is computed as the integral of the product of the factor associated with that node and the incoming beliefs on all other edges. The product of all beliefs impinging on a given variable node yields the posterior pdf for that variable. In cases where the factor graph has no loops, exact marginal posteriors result from two (i.e., forward and backward) passes of the SPA [43], [44]. For loopy factor graphs, exact inference is in general NP hard [45] and so LBP does not guarantee correct posteriors. That said, LBP has shown state-of-the-art performance in many applications, such as inference on Markov random fields [46], turbo decoding [47], LDPC decoding [48], multiuser detection [49], and compressive sensing [30], [31], [33]–[35].

In high-dimensional inference problems, exact implementation of the SPA is impractical, motivating approximations of the

¹Another worthwhile objective could be to compute the joint MAP estimate $\arg \max_{\mathbf{X}, \mathbf{A}} p_{\mathbf{X}, \mathbf{A}|\mathbf{Y}}(\mathbf{X}, \mathbf{A}|\mathbf{Y})$; we leave this to future work.

TABLE I
SPA MESSAGE DEFINITIONS AT ITERATION $t \in \mathbb{Z}$

$\Delta_{m \rightarrow nl}^x(t, \cdot)$	SPA message from node $p_{y_{ml} z_{ml}}$ to node x_{nl}
$\Delta_{m \leftarrow nl}^x(t, \cdot)$	SPA message from node x_{nl} to node $p_{y_{ml} z_{ml}}$
$\Delta_{l \rightarrow mn}^a(t, \cdot)$	SPA message from node $p_{y_{ml} z_{ml}}$ to node a_{mn}
$\Delta_{l \leftarrow mn}^a(t, \cdot)$	SPA message from node a_{mn} to node $p_{y_{ml} z_{ml}}$
$\Delta_{nl}^x(t, \cdot)$	SPA-approximated log posterior pdf of x_{nl}
$\Delta_{mn}^a(t, \cdot)$	SPA-approximated log posterior pdf of a_{mn}

SPA. A notable example is the *generalized approximate message passing* (GAMP) algorithm, developed in [33] to solve the generalized CS problem, which exploits the “blessings of dimensionality” that arise when \mathbf{A} is a sufficiently large and dense and which was rigorously analyzed in [35]. In the sequel, we derive an algorithm for the generalized bilinear inference BiG-AMP algorithm that employs GAMP-like approximations to the SPA on the factor graph in Fig. 1. As we shall see, the approximations are primarily based on central-limit-theorem (CLT) and Taylor-series arguments.

C. Sum-Product Algorithm

In our formulation of the SPA, messages take the form of log-pdfs with arbitrary constant offsets. For example, the iteration- t (where $t \in \mathbb{Z}$) message $\Delta_{m \rightarrow nl}^x(t, \cdot)$ can be converted to the pdf $\frac{1}{C} \exp(\Delta_{m \rightarrow nl}^x(t, \cdot))$, where the choice of scale factor $C = \int_{x_{nl}} \exp(\Delta_{m \rightarrow nl}^x(t, x_{nl}))$ ensures that the pdf integrates to one. Four types of message will be used, as specified in Table I. We also find it convenient to express the (iteration- t SPA-approximated) posterior pdfs $p_{x_{nl}|\mathbf{Y}}(t, \cdot|\mathbf{Y})$ and $p_{a_{mn}|\mathbf{Y}}(t, \cdot|\mathbf{Y})$ in the log domain as $\Delta_{nl}^x(t, \cdot)$ and $\Delta_{mn}^a(t, \cdot)$, respectively, again with arbitrary constant offsets.

Applying the SPA to the factor graph in Fig. 1, we arrive at the following update rules for the four messages in Table I.

$$\begin{aligned} & \Delta_{m \rightarrow nl}^x(t, x_{nl}) \\ &= \log \int_{\{a_{mk}\}_{k=1}^N, \{x_{rl}\}_{r \neq n}} p_{y_{ml}|z_{ml}}(y_{ml} | \sum_{k=1}^N a_{mk}x_{kl}) \\ & \quad \times \prod_{r \neq n} \exp(\Delta_{m \leftarrow rl}^x(t, x_{rl})) \prod_{k=1}^N \exp(\Delta_{l \leftarrow mk}^a(t, a_{mk})) \\ & \quad + \text{const} \end{aligned} \quad (7)$$

$$\begin{aligned} & \Delta_{m \leftarrow nl}^x(t+1, x_{nl}) \\ &= \log p_{x_{nl}}(x_{nl}) + \sum_{k \neq m} \Delta_{k \rightarrow nl}^x(t, x_{nl}) + \text{const} \end{aligned} \quad (8)$$

$$\begin{aligned} & \Delta_{l \rightarrow mn}^a(t, a_{mn}) \\ &= \log \int_{\{a_{mr}\}_{r \neq n}, \{x_{kl}\}_{k=1}^N} p_{y_{ml}|z_{ml}}(y_{ml} | \sum_{k=1}^N a_{mk}x_{kl}) \\ & \quad \times \prod_{k=1}^N \exp(\Delta_{m \leftarrow kl}^x(t, x_{kl})) \prod_{r \neq n} \exp(\Delta_{l \leftarrow mr}^a(t, a_{mr})) \\ & \quad + \text{const} \end{aligned} \quad (9)$$

$$\begin{aligned} & \Delta_{l \leftarrow mn}^a(t+1, a_{mn}) \\ &= \log p_{a_{mn}}(a_{mn}) + \sum_{k \neq l} \Delta_{k \rightarrow mn}^a(t, a_{mn}) + \text{const}, \end{aligned} \quad (10)$$

TABLE II
BIG-AMP VARIABLE SCALINGS IN THE LARGE-SYSTEM LIMIT

$\widehat{z}_{ml}(t)$	$O(1)$	$\nu_{ml}^z(t)$	$O(1)$	$\widehat{x}_{m,nl}(t) - \widehat{x}_{nl}(t)$	$O(\frac{1}{\sqrt{N}})$
$\widehat{x}_{m,nl}(t)$	$O(1)$	$\nu_{m,nl}^x(t)$	$O(1)$	$\widehat{x}_{m,nl}^2(t) - \widehat{x}_{nl}^2(t)$	$O(\frac{1}{\sqrt{N}})$
$\widehat{x}_{nl}(t)$	$O(1)$	$\nu_{nl}^x(t)$	$O(1)$	$\nu_{m,nl}^x(t) - \nu_{nl}^x(t)$	$O(\frac{1}{\sqrt{N}})$
$\widehat{a}_{l,mn}(t)$	$O(\frac{1}{\sqrt{N}})$	$\nu_{l,mn}^a(t)$	$O(\frac{1}{\sqrt{N}})$	$\widehat{a}_{l,mn}(t) - \widehat{a}_{mn}(t)$	$O(\frac{1}{\sqrt{N}})$
$\widehat{a}_{mn}(t)$	$O(\frac{1}{\sqrt{N}})$	$\nu_{mn}^a(t)$	$O(\frac{1}{\sqrt{N}})$	$\widehat{a}_{l,mn}^2(t) - \widehat{a}_{mn}^2(t)$	$O(\frac{1}{N^{3/2}})$
$\widehat{p}_{ml}(t)$	$O(1)$	$\nu_{ml}^p(t)$	$O(1)$	$\nu_{l,mn}^a(t) - \nu_{mn}^a(t)$	$O(\frac{1}{N^{3/2}})$
$\widehat{r}_{m,nl}(t)$	$O(1)$	$\nu_{m,nl}^r(t)$	$O(1)$	$\widehat{r}_{m,nl}(t) - \widehat{r}_{nl}(t)$	$O(\frac{1}{\sqrt{N}})$
$\widehat{r}_{nl}(t)$	$O(1)$	$\nu_{nl}^r(t)$	$O(1)$	$\nu_{m,nl}^r(t) - \nu_{nl}^r(t)$	$O(\frac{1}{\sqrt{N}})$
$\widehat{q}_{l,mn}(t)$	$O(\frac{1}{\sqrt{N}})$	$\nu_{l,mn}^q(t)$	$O(\frac{1}{\sqrt{N}})$	$\widehat{q}_{l,mn}(t) - \widehat{q}_{mn}(t)$	$O(\frac{1}{\sqrt{N}})$
$\widehat{q}_{mn}(t)$	$O(\frac{1}{\sqrt{N}})$	$\nu_{mn}^q(t)$	$O(\frac{1}{\sqrt{N}})$	$\nu_{l,mn}^q(t) - \nu_{mn}^q(t)$	$O(\frac{1}{N^2})$
$\widehat{s}_{ml}(t)$	$O(1)$	$\nu_{ml}^s(t)$	$O(1)$		

where *const* is an arbitrary constant (w.r.t x_{nl} in (7) and (8), and w.r.t a_{mn} in (9) and (10)). In the sequel, we denote the mean and variance of the pdf $\frac{1}{C} \exp(\Delta_m^x \leftarrow_{nl}(t, \cdot))$ by $\widehat{x}_{m,nl}(t)$ and $\nu_{m,nl}^x(t)$, respectively, and we denote the mean and variance of $\frac{1}{C} \exp(\Delta_l^a \leftarrow_{mn}(t, \cdot))$ by $\widehat{a}_{l,mn}(t)$ and $\nu_{l,mn}^a(t)$. For the log-posteriors, the SPA implies

$$\begin{aligned} & \Delta_{nl}^x(t+1, x_{nl}) \\ &= \log p_{x_{nl}}(x_{nl}) + \sum_{m=1}^M \Delta_{m \rightarrow nl}^x(t, x_{nl}) + \text{const} \quad (11) \end{aligned}$$

$$\begin{aligned} & \Delta_{mn}^a(t+1, a_{mn}) \\ &= \log p_{a_{mn}}(a_{mn}) + \sum_{l=1}^L \Delta_{l \rightarrow mn}^a(t, a_{mn}) + \text{const}, \quad (12) \end{aligned}$$

and we denote the mean and variance of $\frac{1}{C} \exp(\Delta_{nl}^x(t, \cdot))$ by $\widehat{x}_{nl}(t)$ and $\nu_{nl}^x(t)$, and the mean and variance of $\frac{1}{C} \exp(\Delta_{mn}^a(t, \cdot))$ by $\widehat{a}_{mn}(t)$ and $\nu_{mn}^a(t)$.

D. Approximated Factor-to-Variable Messages

We now apply AMP approximations to the SPA updates (7)–(12). As we shall see, the approximations are based primarily on central-limit-theorem (CLT) and Taylor-series arguments that become exact in the large-system limit, where $M, L, N \rightarrow \infty$ with fixed ratios M/N and L/N . (Due to the use of finite M, L, N in practice, we still regard them as approximations.) In particular, our derivation will neglect terms that vanish relative to others as $N \rightarrow \infty$, which requires that we establish certain scaling conventions. First, we assume w.l.o.g.² that $E\{z_{ml}^2\}$ and $E\{x_{nl}^2\}$ scale as $O(1)$, i.e., that the magnitudes of these elements stay finite as $N \rightarrow \infty$, and that $E\{a_{mn}\} = 0$. In this case, the relationship $z_{ml} = \sum_{n=1}^N a_{mn} x_{nl}$ implies that $E\{a_{mn}^2\}$ must scale as $O(1/N)$. These scalings are assumed to hold for random variables z_{ml} , a_{mn} , and x_{ml} distributed according to the prior pdfs, according to the pdfs corresponding to the SPA messages (7)–(10), and according to the pdfs corresponding to the SPA posterior approximations (11)–(12). These assumptions lead straightforwardly to the scalings of $\widehat{z}_{ml}(t)$, $\nu_{ml}^z(t)$, $\widehat{x}_{m,nl}(t)$, $\nu_{m,nl}^x(t)$, $\widehat{x}_{nl}(t)$, $\nu_{nl}^x(t)$, $\widehat{a}_{l,mn}(t)$, $\nu_{l,mn}^a(t)$, $\widehat{a}_{mn}(t)$, and $\nu_{mn}^a(t)$ specified in Table II. Furthermore, because $\Delta_{m \rightarrow nl}^x(t, \cdot)$ and $\Delta_{nl}^x(t, \cdot)$ differ by only one

term out of M , it is reasonable to assume [32], [33] that the corresponding difference in means $\widehat{x}_{m,nl}(t) - \widehat{x}_{nl}(t)$ and variances $\nu_{m,nl}^x(t) - \nu_{nl}^x(t)$ are both $O(1/\sqrt{N})$, which then implies that $\widehat{x}_{m,nl}^2(t) - \widehat{x}_{nl}^2(t)$ is also $O(1/\sqrt{N})$. Similarly, because $\Delta_{l \rightarrow mn}^a(t, \cdot)$ and $\Delta_{mn}^a(t, \cdot)$ differ by only one term out of N , where $\widehat{a}_{l,mn}(t)$ and $\widehat{a}_{mn}(t)$ are $O(1/\sqrt{N})$, it is reasonable to assume that $\widehat{a}_{l,mn}(t) - \widehat{a}_{mn}(t)$ is $O(1/N)$ and that both $\nu_{l,mn}^a(t) - \nu_{mn}^a(t)$ and $\widehat{a}_{l,mn}^2(t) - \widehat{a}_{mn}^2(t)$ are $O(1/N^{3/2})$. The remaining entries in Table II will be explained below.

We start by approximating the message $\Delta_{m \rightarrow nl}^x(t, \cdot)$. Expanding (7), we find

$$\begin{aligned} & \Delta_{m \rightarrow nl}^x(t, x_{nl}) \\ &= \log \int_{\{a_{mk}\}_{k=1, \{x_{rl}\}_{r \neq n}}^N} p_{y_{ml}|z_{ml}}(y_{ml} | \overbrace{a_{mn}x_{nl} + \sum_{k=1, k \neq n}^N a_{mk}x_{kl}}^{z_{ml}}) \\ & \times \prod_{r \neq n} \exp(\Delta_m^x \leftarrow_{rl}(t, x_{rl})) \prod_{k=1}^N \exp(\Delta_l^a \leftarrow_{mk}(t, a_{mk})) \\ & + \text{const}. \quad (13) \end{aligned}$$

For large N , the CLT motivates the treatment of z_{ml} , the random variable associated with the z_{ml} identified in (13), conditioned on $x_{nl} = x_{nl}$, as Gaussian and thus completely characterized by a (conditional) mean and variance. Defining the zero-mean r.v.s $\widetilde{a}_{l,mn} \triangleq a_{mn} - \widehat{a}_{l,mn}(t)$ and $\widetilde{x}_{m,kl} = x_{kl} - \widehat{x}_{m,kl}(t)$, where $a_{mn} \sim \frac{1}{C} \exp(\Delta_l^a \leftarrow_{mn}(t, \cdot))$ and $x_{nl} \sim \frac{1}{C} \exp(\Delta_m^x \leftarrow_{nl}(t, \cdot))$, we can write

$$\begin{aligned} z_{ml} &= (\widehat{a}_{l,mn}(t) + \widetilde{a}_{l,mn})x_{nl} + \sum_{k \neq n} (\widehat{a}_{l,mk}(t)\widehat{x}_{m,kl}(t) \\ & + \widetilde{a}_{l,mk}(t)\widetilde{x}_{m,kl} + \widetilde{a}_{l,mk}\widehat{x}_{m,kl}(t) + \widetilde{a}_{l,mk}\widetilde{x}_{m,kl}) \quad (14) \end{aligned}$$

after which it is straightforward to see that

$$E\{z_{ml}|x_{nl} = x_{nl}\} = \widehat{a}_{l,mn}(t)x_{nl} + \widehat{p}_{n,ml}(t) \quad (15)$$

$$\text{var}\{z_{ml}|x_{nl} = x_{nl}\} = \nu_{l,mn}^a(t)x_{nl}^2 + \nu_{n,ml}^p(t) \quad (16)$$

for

$$\widehat{p}_{n,ml}(t) \triangleq \sum_{k \neq n} \widehat{a}_{l,mk}(t)\widehat{x}_{m,kl}(t) \quad (17)$$

$$\begin{aligned} \nu_{n,ml}^p(t) &\triangleq \sum_{k \neq n} (\widehat{a}_{l,mk}^2(t)\nu_{m,kl}^x(t) + \nu_{l,mk}^a(t)\widehat{x}_{m,kl}^2(t) \\ & + \nu_{l,mk}^a(t)\nu_{m,kl}^x(t)). \quad (18) \end{aligned}$$

With this conditional-Gaussian approximation, (13) becomes

$$\begin{aligned} & \Delta_{m \rightarrow nl}^x(t, x_{nl}) \approx \text{const} + \log \int_{z_{ml}} p_{y_{ml}|z_{ml}}(y_{ml}|z_{ml}) \\ & \times \mathcal{N}(z_{ml}; \widehat{a}_{l,mn}(t)x_{nl} + \widehat{p}_{n,ml}(t), \nu_{l,mn}^a(t)x_{nl}^2 + \nu_{n,ml}^p(t)) \\ & = H_{ml}(\widehat{a}_{l,mn}(t)x_{nl} + \widehat{p}_{n,ml}(t), \\ & \nu_{l,mn}^a(t)x_{nl}^2 + \nu_{n,ml}^p(t); y_{ml}) + \text{const} \quad (20) \end{aligned}$$

in terms of the function

$$H_{ml}(\widehat{q}, \nu^q; y) \triangleq \log \int_z p_{y_{ml}|z_{ml}}(y|z) \mathcal{N}(z; \widehat{q}, \nu^q). \quad (21)$$

²Other scalings on $E\{z_{ml}^2\}$, $E\{x_{nl}^2\}$, and $E\{a_{mn}^2\}$ could be used as long as they are consistent with the relationship $z_{ml} = \sum_{n=1}^N a_{mn} x_{nl}$.

Unlike the original SPA message (7), the approximation (20) requires only a single integration. Still, additional simplifications are possible. First, notice that $\widehat{p}_{n,ml}(t)$ and $\nu_{n,ml}^p(t)$ differ from the corresponding n -invariant quantities

$$\widehat{p}_{ml}(t) \triangleq \sum_{k=1}^N \widehat{a}_{l,mk}(t) \widehat{x}_{m,kl}(t) \quad (22)$$

$$\begin{aligned} \nu_{ml}^p(t) &\triangleq \sum_{k=1}^N (\widehat{a}_{l,mk}^2(t) \nu_{m,kl}^x(t) + \nu_{l,mk}^a(t) \widehat{x}_{m,kl}^2(t) \\ &\quad + \nu_{l,mk}^a(t) \nu_{m,kl}^x(t)) \end{aligned} \quad (23)$$

by one term. In the sequel, we will assume that $\widehat{p}_{ml}(t)$ and $\nu_{ml}^p(t)$ are $O(1)$ since these quantities can be recognized as the mean and variance, respectively, of an estimate of z_{ml} , which is $O(1)$. Writing the H_{ml} term in (20) using (22)–(23),

$$\begin{aligned} H_{ml} &\left(\widehat{a}_{l,mn}(t) x_{nl} + \widehat{p}_{n,ml}(t), \nu_{l,mn}^a(t) x_{nl}^2 + \nu_{n,ml}^p(t); y_{ml} \right) \\ &= H_{ml} \left(\widehat{a}_{l,mn}(t) (x_{nl} - \widehat{x}_{m,nl}(t)) + \widehat{p}_{ml}(t), \right. \\ &\quad \left. \nu_{l,mn}^a(t) (x_{nl}^2 - \widehat{x}_{m,nl}^2(t)) - \widehat{a}_{l,mn}^2(t) \nu_{m,nl}^x(t) \right. \\ &\quad \left. - \nu_{l,mn}^a(t) \nu_{m,nl}^x(t) + \nu_{ml}^p(t); y_{ml} \right) \end{aligned} \quad (24)$$

$$= H_{ml} \left(\widehat{a}_{l,mn}(t) (x_{nl} - \widehat{x}_{nl}(t)) + \widehat{p}_{ml}(t) + O(1/N), \right. \\ \left. \nu_{l,mn}^a(t) (x_{nl}^2 - \widehat{x}_{nl}^2(t)) + \nu_{ml}^p(t) + O(1/N); y_{ml} \right) \quad (25)$$

where in (25) we used the facts that $\widehat{a}_{l,mn}(t)(\widehat{x}_{nl}(t) - \widehat{x}_{m,nl}(t))$ and $\nu_{l,mn}^a(t)(\widehat{x}_{m,nl}^2(t) - \widehat{x}_{nl}^2(t)) - \widehat{a}_{l,mn}^2(t) \nu_{m,nl}^x(t) - \nu_{l,mn}^a(t) \nu_{m,nl}^x(t)$ are both $O(1/N)$.

Rewriting (20) using a Taylor series expansion in x_{nl} about the point $\widehat{x}_{nl}(t)$, we get

$$\begin{aligned} \Delta_{m \rightarrow nl}^x(t, x_{nl}) &\approx \text{const} \\ &+ H_{ml}(\widehat{p}_{ml}(t) + O(1/N), \nu_{ml}^p(t) + O(1/N); y_{ml}) \\ &+ \widehat{a}_{l,mn}(t) (x_{nl} - \widehat{x}_{nl}(t)) \\ &\quad \times H'_{ml}(\widehat{p}_{ml}(t) + O(1/N), \nu_{ml}^p(t) + O(1/N); y_{ml}) \\ &+ 2\nu_{l,mn}^a(t) \widehat{x}_{nl}(t) (x_{nl} - \widehat{x}_{nl}(t)) \\ &\quad \times \dot{H}_{ml}(\widehat{p}_{ml}(t) + O(1/N), \nu_{ml}^p(t) + O(1/N); y_{ml}) \\ &+ \nu_{l,mn}^a(t) (x_{nl} - \widehat{x}_{nl}(t))^2 \\ &\quad \times \ddot{H}_{ml}(\widehat{p}_{ml}(t) + O(1/N), \nu_{ml}^p(t) + O(1/N); y_{ml}) \\ &+ \frac{1}{2} \widehat{a}_{l,mn}^2(t) (x_{nl} - \widehat{x}_{nl}(t))^2 \\ &\quad \times H''_{ml}(\widehat{p}_{ml}(t) + O(1/N), \nu_{ml}^p(t) + O(1/N); y_{ml}) \\ &+ O(1/N^{3/2}), \end{aligned} \quad (26)$$

where H'_{ml} and H''_{ml} are the first two derivatives of H_{ml} w.r.t its first argument and \dot{H}_{ml} is the first derivative w.r.t its second argument. Note that, in (26) and elsewhere, the higher-order terms in the Taylor's expansion are written solely in terms of their scaling dependence on N , which is what will eventually allow us to neglect these terms (in the large-system limit).

We now approximate (26) by dropping terms that vanish, relative to the second-to-last term in (26), as $N \rightarrow \infty$. Since this second-to-last term is $O(1/N)$ due to the scalings of $\widehat{a}_{l,mn}^2(t)$,

$\widehat{p}_{ml}(t)$, and $\nu_{ml}^p(t)$, we drop terms that are of order $O(1/N^{3/2})$, such as the final term. We also replace $\nu_{l,mn}^a(t)$ with $\nu_{mn}^a(t)$, and $\widehat{a}_{l,mn}^2(t)$ with $\widehat{a}_{mn}^2(t)$, since in both cases the difference is $O(1/N^{3/2})$. Finally, we drop the $O(1/N)$ terms inside the H_{ml} derivatives, which can be justified by taking a Taylor series expansion of these derivatives with respect to the $O(1/N)$ perturbations and verifying that the higher-order terms in this latter expansion are $O(1/N^{3/2})$. All of these approximations are analogous to those made in previous AMP derivations, e.g., [31], [32], and [33].

Applying these approximations to (26) and absorbing x_{nl} -invariant terms into the *const* term, we obtain

$$\begin{aligned} \Delta_{m \rightarrow nl}^x(t, x_{nl}) &\approx [\widehat{s}_{ml}(t) \widehat{a}_{l,mn}(t) + \nu_{ml}^s(t) \widehat{a}_{mn}^2(t) \widehat{x}_{nl}(t)] \\ &\quad \times x_{nl} - \frac{1}{2} [\nu_{ml}^s(t) \widehat{a}_{mn}^2(t) - \nu_{mn}^a(t) \\ &\quad \times (\widehat{s}_{ml}^2(t) - \nu_{ml}^s(t))] x_{nl}^2 + \text{const}, \end{aligned} \quad (27)$$

where we used the relationship

$$\dot{H}_{ml}(\widehat{q}, \nu^q; y) = \frac{1}{2} \left[(H'_{ml}(\widehat{q}, \nu^q; y))^2 + H''_{ml}(\widehat{q}, \nu^q; y) \right] \quad (28)$$

and defined

$$\widehat{s}_{ml}(t) \triangleq H'_{ml}(\widehat{p}_{ml}(t), \nu_{ml}^p(t); y_{ml}) \quad (29)$$

$$\nu_{ml}^s(t) \triangleq -H''_{ml}(\widehat{p}_{ml}(t), \nu_{ml}^p(t); y_{ml}). \quad (30)$$

Note that (27) is essentially a Gaussian approximation to the pdf $\frac{1}{C} \exp(\Delta_{m \rightarrow nl}^x(t, \cdot))$.

We show in Appendix A that

$$\widehat{s}_{ml}(t) = \frac{1}{\nu_{ml}^p(t)} (\widehat{z}_{ml}(t) - \widehat{p}_{ml}(t)) \quad (31)$$

$$\nu_{ml}^s(t) = \frac{1}{\nu_{ml}^p(t)} \left(1 - \frac{\nu_{ml}^z(t)}{\nu_{ml}^p(t)} \right), \quad (32)$$

for the conditional mean and variance

$$\widehat{z}_{ml}(t) \triangleq \mathbb{E}\{z_{ml} | p_{ml} = \widehat{p}_{ml}(t); \nu_{ml}^p(t)\} \quad (33)$$

$$\nu_{ml}^z(t) \triangleq \text{var}\{z_{ml} | p_{ml} = \widehat{p}_{ml}(t); \nu_{ml}^p(t)\}, \quad (34)$$

computed according to the (conditional) pdf

$$\begin{aligned} p_{z_{ml} | p_{ml}}(z_{ml} | \widehat{p}_{ml}(t); \nu_{ml}^p(t)) \\ \triangleq \frac{1}{C} p_{y_{ml} | z_{ml}}(y_{ml} | z_{ml}) \mathcal{N}(z_{ml}; \widehat{p}_{ml}(t), \nu_{ml}^p(t)), \end{aligned} \quad (35)$$

where here $C = \int_z p_{y_{ml} | z_{ml}}(y_{ml} | z) \mathcal{N}(z; \widehat{p}_{ml}(t), \nu_{ml}^p(t))$. In fact, (35) is BiG-AMP's iteration- t approximation to the true marginal posterior $p_{z_{ml} | \mathbf{Y}}(\cdot | \mathbf{Y})$. We note that (35) can also be interpreted as the (exact) posterior pdf for z_{ml} given the likelihood $p_{y_{ml} | z_{ml}}(y_{ml} | \cdot)$ from (3) and the prior $z_{ml} \sim \mathcal{N}(\widehat{p}_{ml}(t), \nu_{ml}^p(t))$ that is implicitly assumed by iteration- t BiG-AMP.

Since $\mathbf{Z}^\top = \mathbf{X}^\top \mathbf{A}^\top$, the derivation of the BiG-AMP approximation of $\Delta_{l \rightarrow mn}^a(t, \cdot)$ closely follows the derivation for $\Delta_{m \rightarrow nl}^x(t, \cdot)$. In particular, it starts with (similar to (13))

$$\begin{aligned} & \Delta_{l \rightarrow mn}^a(t, a_{mn}) \\ &= \log \int_{\{a_{mk}\}_{k \neq n}, \{x_{rl}\}_{r=1}^N} p_{y_{ml}|z_{ml}}(y_{ml} | \overbrace{a_{mn}x_{nl} + \sum_{k \neq n} a_{mk}x_{kl}}^{z_{ml}}) \\ & \quad \times \prod_{r=1}^N \exp(\Delta_{m \leftarrow rl}^x(t, x_{rl})) \prod_{k \neq n} \exp(\Delta_{l \leftarrow mk}^a(t, a_{mk})) \\ & \quad + \text{const}, \end{aligned} \quad (36)$$

where again the CLT motivates the treatment of z_{ml} , conditioned on $a_{mn} = a_{mn}$, as Gaussian. Eventually we arrive at the Taylor-series approximation (similar to (27))

$$\begin{aligned} \Delta_{l \rightarrow mn}^a(t, a_{mn}) &\approx [\hat{s}_{ml}(t)\hat{x}_{m,nl}(t) + \nu_{ml}^s(t)\hat{x}_{nl}^2(t)\hat{a}_{mn}(t)] \\ & \quad \times a_{mn} - \frac{1}{2}[\nu_{ml}^s(t)\hat{x}_{nl}^2(t) - \nu_{nl}^x(t) \\ & \quad \times (\hat{s}_{ml}^2(t) - \nu_{ml}^s(t))]a_{mn}^2 \\ & \quad + \text{const}. \end{aligned} \quad (37)$$

E. Approximated Variable-to-Factor Messages

We now turn to approximating the messages flowing from the variable nodes to the factor nodes. Starting with (8) and plugging in (27) we obtain

$$\begin{aligned} & \Delta_{m \leftarrow nl}^x(t+1, x_{nl}) \\ &\approx \text{const} + \log p_{x_{nl}}(x_{nl}) + \sum_{k \neq m} \left(\left[\hat{s}_{kl}(t)\hat{a}_{l,kn}(t) \right. \right. \\ & \quad \left. \left. + \nu_{kl}^s(t)\hat{a}_{kn}^2(t)\hat{x}_{nl}(t) \right] x_{nl} - \frac{1}{2} \left[\nu_{kl}^s(t)\hat{a}_{kn}^2(t) \right. \right. \\ & \quad \left. \left. - \nu_{kn}^a(t)(\hat{s}_{kl}^2(t) - \nu_{kl}^s(t)) \right] x_{nl}^2 \right) \end{aligned} \quad (38)$$

$$= \text{const} + \log p_{x_{nl}}(x_{nl}) - \frac{1}{2\nu_{m,nl}^r(t)}(x_{nl} - \hat{r}_{m,nl}(t))^2 \quad (39)$$

$$= \text{const} + \log \left(p_{x_{nl}}(x_{nl}) \mathcal{N}(x_{nl}; \hat{r}_{m,nl}(t), \nu_{m,nl}^r(t)) \right) \quad (40)$$

for

$$\nu_{m,nl}^r(t) \triangleq \left(\sum_{k \neq m} \hat{a}_{kn}^2(t)\nu_{kl}^s(t) - \nu_{kn}^a(t)(\hat{s}_{kl}^2(t) - \nu_{kl}^s(t)) \right)^{-1} \quad (41)$$

$$\begin{aligned} \hat{r}_{m,nl}(t) &\triangleq \hat{x}_{nl}(t) \left(1 + \nu_{m,nl}^r(t) \sum_{k \neq m} \nu_{kn}^a(t)[\hat{s}_{kl}^2(t) - \nu_{kl}^s(t)] \right) \\ & \quad + \nu_{m,nl}^r(t) \sum_{k \neq m} \hat{a}_{l,kn}(t)\hat{s}_{kl}(t). \end{aligned} \quad (42)$$

Since $\hat{a}_{mn}^2(t)$ and $\nu_{mn}^a(t)$ are $O(1/N)$, and recalling $\hat{s}_{ml}^2(t)$ and $\nu_{ml}^s(t)$ are $O(1)$, we take $\nu_{m,nl}^r(t)$ to be $O(1)$. Meanwhile, since $\hat{r}_{m,nl}(t)$ is an estimate of x_{nl} , we reason that it is $O(1)$.

The mean and variance of the pdf associated with the $\Delta_{m \leftarrow nl}^x(t+1, \cdot)$ approximation in (40) are

$$\begin{aligned} & \hat{x}_{m,nl}(t+1) \\ &\triangleq \frac{1}{C} \int_x x p_{x_{nl}}(x) \mathcal{N}(x; \hat{r}_{m,nl}(t), \nu_{m,nl}^r(t)) \\ & \quad \triangleq g_{x_{nl}}(\hat{r}_{m,nl}(t), \nu_{m,nl}^r(t)) \\ & \nu_{m,nl}^x(t+1) \\ &\triangleq \frac{1}{C} \int_x |x - \hat{x}_{m,nl}(t+1)|^2 p_{x_{nl}}(x) \mathcal{N}(x; \hat{r}_{m,nl}(t), \nu_{m,nl}^r(t)) \\ & \quad \nu_{m,nl}^r(t) g'_{x_{nl}}(\hat{r}_{m,nl}(t), \nu_{m,nl}^r(t)) \end{aligned} \quad (43)$$

where here $C = \int_x p_{x_{nl}}(x) \mathcal{N}(x; \hat{r}_{m,nl}(t), \nu_{m,nl}^r(t))$ and $g'_{x_{nl}}$ denotes the derivative of $g_{x_{nl}}$ with respect to the first argument. The fact that (43) and (44) are related through a derivative was shown in [33].

We now derive approximations of $\hat{x}_{m,nl}(t)$ and $\nu_{m,nl}^x(t)$ that avoid the dependence on the destination node m . For this, we introduce m -invariant versions of $\hat{r}_{m,nl}(t)$ and $\nu_{m,nl}^r(t)$:

$$\nu_{nl}^r(t) \triangleq \left(\sum_{m=1}^M \hat{a}_{mn}^2(t)\nu_{ml}^s(t) - \nu_{mn}^a(t)(\hat{s}_{ml}^2(t) - \nu_{ml}^s(t)) \right)^{-1} \quad (45)$$

$$\begin{aligned} \hat{r}_{nl}(t) &\triangleq \hat{x}_{nl}(t) \left(1 + \nu_{nl}^r(t) \sum_{m=1}^M \nu_{mn}^a(t)[\hat{s}_{ml}^2(t) - \nu_{ml}^s(t)] \right) \\ & \quad + \nu_{nl}^r(t) \sum_{m=1}^M \hat{a}_{l,mn}(t)\hat{s}_{ml}(t). \end{aligned} \quad (46)$$

Comparing (45)–(46) with (41)–(42) and applying previously established scalings from Table II reveals that $\nu_{m,nl}^r(t) - \nu_{nl}^r(t)$ is $O(1/N)$ and that $\hat{r}_{m,nl}(t) = \hat{r}_{nl}(t) - \nu_{nl}^r(t)\hat{a}_{mn}(t)\hat{s}_{ml}(t) + O(1/N)$, so that (43) implies

$$\begin{aligned} & \hat{x}_{m,nl}(t+1) \\ &= g_{x_{nl}}(\hat{r}_{nl}(t) - \nu_{nl}^r(t)\hat{a}_{mn}(t)\hat{s}_{ml}(t) + O(1/N), \\ & \quad \nu_{nl}^r(t) + O(1/N)) \end{aligned} \quad (47)$$

$$= g_{x_{nl}}(\hat{r}_{nl}(t) - \nu_{nl}^r(t)\hat{a}_{mn}(t)\hat{s}_{ml}(t), \nu_{nl}^r(t)) + O(1/N) \quad (48)$$

$$= g_{x_{nl}}(\hat{r}_{nl}(t), \nu_{nl}^r(t)) \quad (49)$$

$$\begin{aligned} & \quad - \nu_{nl}^r(t)\hat{a}_{mn}(t)\hat{s}_{ml}(t)g'_{x_{nl}}(\hat{r}_{nl}(t), \nu_{nl}^r(t)) + O(1/N) \\ & \approx \hat{x}_{nl}(t+1) - \hat{a}_{mn}(t)\hat{s}_{ml}(t)\nu_{nl}^x(t+1). \end{aligned} \quad (50)$$

Above, (48) follows from taking Taylor series expansions around each of the $O(1/N)$ perturbations in (47); (49) follows from a Taylor series expansion in the first argument of (48) about the point $\hat{r}_{nl}(t)$; and (50) follows by neglecting the $O(1/N)$ term (which vanishes relative to the others in the large-system limit) and applying the definitions

$$\hat{x}_{nl}(t+1) \triangleq g_{x_{nl}}(\hat{r}_{nl}(t), \nu_{nl}^r(t)) \quad (51)$$

$$\nu_{nl}^x(t+1) \triangleq \nu_{nl}^r(t)g'_{x_{nl}}(\hat{r}_{nl}(t), \nu_{nl}^r(t)), \quad (52)$$

which match (43)–(44) sans the m dependence. Note that (50) confirms that the difference $\hat{x}_{m,nl}(t) - \hat{x}_{nl}(t)$ is $O(1/\sqrt{N})$, as

was assumed at the start of the BiG-AMP derivation. Likewise, taking Taylor series expansions of $g'_{x_{nl}}$ in (44) about the point $\hat{r}_{nl}(t)$ in the first argument and about the point $\nu_{nl}^r(t)$ in the second argument and then comparing the result with (52) confirms that $\nu_{m,nl}^x(t) - \nu_{nl}^x(t)$ is $O(1/\sqrt{N})$.

We then repeat the above procedure to derive an approximation to $\Delta_{l \leftarrow mn}^a(t+1, \cdot)$ analogous to (40), whose corresponding mean is then further approximated as

$$\hat{a}_{l,mn}(t+1) \approx \hat{a}_{mn}(t+1) - \hat{x}_{nl}(t)\hat{s}_{ml}(t)\nu_{mn}^a(t+1), \quad (53)$$

for

$$\hat{a}_{mn}(t+1) \triangleq g_{a_{mn}}(\hat{q}_{mn}(t), \nu_{mn}^q(t)) \quad (54)$$

$$\nu_{mn}^a(t+1) \triangleq \nu_{mn}^q(t)g'_{a_{mn}}(\hat{q}_{mn}(t), \nu_{mn}^q(t)) \quad (55)$$

$$g_{a_{mn}}(\hat{q}, \nu^q) \triangleq \frac{\int_a a p_{a_{mn}}(a)\mathcal{N}(a; \hat{q}, \nu^q)}{\int_a p_{a_{mn}}(a)\mathcal{N}(a; \hat{q}, \nu^q)} \quad (56)$$

where

$$\nu_{mn}^q(t) \triangleq \left(\sum_{l=1}^L \hat{x}_{nl}^2(t)\nu_{ml}^s(t) - \nu_{nl}^x(t)(\hat{s}_{ml}^2(t) - \nu_{ml}^s(t)) \right)^{-1} \quad (57)$$

$$\begin{aligned} \hat{q}_{mn}(t) &\triangleq \hat{a}_{mn}(t) \left(1 + \nu_{mn}^q(t) \sum_{l=1}^L \nu_{nl}^x(t)[\hat{s}_{ml}^2(t) - \nu_{ml}^s(t)] \right) \\ &+ \nu_{mn}^q(t) \sum_{l=1}^L \hat{x}_{m,nl}(t)\hat{s}_{ml}(t). \end{aligned} \quad (58)$$

Arguments analogous to the discussion following (42) justify the remaining scalings in Table II.

F. Closing the Loop

The penultimate step in the derivation of BiG-AMP is to approximate earlier steps that use $\hat{a}_{l,mn}(t)$ and $\hat{x}_{m,nl}(t)$ in place of $\hat{a}_{mn}(t)$ and $\hat{x}_{nl}(t)$. For this, we start by plugging (50) and (53) into (22), which yields³

$$\begin{aligned} \hat{p}_{ml}(t) &\triangleq \overbrace{\bar{p}_{ml}(t)} \\ &= O(1/\sqrt{N}) + \sum_{n=1}^N \hat{a}_{mn}(t)\hat{x}_{nl}(t) - \hat{s}_{ml}(t-1) \\ &\quad \times \sum_{n=1}^N \left(\nu_{mn}^a(t)\hat{x}_{nl}(t)\hat{x}_{nl}(t-1) + \hat{a}_{mn}(t)\hat{a}_{mn}(t-1)\nu_{nl}^x(t) \right) \\ &\quad + \hat{s}_{ml}^2(t-1) \sum_{n=1}^N \hat{a}_{mn}(t-1)\nu_{mn}^a(t)\nu_{nl}^x(t)\hat{x}_{nl}(t-1) \quad (59) \\ &\approx \bar{p}_{ml}(t) - \hat{s}_{ml}(t-1) \underbrace{\sum_{n=1}^N \left(\nu_{mn}^a(t)\hat{x}_{nl}^2(t) + \hat{a}_{mn}^2(t)\nu_{nl}^x(t) \right)}_{\triangleq \bar{\nu}_{ml}^p(t)}, \end{aligned} \quad (60)$$

where, for (60), we used $\hat{a}_{mn}^2(t)$ in place of $\hat{a}_{mn}(t)\hat{a}_{mn}(t-1)$, used $\hat{x}_{nl}^2(t)$ in place of $\hat{x}_{nl}(t)\hat{x}_{nl}(t-1)$, and neglected terms that are $O(1/\sqrt{N})$, since they vanish relative to the remaining $O(1)$ terms in the large-system limit.

³Recall that the error of the approximation in (50) is $O(1/N)$ and the error in (53) is $O(1/N^{3/2})$.

Next we plug (50), (53), $\nu_{m,nl}^x(t) = \nu_{nl}^x(t) + O(1/\sqrt{N})$, and $\nu_{l,mn}^a(t) = \nu_{mn}^a(t) + O(1/N^{3/2})$ into (23), giving

$$\nu_{ml}^p(t) = \bar{\nu}_{ml}^p(t) + \sum_{n=1}^N \nu_{mn}^a(t)\nu_{nl}^x(t) \quad (61)$$

$$- 2\hat{s}_{ml}(t-1) \sum_{n=1}^N \left(\nu_{mn}^a(t)\hat{a}_{mn}(t)\hat{x}_{nl}(t-1)\nu_{nl}^x(t) + \nu_{mn}^a(t)\hat{a}_{mn}(t-1)\hat{x}_{nl}(t)\nu_{nl}^x(t) \right)$$

$$+ \hat{s}_{ml}(t-1)^2 \sum_{n=1}^N \left((\nu_{mn}^a(t))^2 \hat{x}_{nl}^2(t-1)\nu_{nl}^x(t) + \nu_{mn}^a(t)\hat{a}_{mn}^2(t-1)(\nu_{nl}^x(t))^2 \right) + O(1/\sqrt{N})$$

$$\approx \bar{\nu}_{ml}^p(t) + \sum_{n=1}^N \nu_{mn}^a(t)\nu_{nl}^x(t), \quad (62)$$

where (62) retains only the $O(1)$ terms from (61).

Similarly, we plug (53) into (46) and (50) into (58) to obtain

$$\begin{aligned} \hat{r}_{nl}(t) &\approx \hat{x}_{nl}(t) \left(1 - \frac{\sum_{m=1}^M \nu_{mn}^a(t)\nu_{ml}^s(t)}{\sum_{m=1}^M \hat{a}_{mn}(t)\nu_{ml}^s(t)} \right) \\ &+ \nu_{nl}^r(t) \sum_{m=1}^M \hat{a}_{mn}(t)\hat{s}_{ml}(t) \end{aligned} \quad (63)$$

$$\begin{aligned} \hat{q}_{mn}(t) &\approx \hat{a}_{mn}(t) \left(1 - \frac{\sum_{l=1}^L \nu_{nl}^x(t)\nu_{ml}^s(t)}{\sum_{l=1}^L \hat{x}_{nl}^2(t)\nu_{ml}^s(t)} \right) \\ &+ \nu_{mn}^q(t) \sum_{l=1}^L \hat{x}_{nl}(t)\hat{s}_{ml}(t), \end{aligned} \quad (64)$$

where the approximations involve the use of $\hat{s}_{ml}^2(t)$ in place of $\hat{s}_{ml}(t)\hat{s}_{ml}(t-1)$, of $\hat{a}_{mn}(t)$ in place of $\hat{a}_{mn}(t-1)$, of $\hat{x}_{nl}(t)$ in place of $\hat{x}_{nl}(t-1)$, and the dropping of terms that vanish in the large-system limit. Finally, we make the approximations

$$\nu_{nl}^r(t) \approx \left(\sum_{m=1}^M \hat{a}_{mn}^2(t)\nu_{ml}^s(t) \right)^{-1} \quad (65)$$

$$\nu_{mn}^q(t) \approx \left(\sum_{l=1}^L \hat{x}_{nl}^2(t)\nu_{ml}^s(t) \right)^{-1}, \quad (66)$$

by neglecting the $\hat{s}_{ml}^2(t) - \nu_{ml}^s(t)$ terms in (45) and (57), as explained in Appendix B.

G. Approximated Posteriors

The final step in the BiG-AMP derivation is to approximate the SPA posterior log-pdfs in (11) and (12). Plugging (27) and (37) into those expressions, we get

$$\begin{aligned} \Delta_{nl}^x(t+1, x_{nl}) &\approx \text{const} + \log \left(p_{x_{nl}}(x_{nl})\mathcal{N}(x_{nl}; \hat{r}_{nl}(t), \nu_{nl}^r(t)) \right) \end{aligned} \quad (67)$$

$$\begin{aligned} \Delta_{mn}^a(t+1, a_{mn}) &\approx \text{const} + \log \left(p_{a_{mn}}(a_{mn})\mathcal{N}(a_{mn}; \hat{q}_{mn}(t), \nu_{mn}^q(t)) \right) \end{aligned} \quad (68)$$

using steps similar to (40). The associated pdfs are

$$\begin{aligned} p_{x_{nl}|r_{nl}}(x_{nl}|\hat{r}_{nl}(t); \nu_{nl}^r(t)) &\triangleq \frac{1}{C_x} p_{x_{nl}}(x_{nl})\mathcal{N}(x_{nl}; \hat{r}_{nl}(t), \nu_{nl}^r(t)) \end{aligned} \quad (69)$$

TABLE III
THE BiG-AMP ALGORITHM

definitions:	
$p_{z_{ml} p_{ml}}(z \hat{p}; \nu^p) \triangleq \frac{p_{y_{ml} z_{ml}}(y_{ml} z) \mathcal{N}(z; \hat{p}, \nu^p)}{\int_{z'} p_{y_{ml} z'_{ml}}(y_{ml} z') \mathcal{N}(z'; \hat{p}, \nu^p)}$	(D1)
$p_{x_{nl} r_{nl}}(x \hat{r}; \nu^r) \triangleq \frac{p_{x_{nl}}(x) \mathcal{N}(x; \hat{r}, \nu^r)}{\int_{x'} p_{x_{nl}}(x') \mathcal{N}(x'; \hat{r}, \nu^r)}$	(D2)
$p_{a_{mn} q_{mn}}(a \hat{q}; \nu^q) \triangleq \frac{p_{a_{mn}}(a) \mathcal{N}(a; \hat{q}, \nu^q)}{\int_{a'} p_{a_{mn}}(a') \mathcal{N}(a'; \hat{q}, \nu^q)}$	(D3)
initialization:	
$\forall m, l : \hat{s}_{ml}(0) = 0$	(I1)
$\forall m, n, l : \text{choose } \nu_{nl}^x(1), \hat{x}_{nl}(1), \nu_{mn}^a(1), \hat{a}_{mn}(1)$	(I2)
for $t = 1, \dots, T_{\max}$	
$\forall m, l : \bar{\nu}_{ml}^p(t) = \sum_{n=1}^N \hat{a}_{mn}(t) ^2 \nu_{nl}^x(t) + \nu_{mn}^a(t) \hat{x}_{nl}(t) ^2$	(R1)
$\forall m, l : \bar{p}_{ml}(t) = \sum_{n=1}^N \hat{a}_{mn}(t) \hat{x}_{nl}(t)$	(R2)
$\forall m, l : \nu_{ml}^p(t) = \bar{\nu}_{ml}^p(t) + \sum_{n=1}^N \nu_{mn}^a(t) \nu_{nl}^x(t)$	(R3)
$\forall m, l : \hat{p}_{ml}(t) = \bar{p}_{ml}(t) - \hat{s}_{ml}(t-1) \bar{\nu}_{ml}^p(t)$	(R4)
$\forall m, l : \nu_{ml}^z(t) = \text{var}\{\mathbf{Z}_{ml} p_{ml} = \hat{p}_{ml}(t); \nu_{ml}^p(t)\}$	(R5)
$\forall m, l : \hat{z}_{ml}(t) = \mathbf{E}\{\mathbf{Z}_{ml} p_{ml} = \hat{p}_{ml}(t); \nu_{ml}^p(t)\}$	(R6)
$\forall m, l : \nu_{ml}^s(t) = (1 - \nu_{ml}^z(t) / \nu_{ml}^p(t)) / \nu_{ml}^p(t)$	(R7)
$\forall m, l : \hat{s}_{ml}(t) = (\hat{z}_{ml}(t) - \bar{p}_{ml}(t)) / \nu_{ml}^p(t)$	(R8)
$\forall n, l : \nu_{nl}^r(t) = (\sum_{m=1}^M \hat{a}_{mn}(t) ^2 \nu_{ml}^s(t))^{-1}$	(R9)
$\forall n, l : \hat{r}_{nl}(t) = \hat{x}_{nl}(t) (1 - \nu_{nl}^r(t) \sum_{m=1}^M \nu_{mn}^a(t) \nu_{ml}^s(t)) + \nu_{nl}^r(t) \sum_{m=1}^M \hat{a}_{mn}(t) \hat{s}_{ml}(t)$	(R10)
$\forall m, n : \nu_{mn}^q(t) = (\sum_{l=1}^L \hat{x}_{nl}(t) ^2 \nu_{ml}^r(t))^{-1}$	(R11)
$\forall m, n : \hat{q}_{mn}(t) = \hat{a}_{mn}(t) (1 - \nu_{mn}^q(t) \sum_{l=1}^L \nu_{nl}^r(t) \nu_{ml}^s(t)) + \nu_{mn}^q(t) \sum_{l=1}^L \hat{x}_{nl}(t) \hat{s}_{ml}(t)$	(R12)
$\forall n, l : \nu_{nl}^x(t+1) = \text{var}\{x_{nl} r_{nl} = \hat{r}_{nl}(t); \nu_{nl}^r(t)\}$	(R13)
$\forall n, l : \hat{x}_{nl}(t+1) = \mathbf{E}\{x_{nl} r_{nl} = \hat{r}_{nl}(t); \nu_{nl}^r(t)\}$	(R14)
$\forall m, n : \nu_{mn}^a(t+1) = \text{var}\{a_{mn} q_{mn} = \hat{q}_{mn}(t); \nu_{mn}^q(t)\}$	(R15)
$\forall m, n : \hat{a}_{mn}(t+1) = \mathbf{E}\{a_{mn} q_{mn} = \hat{q}_{mn}(t); \nu_{mn}^q(t)\}$	(R16)
if $\sum_{m,l} \bar{p}_{ml}(t) - \bar{p}_{ml}(t-1) ^2 \leq \tau_{\text{BiG-AMP}} \sum_{m,l} \bar{p}_{ml}(t) ^2$, stop end	(R17)

$$p_{a_{mn}|q_{mn}}(a_{mn}|\hat{q}_{mn}(t); \nu_{mn}^q(t)) \triangleq \frac{1}{C_a} p_{a_{mn}}(a_{mn}) \mathcal{N}(a_{mn}; \hat{q}_{mn}(t), \nu_{mn}^q(t)) \quad (70)$$

for $C_x \triangleq \int_x p_{x_{nl}}(x) \mathcal{N}(x; \hat{r}_{nl}(t), \nu_{nl}^r(t))$ and $C_a \triangleq \int_a p_{a_{mn}}(a) \mathcal{N}(a; \hat{q}_{mn}(t), \nu_{mn}^q(t))$, which are iteration- t BiG-AMP's approximations to the true marginal posteriors $p_{x_{nl}|\mathbf{Y}}(x_{nl}|\mathbf{Y})$ and $p_{a_{mn}|\mathbf{Y}}(a_{mn}|\mathbf{Y})$, respectively.

Note that $\hat{x}_{nl}(t+1)$ and $\nu_{nl}^x(t+1)$ from (51)–(52) are the mean and variance, respectively, of the posterior pdf in (69). Note also that (69) can be interpreted as the (exact) posterior pdf of x_{nl} given the observation $r_{nl} = \hat{r}_{nl}(t)$ under the prior model $x_{nl} \sim p_{x_{nl}}$ and the likelihood model $p_{r_{nl}|x_{nl}}(\hat{r}_{nl}(t)|x_{nl}; \nu_{nl}^r(t)) = \mathcal{N}(\hat{r}_{nl}(t); x_{nl}, \nu_{nl}^r(t))$ implicitly assumed by iteration- t BiG-AMP. Analogous statements can be made about the posterior pdf of a_{mn} in (70).

This completes the derivation of BiG-AMP. As a final comment, we note that some of the approximations used in (60), (63), and (64) are not exact in the large-system limit. Empirical experiments, however, suggest that the effect of these approximations are very minor. Although we would have liked to first present the “exact” version of the algorithm and then simplify it to yield the BiG-AMP in Table III, this was prevented by page constraints.

Algorithm Summary

The BiG-AMP algorithm derived in Sections II-C to II-G is summarized in Table III. There, we have included a maximum number of iterations, T_{\max} and a stopping condition (R17) based on the (normalized) change in the residual and a user-defined parameter $\tau_{\text{BiG-AMP}}$. We have also written the algorithm

in a more general form that allows the use of complex-valued quantities [note the complex conjugates in (R10) and (R12)], in which case \mathcal{N} in (D1)–(D3) would be circular complex Gaussian. For ease of interpretation, Table III does not include the important damping modifications that will be detailed in Section IV-A. Suggestions for the initializations in (I2) will be given in the sequel.

We note that BiG-AMP avoids the use of SVD or QR decompositions, lending itself to simple and potentially parallel implementations. Its complexity order is dominated⁴ by ten matrix multiplications per iteration [in steps (R1)–(R3) and (R9)–(R12)], each requiring MNL multiplications, although simplifications will be discussed in Section III.

The steps in Table III can be interpreted as follows. (R1)–(R2) compute a “plug-in” estimate $\bar{\mathbf{P}}$ of the matrix product $\mathbf{Z} = \mathbf{A}\mathbf{X}$ and a corresponding set of element-wise variances $\{\bar{\nu}_{ml}^p\}$. (R3)–(R4) then apply “Onsager” correction (see [32] and [33] for discussions in the contexts of AMP and GAMP, respectively) to obtain the corresponding quantities $\hat{\mathbf{P}}$ and $\{\nu_{ml}^p\}$. Using these quantities, (R5)–(R6) compute the (approximate) marginal posterior means $\hat{\mathbf{Z}}$ and variances $\{\nu_{ml}^z\}$ of \mathbf{Z} . Steps (R7)–(R8) then use these posterior moments to compute the scaled residual $\hat{\mathbf{S}}$ and a set of inverse-residual-variances $\{\nu_{ml}^s\}$. This interpretation becomes clear in the case of AWGN observations with noise variance ν^w , where

$$p_{y_{ml}|z_{ml}}(y_{ml}|z_{ml}) = \mathcal{N}(y_{ml}; z_{ml}, \nu^w) \quad (71)$$

and hence

$$\nu_{ml}^s = \frac{1}{\nu_{ml}^p + \nu^w} \text{ and } \hat{s}_{ml} = \frac{y_{ml} - \hat{p}_{ml}}{\nu_{ml}^p + \nu^w}. \quad (72)$$

Steps (R9)–(R10) then use the residual terms $\hat{\mathbf{S}}$ and $\{\nu_{ml}^s\}$ to compute $\hat{\mathbf{R}}$ and $\{\nu_{nl}^r\}$, where \hat{r}_{nl} can be interpreted as a ν_{nl}^r -variance-AWGN corrupted observation of the true x_{nl} . Similarly, (R11)–(R12) compute $\hat{\mathbf{Q}}$ and $\{\nu_{mn}^q\}$, where \hat{q}_{mn} can be interpreted as a ν_{mn}^q -variance-AWGN corrupted observation of the true a_{mn} . Finally, (R13)–(R14) merge these AWGN-corrupted observations with the priors $\{p_{x_{nl}}\}$ to produce the posterior means $\hat{\mathbf{X}}$ and variances $\{\nu_{nl}^x\}$; (R15)–(R16) do the same for the a_{mn} quantities.

The BiG-AMP algorithm in Table III is a direct (although non-trivial) extension of the GAMP algorithm for compressive sensing [33], which estimates \mathbf{X} assuming perfectly known \mathbf{A} , and even stronger similarities to the \mathbf{A} -uncertain GAMP from [50], which estimates \mathbf{X} assuming knowledge of the marginal means and variances of unknown random \mathbf{A} , but which makes no attempt to estimate \mathbf{A} itself. In Section III-B, a simplified version of BiG-AMP will be developed that is similar to the Bayesian-AMP algorithm [31] for compressive sensing.

III. BiG-AMP SIMPLIFICATIONS

We now describe simplifications of the BiG-AMP algorithm from Table III that result from additional approximations and from the use of specific priors $p_{y_{ml}|z_{ml}}$, $p_{x_{nl}}$, and $p_{a_{mn}}$ that arise in practical applications of interest.

⁴The computations in steps (R4)–(R8) are $O(ML)$, while the remainder of the algorithm is $O(MN + NL)$. Thus, as N grows, the matrix multiplies dominate the complexity.

A. Scalar Variances

The BiG-AMP algorithm in Table III stores and processes a number of element-wise variance terms whose values vary across the elements (e.g., ν_{nl}^x can vary across n and l). The use of scalar variances (i.e., uniform across m, n, l) significantly reduces the memory and complexity of the algorithm.

To derive scalar-variance BiG-AMP, we first assume $\forall n, l : \nu_{nl}^x(t) \approx \nu^x(t) \triangleq \frac{1}{NL} \sum_{n=1}^N \sum_{l=1}^L \nu_{nl}^x(t)$ and $\forall m, n : \nu_{mn}^a(t) \approx \nu^a(t) \triangleq \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N \nu_{mn}^a(t)$, so from (R1)

$$\bar{\nu}_{ml}^p(t) \approx \nu^x(t) \sum_{n=1}^N |\hat{a}_{mn}(t)|^2 + \nu^a(t) \sum_{n=1}^N |\hat{x}_{nl}(t)|^2 \quad (73)$$

$$\approx \frac{\|\hat{\mathbf{A}}(t)\|_F^2}{M} \nu^x(t) + \frac{\|\hat{\mathbf{X}}(t)\|_F^2}{L} \nu^a(t) \triangleq \bar{\nu}^p(t). \quad (74)$$

Note that using (74) in place of (R1) avoids two matrix multiplies. Plugging these approximations into (R3) gives

$$\nu_{ml}^p(t) \approx \bar{\nu}^p(t) + N\nu^a(t)\nu^x(t) \triangleq \nu^p(t) \quad (75)$$

which, when used in place of (R3), avoids another matrix multiply. Even with the above scalar-variance approximations, $\{\nu_{ml}^s(t)\}$ from (R5) are not guaranteed to be equal (except in special cases like AWGN $p_{y_{ml}|z_{ml}}$). Still, they can be approximated as such using $\nu^s(t) \triangleq \frac{1}{ML} \sum_{m=1}^M \sum_{l=1}^L \nu_{ml}^s(t)$, in which case

$$\nu_{nl}^r(t) \approx \frac{1}{\nu^s(t) \sum_{m=1}^M |\hat{a}_{mn}(t)|^2} \approx \frac{N}{\nu^s(t) \|\hat{\mathbf{A}}(t)\|_F^2} \triangleq \nu^r(t) \quad (76)$$

$$\nu_{mn}^q(t) \approx \frac{1}{\nu^s(t) \sum_{l=1}^L |\hat{x}_{nl}(t)|^2} \approx \frac{N}{\nu^s(t) \|\hat{\mathbf{X}}(t)\|_F^2} \triangleq \nu^q(t). \quad (77)$$

Using (76) in place of (R9) and (77) in place of (R11) avoids two matrix multiplies and $NL + MN - 2$ scalar divisions, and furthermore allows (R10) and (R12) to be implemented as

$$\hat{r}_{nl}(t) = \hat{x}_{nl}(t) \left(1 - \frac{MN\nu^a(t)}{\|\hat{\mathbf{A}}(t)\|_F^2} \right) + \nu^r(t) \sum_{m=1}^M \hat{a}_{mn}(t) \hat{s}_{ml}(t) \quad (78)$$

$$\hat{q}_{mn}(t) = \hat{a}_{mn}(t) \left(1 - \frac{NL\nu^x(t)}{\|\hat{\mathbf{X}}(t)\|_F^2} \right) + \nu^q(t) \sum_{l=1}^L \hat{x}_{nl}(t) \hat{s}_{ml}(t), \quad (79)$$

saving two more matrix multiplies, and leaving a total of only three matrix multiplies per iteration.

B. Possibly Incomplete AWGN Observations

We now consider a particular observation model wherein the elements of $\mathbf{Z} = \mathbf{A}\mathbf{X}$ are AWGN-corrupted at a subset of indices $\Omega \subset (1 \dots M) \times (1 \dots L)$ and unobserved at the remaining indices, noting that the standard AWGN model (71)

is the special case where $|\Omega| = ML$. This ‘‘possibly incomplete AWGN’’ (PIAWGN) model arises in a number of important applications, such as matrix completion and dictionary learning.

We can state the PIAWGN model probabilistically as

$$p_{y_{ml}|z_{ml}}(y_{ml}|z_{ml}) = \begin{cases} \mathcal{N}(y_{ml}; z_{ml}, \nu^w) & (m, l) \in \Omega \\ \mathbb{1}_{y_{ml}} & (m, l) \notin \Omega, \end{cases} \quad (80)$$

where ν^w is the noise variance on the non-missing observations and $\mathbb{1}_y$ denotes a point mass at $y = 0$. Thus, at the observed entries $(m, l) \in \Omega$, the quantities \hat{s}_{ml} and ν_{ml}^s calculated using the AWGN expressions (72), while at the ‘‘missing’’ entries $(m, l) \notin \Omega$, where y_{ml} is invariant to z_{ml} , we have $E\{z_{ml}|p_{ml} = \hat{p}_{ml}; \nu_{ml}^p\} = \hat{p}_{ml}$ and $\text{var}\{z_{ml}|p_{ml} = \hat{p}_{ml}; \nu_{ml}^p\} = \nu_{ml}^p$, so that $\hat{s}_{ml} = 0$ and $\nu_{ml}^s = 0$. This is expected, given that ν^s can be interpreted as an inverse residual variance and \hat{s} as a ν^s -scaled residual. In summary, the PIAWGN model yields

$$\hat{s}_{ml}(t) = \begin{cases} \frac{y_{ml} - \hat{p}_{ml}(t)}{\nu_{ml}^p(t) + \nu^w} & (m, l) \in \Omega \\ 0 & (m, l) \notin \Omega \end{cases} \quad (81)$$

$$\nu_{ml}^s(t) = \begin{cases} \frac{1}{\nu_{ml}^p(t) + \nu^w} & (m, l) \in \Omega \\ 0 & (m, l) \notin \Omega \end{cases}. \quad (82)$$

When the PIAWGN model is combined with the scalar-variance approximations from Section III-A, BiG-AMP simplifies considerably. To see this, we start by using $\nu^p(t)$ from (75) in place of $\nu_{ml}^p(t)$ in (81)–(82), resulting in

$$\hat{\mathbf{S}}(t) = P_\Omega \left(\frac{\mathbf{Y} - \hat{\mathbf{P}}(t)}{\nu^p(t) + \nu^w} \right) \quad (83)$$

$$\nu^s(t) = \frac{\delta}{\nu^w + \nu^p(t)}, \quad (84)$$

where $\delta \triangleq \frac{|\Omega|}{ML}$ denotes the fraction of observed entries and $P_\Omega : \mathbb{R}^{M \times L} \rightarrow \mathbb{R}^{M \times L}$ is the projection operator defined by

$$[P_\Omega(\mathbf{Z})]_{ml} \triangleq \begin{cases} z_{ml} & (m, l) \in \Omega \\ 0 & (m, l) \notin \Omega \end{cases}. \quad (85)$$

We can then write (R10) and (R12) as

$$\hat{\mathbf{R}}(t) = \hat{\mathbf{X}}(t) \left(1 - \frac{MN\nu^a(t)}{\|\hat{\mathbf{A}}(t)\|_F^2} \right) + \frac{N}{\delta \|\hat{\mathbf{A}}(t)\|_F^2} \hat{\mathbf{A}}^\top(t) \hat{\mathbf{V}}(t) \quad (86)$$

$$\hat{\mathbf{Q}}(t) = \hat{\mathbf{A}}(t) \left(1 - \frac{NL\nu^x(t)}{\|\hat{\mathbf{X}}(t)\|_F^2} \right) + \frac{N}{\delta \|\hat{\mathbf{X}}(t)\|_F^2} \hat{\mathbf{V}}(t) \hat{\mathbf{X}}^\top(t) \quad (87)$$

using (78)–(79) and (83)–(85) with

$$\hat{\mathbf{V}}(t) \triangleq P_\Omega(\mathbf{Y} - \hat{\mathbf{P}}(t)) \quad (88)$$

$$= P_\Omega(\mathbf{Y} - \bar{\mathbf{P}}(t)) + \bar{\nu}^p(t) \hat{\mathbf{S}}(t-1) \quad (89)$$

$$= P_\Omega(\mathbf{Y} - \bar{\mathbf{P}}(t)) + \frac{\bar{\nu}^p(t)}{\nu^p(t-1) + \nu^w} \hat{\mathbf{V}}(t-1), \quad (90)$$

since P_Ω is a projection operator, and using (R4) and (83).

Scalar-variance BiG-AMP under PIAWGN observations is summarized in Table IV. Note that the residual matrix $\hat{\mathbf{U}}(t) \triangleq$

TABLE IV
SCALAR-VARIANCE BiG-AMP WITH PIAWGN $p_{Y|Z}$

initialization:	
	$\hat{\mathbf{V}}(0) = \mathbf{0}$ (I1p)
	choose $\nu^x(1), \hat{\mathbf{X}}(1), \nu^a(1), \hat{\mathbf{A}}(1)$ (I2p)
for $t = 1, \dots, T_{\max}$	
	$G_a(t) = \frac{N}{\delta \ \hat{\mathbf{A}}(t)\ _F^2}$ (R1p)
	$G_x(t) = \frac{N}{\delta \ \hat{\mathbf{X}}(t)\ _F^2}$ (R2p)
	$\hat{\mathbf{U}}(t) = P_{\Omega}(\mathbf{Y} - \hat{\mathbf{A}}(t)\hat{\mathbf{X}}(t))$ (R3p)
	$\bar{\nu}^p(t) = \left(\frac{\nu^x(t)}{M G_a(t)} + \frac{\nu^a(t)}{L G_x(t)} \right) \frac{N}{\delta}$ (R4p)
	$\nu^p(t) = \bar{\nu}^p(t) + N \nu^a(t) \nu^x(t)$ (R5p)
	$\hat{\mathbf{V}}(t) = \hat{\mathbf{U}}(t) + \frac{\bar{\nu}^p(t)}{\nu^p(t-1) + \nu^w} \hat{\mathbf{V}}(t-1)$ (R6p)
	$\nu^r(t) = G_a(t)(\nu^p(t) + \nu^w)$ (R7p)
	$\hat{\mathbf{R}}(t) = (1 - M \delta \nu^a(t) G_a(t)) \hat{\mathbf{X}}(t) + G_a(t) \hat{\mathbf{A}}^H(t) \hat{\mathbf{V}}(t)$ (R8p)
	$\nu^q(t) = G_x(t)(\nu^p(t) + \nu^w)$ (R9p)
	$\hat{\mathbf{Q}}(t) = (1 - L \delta \nu^x(t) G_x(t)) \hat{\mathbf{A}}(t) + G_x(t) \hat{\mathbf{V}}(t) \hat{\mathbf{X}}^H(t)$ (R10p)
	$\nu^x(t+1) = \frac{1}{N L} \sum_{n=1}^N \sum_{l=1}^L \text{var}\{x_{nl} \mathbf{Y}; \hat{r}_{nl}(t), \nu^r(t)\}$ (R11p)
$\forall n, l$	$\hat{x}_{nl}(t+1) = \mathbb{E}\{x_{nl} \mathbf{Y}; \hat{r}_{nl}(t), \nu^r(t)\}$ (R12p)
	$\nu^a(t+1) = \frac{1}{M N} \sum_{m=1}^M \sum_{n=1}^N \text{var}\{a_{mn} \mathbf{Y}; \hat{q}_{mn}(t), \nu^q(t)\}$ (R13p)
$\forall m, n$	$\hat{a}_{mn}(t+1) = \mathbb{E}\{a_{mn} \mathbf{Y}; \hat{q}_{mn}(t), \nu^q(t)\}$ (R14p)
	if $\ \hat{\mathbf{U}}(t) - \hat{\mathbf{U}}(t-1)\ _F^2 \leq \tau_{\text{BiG-AMP}} \ \hat{\mathbf{U}}(t)\ _F^2$, stop (R15p)
end	

TABLE V
BiG-AMP-LITE: SCALAR-VARIANCE, PIAWGN, GAUSSIAN $p_{\mathbf{X}}$ AND $p_{\mathbf{A}}$

initialization:	
	$\hat{\mathbf{V}}(0) = \mathbf{0}$ (I1i)
	choose $\nu^x(1), \hat{\mathbf{X}}(1), \nu^a(1), \hat{\mathbf{A}}(1)$ (I2i)
for $t = 1, \dots, T_{\max}$	
	$G_a(t) = \frac{N}{\delta \ \hat{\mathbf{A}}(t)\ _F^2}$ (R1i)
	$G_x(t) = \frac{N}{\delta \ \hat{\mathbf{X}}(t)\ _F^2}$ (R2i)
	$\hat{\mathbf{U}}(t) = P_{\Omega}(\mathbf{Y} - \hat{\mathbf{A}}(t)\hat{\mathbf{X}}(t))$ (R3i)
	$\bar{\nu}^p(t) = \left(\frac{\nu^x(t)}{M G_a(t)} + \frac{\nu^a(t)}{L G_x(t)} \right) \frac{N}{\delta}$ (R4i)
	$\nu^p(t) = \bar{\nu}^p(t) + N \nu^a(t) \nu^x(t)$ (R5i)
	$\hat{\mathbf{V}}(t) = \hat{\mathbf{U}}(t) + \frac{\bar{\nu}^p(t)}{\nu^p(t-1) + \nu^w} \hat{\mathbf{V}}(t-1)$ (R6i)
	$\nu^r(t) = G_a(t)(\nu^p(t) + \nu^w)$ (R7i)
	$\nu^q(t) = G_x(t)(\nu^p(t) + \nu^w)$ (R8i)
	$\nu^x(t+1) = \left(\frac{1}{\nu^r(t)} + \frac{1}{\nu^x_0} \right)^{-1}$ (R9i)
	$\hat{\mathbf{X}}(t+1) = \frac{\nu^x(t+1)}{\nu^r(t)} \left((1 - M \delta \nu^a(t) G_a(t)) \hat{\mathbf{X}}(t) + G_a(t) \hat{\mathbf{A}}^H(t) \hat{\mathbf{V}}(t) \right)$ (R10i)
	$\nu^a(t+1) = \left(\frac{1}{\nu^q(t)} + \frac{1}{\nu^a_0} \right)^{-1}$ (R11i)
	$\hat{\mathbf{A}}(t+1) = \frac{\nu^a(t+1)}{\nu^q(t)} \left((1 - L \delta \nu^x(t) G_x(t)) \hat{\mathbf{A}}(t) + G_x(t) \hat{\mathbf{V}}(t) \hat{\mathbf{X}}^H(t) \right)$ (R12i)
	if $\ \hat{\mathbf{U}}(t) - \hat{\mathbf{U}}(t-1)\ _F^2 \leq \tau_{\text{BiG-AMP}} \ \hat{\mathbf{U}}(t)\ _F^2$, stop (R13i)
end	

$P_{\Omega}(\mathbf{Y} - \hat{\mathbf{A}}(t)\hat{\mathbf{X}}(t))$ needs to be computed and stored only at the observed entries $(m, l) \in \Omega$, leading to significant savings⁵ when the observations are highly incomplete (i.e., $|\Omega| \ll ML$). The same is true for the Onsager-corrected residual, $\hat{\mathbf{V}}(t)$. Thus, the algorithm in Table IV involves only three (partial) matrix multiplies [in steps (R3p), (R8p), and (R10p), respectively], each of which can be computed using only $N|\Omega|$ scalar multiplies.

⁵Similar computational savings also occur with incomplete non-Gaussian observations.

We note that Krzakala, Mézard, and Zdeborová recently proposed an AMP-based approach to blind calibration and dictionary learning [42] that bears close similarity⁶ to BiG-AMP under the special case of AWGN-corrupted observations (i.e., $|\Omega| = ML$) and scalar variances. Their derivation differs significantly from that in Section II due to the many simplifications offered by this special case.

C. Zero-Mean iid Gaussian Priors on \mathbf{A} and \mathbf{X}

In this section we will investigate the simplifications that result in the case that both $p_{a_{mn}}$ and $p_{x_{nl}}$ are zero-mean iid Gaussian, i.e.,

$$p_{x_{nl}}(x) = \mathcal{N}(x; 0, \nu_0^x) \quad \forall n, l \quad (91)$$

$$p_{a_{mn}}(a) = \mathcal{N}(a; 0, \nu_0^a) \quad \forall m, n, \quad (92)$$

which, as will be discussed later, is appropriate for matrix completion. In this case, straightforward calculations reveal that $\mathbb{E}\{x_{nl} | r_{nl} = \hat{r}_{nl}; \nu_{nl}^r\} = \hat{r}_{nl} \nu_0^x / (\nu_{nl}^r + \nu_0^x)$ and $\text{var}\{x_{nl} | r_{nl} = \hat{r}_{nl}; \nu_{nl}^r\} = \nu_0^x \nu_{nl}^r / (\nu_{nl}^r + \nu_0^x)$ and, similarly, that $\mathbb{E}\{a_{mn} | q_{mn} = \hat{q}_{mn}; \nu_{mn}^q\} = \hat{q}_{mn} \nu_0^a / (\nu_{mn}^q + \nu_0^a)$ and $\text{var}\{a_{mn} | q_{mn} = \hat{q}_{mn}; \nu_{mn}^q\} = \nu_0^a \nu_{mn}^q / (\nu_{mn}^q + \nu_0^a)$. Combining these iid Gaussian simplifications with the scalar-variance simplifications from Section III-A yields an algorithm whose computational cost is dominated by three matrix multiplies per iteration, each with a cost of MNL scalar multiplies. The precise number of multiplies it consumes depends on the assumed likelihood model that determines steps (R7g)-(R8g).

Additionally incorporating the PIAWGN observations from Section III-B reduces the cost of the three matrix multiplies to only $N|\Omega|$ scalar multiplies each, and yields the ‘‘BiG-AMP-Lite’’ algorithm summarized in Table V, consuming $(3N + 5)|\Omega| + 3(MN + NL) + 29$ multiplies per iteration.

IV. ADAPTIVE DAMPING

The approximations made in the BiG-AMP derivation presented in Section II were well-justified in the large system limit, i.e., the case where $M, N, L \rightarrow \infty$ with fixed $\frac{M}{N}$ and $\frac{L}{N}$. In practical applications, however, these dimensions (especially N) are finite, and hence the algorithm presented in Section II may diverge. In case of compressive sensing, the use of ‘‘damping’’ with GAMP yields provable convergence guarantees with arbitrary matrices [36]. Here, we propose to incorporate damping into BiG-AMP. Moreover, we propose to *adapt* the damping of these variables to ensure that a particular cost criterion decreases monotonically (or near-monotonically), as described in the sequel. The specific damping strategy that we adopt is similar to that described in [51] and coded in [52].

A. Damping

In BiG-AMP, the iteration- t damping factor $\beta(t) \in (0, 1]$ is used to slow the evolution of certain variables, namely $\bar{\nu}_{ml}^p, \nu_{ml}^p$,

⁶The approach in [42] does not compute (or use) $\nu^p(t)$ as given in lines (R4p)-(R5p) of Table IV, but rather uses an empirical average of the squared Onsager-corrected residual in place of our $\nu^p(t) + \nu^w$ throughout their algorithm.

ν_{ml}^s , \hat{s}_{ml} , \hat{x}_{nl} , and \hat{a}_{mn} . To do this, steps (R1), (R3), (R7), and (R8) in Table III are replaced with

$$\bar{\nu}_{ml}^p(t) = \beta(t) \left(\sum_{n=1}^N |\hat{a}_{mn}(t)|^2 \nu_{nl}^x(t) + \nu_{mn}^a(t) |\hat{x}_{nl}(t)|^2 \right) + (1 - \beta(t)) \bar{\nu}_{ml}^p(t-1) \quad (93)$$

$$\nu_{ml}^p(t) = \beta(t) \left(\bar{\nu}_{ml}^p(t) + \sum_{n=1}^N \nu_{mn}^a(t) \nu_{nl}^x(t) \right) + (1 - \beta(t)) \nu_{ml}^p(t-1) \quad (94)$$

$$\nu_{ml}^s(t) = \beta(t) (1 - \nu_{ml}^z(t) / \nu_{ml}^p(t)) / \nu_{ml}^p(t) + (1 - \beta(t)) \nu_{ml}^s(t-1) \quad (95)$$

$$\hat{s}_{ml}(t) = \beta(t) (\hat{z}_{ml}(t) - \hat{p}_{ml}(t)) / \nu_{ml}^p(t) + (1 - \beta(t)) \hat{s}_{ml}(t-1), \quad (96)$$

and the following are inserted between (R8) and (R9):

$$\bar{x}_{nl}(t) = \beta(t) \hat{x}_{nl}(t) + (1 - \beta(t)) \bar{x}_{nl}(t-1) \quad (97)$$

$$\bar{a}_{mn}(t) = \beta(t) \hat{a}_{mn}(t) + (1 - \beta(t)) \bar{a}_{mn}(t-1). \quad (98)$$

The newly defined state variables $\bar{x}_{nl}(t)$ and $\bar{a}_{mn}(t)$ are then used in place of $\hat{x}_{nl}(t)$ and $\hat{a}_{mn}(t)$ in steps (R9)–(R12) [but not (R1)–(R2)] of Table III. A similar approach can be used for the algorithm in Table IV (with the damping applied to $\hat{\mathbf{V}}(t)$ instead of $\hat{\mathbf{S}}(t)$) and those in Table V. Notice that, when $\beta(t) = 1$, the damping has no effect, whereas when $\beta(t) = 0$, all quantities become frozen in t .

B. Adaptive Damping

The idea behind adaptive damping is to monitor a chosen cost criterion $J(t)$ and decrease $\beta(t)$ when the cost has not decreased sufficiently⁷ relative to $\{J(\tau)\}_{\tau=t-1-T}^{t-1}$ for some “step window” $T \geq 0$. This mechanism allows the cost criterion to increase over short intervals of T iterations and in this sense is similar to the procedure used by SpARSA [53]. We now describe how the cost criterion $J(t)$ is constructed, building on ideas in [54].

Notice that, for fixed observations \mathbf{Y} , the joint posterior pdf solves the (trivial) KL-divergence minimization problem

$$p_{\mathbf{X}, \mathbf{A} | \mathbf{Y}} = \arg \min_{b_{\mathbf{X}, \mathbf{A}}} D(b_{\mathbf{X}, \mathbf{A}} \| p_{\mathbf{X}, \mathbf{A} | \mathbf{Y}}). \quad (99)$$

The factorized form (5) of the posterior allows us to write

$$\begin{aligned} & D(b_{\mathbf{X}, \mathbf{A}} \| p_{\mathbf{X}, \mathbf{A} | \mathbf{Y}}) - \log p_{\mathbf{Y}}(\mathbf{Y}) \\ &= \int_{\mathbf{A}, \mathbf{X}} b_{\mathbf{X}, \mathbf{A}}(\mathbf{A}, \mathbf{X}) \log \frac{b_{\mathbf{X}, \mathbf{A}}(\mathbf{A}, \mathbf{X})}{p_{\mathbf{Y} | \mathbf{Z}}(\mathbf{Y} | \mathbf{A}, \mathbf{X}) p_{\mathbf{X}}(\mathbf{X}) p_{\mathbf{A}}(\mathbf{A})} \quad (100) \\ &= D(b_{\mathbf{X}, \mathbf{A}} \| p_{\mathbf{A}} p_{\mathbf{X}}) - \int_{\mathbf{A}, \mathbf{X}} b_{\mathbf{X}, \mathbf{A}}(\mathbf{A}, \mathbf{X}) \log p_{\mathbf{Y} | \mathbf{Z}}(\mathbf{Y} | \mathbf{A}, \mathbf{X}) \quad (101) \end{aligned}$$

⁷The following adaptation procedure is borrowed from GAMPmatlab[52], where it has been established to work well in the context of GAMP-based compressive sensing. When the current cost $J(t)$ is not smaller than the largest cost in the most recent `stepWindow` iterations, then the “step” is deemed unsuccessful, the damping factor $\beta(t)$ is reduced by the factor `stepDec`, and the step is attempted again. These attempts continue until either the cost criterion decreases or the damping factor reaches `stepMin`, at which point the step is considered successful, or the iteration count exceeds T_{\max} or the damping factor reaches `stepTo1`, at which point the algorithm terminates. When a step is deemed successful, the damping factor is increased by the factor `stepInc`, up to the allowed maximum value `stepMax`.

Equations (99) and (101) then imply that

$$p_{\mathbf{X}, \mathbf{A} | \mathbf{Y}} = \arg \min_{b_{\mathbf{X}, \mathbf{A}}} J(b_{\mathbf{X}, \mathbf{A}}) \quad (102)$$

$$J(b_{\mathbf{X}, \mathbf{A}}) \triangleq D(b_{\mathbf{X}, \mathbf{A}} \| p_{\mathbf{A}} p_{\mathbf{X}}) - \mathbb{E}_{b_{\mathbf{X}, \mathbf{A}}} \{ \log p_{\mathbf{Y} | \mathbf{Z}}(\mathbf{Y} | \mathbf{A}, \mathbf{X}) \}. \quad (103)$$

To judge whether a given time- t BiG-AMP approximation “ $b_{\mathbf{X}, \mathbf{A}}(t)$ ” of the joint posterior $p_{\mathbf{X}, \mathbf{A} | \mathbf{Y}}$ is better than the previous approximation $b_{\mathbf{X}, \mathbf{A}}(t-1)$, one could in principle plug the posterior approximation expressions (69)–(70) into (103) and then check whether $J(b_{\mathbf{X}, \mathbf{A}}(t)) < J(b_{\mathbf{X}, \mathbf{A}}(t-1))$. But, since the expectation in (103) is difficult to evaluate, we approximate the cost (103) by using, in place of \mathbf{A}, \mathbf{X} , an independent Gaussian matrix⁸ whose component means and variances are matched to those of \mathbf{A}, \mathbf{X} . Taking the joint BiG-AMP posterior approximation $b_{\mathbf{X}, \mathbf{A}}(t)$ to be the product of the marginals from (69)–(70), the resulting component means and variances are

$$\mathbb{E}_{b_{\mathbf{X}, \mathbf{A}}(t)} \{ [\mathbf{A}, \mathbf{X}]_{ml} \} = \sum_n \mathbb{E}_{b_{\mathbf{X}, \mathbf{A}}(t)} \{ a_{mn} \} \mathbb{E}_{b_{\mathbf{X}, \mathbf{A}}(t)} \{ x_{nl} \} \quad (104)$$

$$= \sum_n \hat{a}_{mn}(t) \hat{x}_{nl}(t) = \bar{p}_{ml}(t) \quad (105)$$

$$\begin{aligned} \text{var}_{b_{\mathbf{X}, \mathbf{A}}(t)} \{ [\mathbf{A}, \mathbf{X}]_{ml} \} &= \sum_n \hat{a}_{mn}^2(t) \nu_{nl}^x(t) + \nu_{mn}^a(t) \hat{x}_{nl}^2(t) \\ &\quad + \nu_{mn}^a(t) \nu_{nl}^x(t) \quad (106) \end{aligned}$$

$$= \nu_{ml}^p(t). \quad (107)$$

In this way, the approximate iteration- t cost becomes

$$\begin{aligned} \hat{J}(t) &= \sum_{n,l} D \left(p_{x_{nl} | r_{nl}}(\cdot | \hat{r}_{nl}(t); \nu_{nl}^r(t)) \middle\| p_{x_{nl}}(\cdot) \right) \\ &\quad + \sum_{m,n} D \left(p_{a_{mn} | q_{mn}}(\cdot | \hat{q}_{mn}(t); \nu_{mn}^q(t)) \middle\| p_{a_{mn}}(\cdot) \right) \\ &\quad - \sum_{m,l} \mathbb{E}_{z_{ml} \sim \mathcal{N}(\bar{p}_{ml}(t); \nu_{ml}^p(t))} \{ \log p_{y_{ml} | z_{ml}}(y_{ml} | z_{ml}) \}. \quad (108) \end{aligned}$$

Intuitively, the first term in (108) penalizes the deviation between the (BiG-AMP approximated) posterior and the assumed prior on \mathbf{X} , the second penalizes the deviation between the (BiG-AMP approximated) posterior and the assumed prior on \mathbf{A} , and the third term rewards highly likely estimates \mathbf{Z} .

V. PARAMETER TUNING AND RANK SELECTION

A. Parameter Tuning via Expectation Maximization

Recall that BiG-AMP requires the specification of priors $p_{\mathbf{X}}(\mathbf{X}) = \prod_{n,l} p_{x_{nl}}(x_{nl})$, $p_{\mathbf{A}}(\mathbf{A}) = \prod_{m,n} p_{a_{mn}}(a_{mn})$, and $p_{\mathbf{Y} | \mathbf{Z}}(\mathbf{Y} | \mathbf{Z}) = \prod_{m,l} p_{y_{ml} | z_{ml}}(y_{ml} | z_{ml})$. In practice, although one may know appropriate families for these distributions, the exact parameters that govern them are generally unknown. For example, one may have good reason to believe apriori that the observations are AWGN corrupted, justifying the choice

⁸The GAMP work [54] uses a similar approximation.

$p_{y_{ml}|z_{ml}}(y_{ml}|z_{ml}) = \mathcal{N}(y_{ml}; z_{ml}, \nu^w)$, but the noise variance ν^w may be unknown. In this section, we outline a methodology that takes a given set of BiG-AMP parameterized priors $\{p_{x_{nl}}(\cdot; \boldsymbol{\theta}), p_{a_{mn}}(\cdot; \boldsymbol{\theta}), p_{y_{ml}|z_{ml}}(y_{ml}|\cdot; \boldsymbol{\theta})\}_{\forall m,n,l}$ and tunes the parameter vector $\boldsymbol{\theta}$ using an expectation-maximization (EM) [37] based approach, with the goal of maximizing the likelihood, i.e., finding $\hat{\boldsymbol{\theta}} \triangleq \arg \max_{\boldsymbol{\theta}} p_{\mathbf{Y}}(\mathbf{Y}; \boldsymbol{\theta})$. The approach presented here can be considered as a generalization of the GAMP-based work [38] to BiG-AMP.

Taking \mathbf{X} , \mathbf{A} , and \mathbf{Z} to be the hidden variables, the EM recursion can be written as [37]

$$\begin{aligned} \hat{\boldsymbol{\theta}}^{k+1} &= \arg \max_{\boldsymbol{\theta}} \mathbb{E} \left\{ \log p_{\mathbf{X}, \mathbf{A}, \mathbf{Z}, \mathbf{Y}}(\mathbf{X}, \mathbf{A}, \mathbf{Z}, \mathbf{Y}; \boldsymbol{\theta}) \middle| \mathbf{Y}; \hat{\boldsymbol{\theta}}^k \right\} \\ &= \arg \max_{\boldsymbol{\theta}} \left\{ \sum_{n,l} \mathbb{E} \left\{ \log p_{x_{nl}}(x_{nl}; \boldsymbol{\theta}) \middle| \mathbf{Y}; \hat{\boldsymbol{\theta}}^k \right\} \right. \\ &\quad + \sum_{m,n} \mathbb{E} \left\{ \log p_{a_{mn}}(a_{mn}; \boldsymbol{\theta}) \middle| \mathbf{Y}; \hat{\boldsymbol{\theta}}^k \right\} \\ &\quad \left. + \sum_{m,l} \mathbb{E} \left\{ \log p_{y_{ml}|z_{ml}}(y_{ml}|z_{ml}; \boldsymbol{\theta}) \middle| \mathbf{Y}; \hat{\boldsymbol{\theta}}^k \right\} \right\} \quad (109) \end{aligned}$$

where for (109) we used the fact $p_{\mathbf{X}, \mathbf{A}, \mathbf{Z}, \mathbf{Y}}(\mathbf{X}, \mathbf{A}, \mathbf{Z}, \mathbf{Y}; \boldsymbol{\theta}) = p_{\mathbf{X}}(\mathbf{X}; \boldsymbol{\theta}) p_{\mathbf{A}}(\mathbf{A}; \boldsymbol{\theta}) p_{\mathbf{Y}|\mathbf{Z}}(\mathbf{Y}|\mathbf{Z}; \boldsymbol{\theta}) \mathbb{1}_{\mathbf{Z}=\mathbf{A}\mathbf{X}}$ and the factorizability of $p_{\mathbf{X}}$, $p_{\mathbf{A}}$, and $p_{\mathbf{Y}|\mathbf{Z}}$. As can be seen from (109), knowledge of the marginal posteriors $\{p_{x_{nl}|\mathbf{Y}}, p_{a_{mn}|\mathbf{Y}}, p_{z_{ml}|\mathbf{Y}}\}_{\forall m,n,l}$ is sufficient to compute the EM update. Since the exact marginal posteriors are unknown, we employ BiG-AMP's approximations from (69), (70), and (35) for approximate EM. In addition, we adopt the "incremental" update strategy from [55], where the maximization over $\boldsymbol{\theta}$ is performed one element at a time while holding the others fixed.

As a concrete example, consider updating the noise variance ν^w under the PIAWGN model (80). Equation (109) suggests

$$(\nu^w)^{k+1} = \arg \max_{\nu^w} \sum_{(m,l) \in \Omega} \int_{z_{ml}} p_{z_{ml}|\mathbf{Y}}(z_{ml}|\mathbf{Y}) \times \log \mathcal{N}(y_{ml}; z_{ml}, \nu^w), \quad (110)$$

where the true marginal posterior $p_{z_{ml}|\mathbf{Y}}(\cdot|\mathbf{Y})$ is replaced with the most recent BiG-AMP approximation $p_{z_{ml}|\rho_{ml}}(\cdot|\hat{\rho}_{ml}(T_{\max}); \nu_{ml}^p(T_{\max}), \hat{\boldsymbol{\theta}}^k)$, where "most recent" is with respect to both EM and BiG-AMP iterations. Zeroing the derivative of the sum in (110) with respect to ν^w ,

$$(\nu^w)^{k+1} = \frac{1}{|\Omega|} \sum_{(m,l) \in \Omega} (y_{ml} - \hat{z}_{ml}(T_{\max}))^2 + \nu_{ml}^z(T_{\max}), \quad (111)$$

where $\hat{z}_{ml}(t)$ and $\nu_{ml}^z(t)$ are the BiG-AMP approximated posterior mean and variance from (33)–(34).

The overall procedure can be summarized as follows. From a suitable initialization $\hat{\boldsymbol{\theta}}^0$, BiG-AMP is run using the priors $\{p_{x_{nl}}(\cdot; \hat{\boldsymbol{\theta}}^0), p_{a_{mn}}(\cdot; \hat{\boldsymbol{\theta}}^0), p_{y_{ml}|z_{ml}}(y_{ml}|\cdot; \hat{\boldsymbol{\theta}}^0)\}_{\forall m,n,l}$ and iterated to completion, yielding approximate marginal posteriors on $\{x_{nl}, a_{mn}, z_{ml}\}_{\forall m,n,l}$. These posteriors are

used in (109) to update the parameters $\boldsymbol{\theta}$ one element at a time, yielding $\hat{\boldsymbol{\theta}}^1$. BiG-AMP is then run using the priors $\{p_{x_{nl}}(\cdot; \hat{\boldsymbol{\theta}}^1), p_{a_{mn}}(\cdot; \hat{\boldsymbol{\theta}}^1), p_{y_{ml}|z_{ml}}(y_{ml}|\cdot; \hat{\boldsymbol{\theta}}^1)\}_{\forall m,n,l}$, and so on. A detailed discussion in the context of GAMP, along with explicit update equations for the parameters of Bernoulli-Gaussian-mixture pdfs, can be found in [38].

B. Rank Selection

BiG-AMP and EM-BiG-AMP, as described up to this point, require the specification of the rank N , i.e., the number of columns in \mathbf{A} (and rows in \mathbf{X}) in the matrix factorization $\mathbf{Z} = \mathbf{A}\mathbf{X}$. Since, in many applications, the best choice of N is difficult to specify in advance, we now describe two procedures to estimate N from the data \mathbf{Y} , building on well-known rank-selection procedures.

1) *Penalized Log-Likelihood Maximization*: Consider a set of possible models $\{\mathcal{H}_N\}_{N=1}^{\bar{N}}$ for the observation \mathbf{Y} where, under \mathcal{H}_N , EM-BiG-AMP estimates $\boldsymbol{\Theta}_N = \{\mathbf{A}_N, \mathbf{X}_N, \boldsymbol{\theta}\}$. Here, the subscripts on \mathbf{A}_N and \mathbf{X}_N indicate the restriction to N columns and rows, $\boldsymbol{\theta}$ refers to the vector of parameters defined in Section V-A, and the subscript on $\boldsymbol{\Theta}_N$ indicates the dependence of the overall number of parameters in $\boldsymbol{\Theta}_N$ with the rank N . Because the selection rule $\hat{N} = \arg \max_N p_{\mathbf{Y}}(\mathbf{Y}; \mathcal{H}_N)$ is typically intractable, several well-known rules of the form

$$\hat{N} = \arg \max_{N=1, \dots, \bar{N}} 2 \log p_{\mathbf{Y}|\boldsymbol{\Theta}_N}(\mathbf{Y}|\hat{\boldsymbol{\Theta}}_N) - \eta(N) \quad (112)$$

have been developed, such as the Bayesian Information Criterion (BIC) and Akaike's Information Criterion (AIC) [56]. In (112), $\hat{\boldsymbol{\Theta}}_N$ is the ML estimate of $\boldsymbol{\Theta}_N$ under \mathbf{Y} , and $\eta(\cdot)$ is a penalty function that depends on the *effective* number of scalar parameters N_{eff} estimated under model \mathcal{H}_N (which depends on N) and possibly on the number of scalar parameters $|\Omega|$ that make up the observation \mathbf{Y} .

Applying this methodology to EM-BiG-AMP, where $p_{\mathbf{Y}|\boldsymbol{\Theta}_N}(\mathbf{Y}|\boldsymbol{\Theta}_N) = p_{\mathbf{Y}|\mathbf{Z}}(\mathbf{Y}|\mathbf{A}_N \mathbf{X}_N; \boldsymbol{\theta})$, we obtain the rank-selection rule

$$\hat{N} = \arg \max_{N=1, \dots, \bar{N}} 2 \log p_{\mathbf{Y}|\mathbf{Z}}(\mathbf{Y}|\hat{\mathbf{A}}_N \hat{\mathbf{X}}_N; \hat{\boldsymbol{\theta}}) - \eta(N). \quad (113)$$

Since N_{eff} depends on the application (e.g., matrix completion, robust PCA, dictionary learning), detailed descriptions of $\eta(\cdot)$ are postponed to [1].

To perform the maximization over N in (113), we start with a small hypothesis N_1 and run EM-BiG-AMP to completion, generating the (approximate) MMSE estimates $\hat{\mathbf{A}}_{N_1}$, $\hat{\mathbf{X}}_{N_1}$ and ML estimate $\hat{\boldsymbol{\theta}}$, which are then used to evaluate⁹ the penalized log-likelihood in (113). The N hypothesis is then increased by a fixed value (i.e., $N_2 = N_1 + \text{rankStep}$), initializations of $(\mathbf{A}_{N_2}, \mathbf{X}_{N_2}, \boldsymbol{\theta})$ are chosen based on the previously computed $(\hat{\mathbf{A}}_{N_1}, \hat{\mathbf{X}}_{N_1}, \hat{\boldsymbol{\theta}})$, and EM-BiG-AMP is run to completion,

⁹Since we compute approximate MMSE estimates rather than ML estimates, we are in fact evaluating a lower bound on the penalized log-likelihood.

yielding estimates $(\hat{\mathbf{A}}_{N_2}, \hat{\mathbf{X}}_{N_2}, \hat{\boldsymbol{\theta}})$ with which the penalized likelihood is again evaluated. This process continues until either the value of the penalized log-likelihood decreases, in which case \hat{N} is set at the previous (i.e., maximizing) hypothesis of N , or the maximum-allowed rank \bar{N} is reached.

2) *Rank Contraction*: We now describe an alternative rank-selection procedure that is appropriate when \mathbf{Z} has a “cliff” in its singular value profile and which is reminiscent of that used in LMaFit [13]. In this approach, EM-BiG-AMP is initially configured to use the maximum-allowed rank, i.e., $N = \bar{N}$. After the first EM iteration, the singular values $\{\sigma_n\}$ of the estimate $\hat{\mathbf{X}}$ and the corresponding pairwise ratios $R_n = \sigma_n/\sigma_{n+1}$ are computed,¹⁰ from which a candidate rank estimate $\hat{N} = \arg \max_n R_n$ is identified, corresponding to the largest gap in successive singular values. However, this candidate is accepted only if this maximizing ratio exceeds the average ratio by the user-specified parameter τ_{MOS} (e.g., $\tau_{\text{MOS}} = 5$), i.e., if

$$R_{\hat{N}} > \frac{\tau_{\text{MOS}}}{\bar{N} - 2} \sum_{i \neq \hat{N}} R_i, \quad (114)$$

and if \hat{N}/\bar{N} is sufficiently small. Increasing τ_{MOS} makes the approach less prone to selecting an erroneous rank during the first few iterations, but making the value too large prevents the algorithm from detecting small gaps between the singular values. If \hat{N} is accepted, then the matrices \mathbf{A} and \mathbf{X} are pruned to size \hat{N} and EM-BiG-AMP is run to convergence. If not, EM-BiG-AMP is run for one more iteration, after which a new candidate \hat{N} is identified and checked for acceptance, and so on.

In many cases, a rank candidate is accepted after a small number of iterations, and thus only a few SVDs need be computed. This procedure has the advantage of running EM-BiG-AMP to convergence only once, rather than several times under different hypothesized ranks. However, when the singular values of \mathbf{Z} decay smoothly, this procedure can mis-estimate the rank, as discussed in [13].

VI. CONCLUSION

In this work, we proposed and derived BiG-AMP, an extension of the G-AMP algorithm [33] originally proposed for high-dimensional generalized-*linear* regression in the context of compressive sensing, to generalized-*bilinear* regression, with applications in matrix completion, robust PCA, dictionary learning, and related matrix-factorization problems. In addition, we proposed an adaptive damping mechanism to aid convergence under realistic problem sizes, an expectation-maximization (EM)-based method to automatically tune the parameters of the assumed priors, and two rank-selection strategies. In Part II [1] of this two-part work, we detail the application of BiG-AMP to matrix completion, robust PCA, and dictionary learning, and we present the results of an extensive numerical investigation into the performance of BiG-AMP on both synthetic and real-world datasets. The results in [1] demonstrate that BiG-AMP yields excellent reconstruction

accuracy (often best in class) while maintaining competitive runtimes, and that the proposed EM and rank-selection strategies successfully avoid the need to tune algorithmic parameters.

APPENDIX A

Here we derive (31)–(32), which are stated without a detailed derivation in [33]. Recalling (21) and omitting the ml subscripts for brevity, it can be seen that

$$\begin{aligned} H'(\hat{q}, \nu^q; y) &= \frac{\partial}{\partial \hat{q}} \log \int p_{y|z}(y|z) \frac{1}{\sqrt{2\pi\nu^q}} \exp\left(-\frac{1}{2\nu^q}(z - \hat{q})^2\right) dz \\ &= \frac{\partial}{\partial \hat{q}} \left\{ \log \int \exp\left(\log p_{y|z}(y|z) - \frac{z^2}{2\nu^q} + \frac{\hat{q}z}{\nu^q}\right) dz - \frac{\hat{q}^2}{2\nu^q} \right\} \\ &= -\frac{\hat{q}}{\nu^q} + \frac{\partial}{\partial \hat{q}} \log \int \exp(\phi(u) + \hat{q}u) du \text{ via } u \triangleq \frac{z}{\nu^q} \end{aligned} \quad (115)$$

for an appropriately defined function $\phi(\cdot)$. Now, defining $p_{u|q}(u|\hat{q}) \triangleq Z(\hat{q})^{-1} \exp(\phi(u) + \hat{q}u)$ with normalization term $Z(\hat{q}) \triangleq \int \exp(\phi(u) + \hat{q}u) du$, simple calculus yields

$$\frac{\partial}{\partial \hat{q}} \log Z(\hat{q}) = \mathbb{E}\{u|q = \hat{q}\} \quad (116)$$

$$\frac{\partial^2}{\partial \hat{q}^2} \log Z(\hat{q}) = \text{var}\{u|q = \hat{q}\}. \quad (117)$$

Thus, from (115) and (116) it follows that

$$\begin{aligned} H'(\hat{q}, \nu^q; y) &= -\frac{\hat{q}}{\nu^q} + \int u \frac{\exp(\phi(u) + \hat{q}u)}{Z(\hat{q})} du \\ &= -\frac{\hat{q}}{\nu^q} + \int \frac{z}{\nu^q} \frac{\exp\left(\log p_{y|z}(y|z) - \frac{z^2}{2\nu^q} + \frac{\hat{q}z}{\nu^q}\right) dz}{Z(\hat{q})} \frac{1}{\nu^q} \\ &= -\frac{\hat{q}}{\nu^q} + \frac{1}{\nu^q} \int z \frac{p_{y|z}(y|z) \mathcal{N}(z; \hat{q}, \nu^q)}{\int p_{y|z}(y|\bar{z}) \mathcal{N}(\bar{z}; \hat{q}, \nu^q) d\bar{z}} dz. \end{aligned} \quad (118)$$

Equation (31) is then established by applying definitions (33) and (35) to (118).

Similarly, from (115) and (117),

$$\begin{aligned} -H''(\hat{q}, \nu^q; y) &= \frac{\partial}{\partial \hat{q}} \left\{ \frac{\hat{q}}{\nu^q} - \frac{\partial}{\partial \hat{q}} \log Z(\hat{q}) \right\} = \frac{1}{\nu^q} - \text{var}\{u|q = \hat{q}\} \\ &= \frac{1}{\nu^q} - \int (u - \mathbb{E}\{u|q = \hat{q}\})^2 \frac{\exp(\phi(u) + \hat{q}u)}{Z(\hat{q})} du \\ &= \frac{1}{\nu^q} - \frac{1}{(\nu^q)^2} \int (z - \hat{z})^2 \frac{p_{y|z}(y|z) \mathcal{N}(z; \hat{q}, \nu^q)}{\int p_{y|z}(y|\bar{z}) \mathcal{N}(\bar{z}; \hat{q}, \nu^q) d\bar{z}} dz, \end{aligned} \quad (119)$$

¹⁰In some cases the singular values of $\hat{\mathbf{A}}$ could be used instead.

where \hat{z} is the expectation from (33). Equation (32) is then established by applying the definitions (34) and (35) to (119).

APPENDIX B

Here we explain the approximations (65)–(66). The term neglected in going from (45) to (65) can be written using (31)–(32) as

$$\begin{aligned} & \sum_{m=1}^M \nu_{mn}^a(t) (\hat{s}_{ml}^2(t) - \nu_{ml}^s(t)) \\ &= \sum_{m=1}^M \nu_{mn}^a(t) \left[\frac{(\hat{z}_{ml}(t) - \hat{p}_{ml}(t))^2 + \nu_{ml}^z(t)}{\nu_{ml}^p(t)^2} - \frac{1}{\nu_{ml}^p(t)} \right] \end{aligned} \quad (120)$$

$$= \sum_{m=1}^M \frac{\nu_{mn}^a(t)}{\nu_{ml}^p(t)} \left[\mathbb{E} \left\{ \frac{(z_{ml} - \hat{p}_{ml}(t))^2}{\nu_{ml}^p(t)} \right\} - 1 \right] \quad (121)$$

where the expectations are taken over $z_{ml} \sim p_{z_{ml}|p_{ml}}(\cdot | \hat{p}_{ml}(t); \nu_{ml}^p(t))$ from (35). For GAMP, [57, Sec. VI.D] clarifies that, in the large system limit, under i.i.d. priors and scalar variances, the true z_{ml} and the iterates $\hat{p}_m(t)$ converge empirically to a pair of random variables (z, p) that satisfy $p_{z|p}(z | \hat{p}(t)) = \mathcal{N}(z; \hat{p}(t), \nu^p(t))$. This result leads us to believe that the expectation in (121) is approximately unit-valued when averaged over m , and thus (121) is approximately zero-valued. Similar reasoning applies to (66).

ACKNOWLEDGMENT

The authors would like to thank Sundeep Rangan, Florent Krzakala, and Lenka Zdeborová for insightful discussions on various aspects of AMP-based inference. We would also like to thank Subhojit Som, Jeremy Vila, and Justin Ziniel for helpful discussions about EM and turbo methods for AMP.

REFERENCES

- [1] J. T. Parker, P. Schniter, and V. Cevher, "Bilinear generalized approximate message passing—Part II: Applications," *IEEE Trans. Signal Process.*, vol. 62, no. 22, pp. 5854–5867, 2014.
- [2] P. Schniter and V. Cevher, "Approximate message passing for bilinear models," in *Proc. Workshop Signal Process. Adapt. Sparse Struct. Repr. (SPARS)*, Edinburgh, Scotland, Jun. 2011, pp. 68–68.
- [3] P. Schniter, J. T. Parker, and V. Cevher, "Bilinear generalized approximate message passing (BiG-AMP) for matrix recovery problems," presented at the Inf. Theory Appl. Workshop (ITA), La Jolla, CA, USA, Feb. 2012.
- [4] J. T. Parker and P. Schniter, "Bilinear generalized approximate message passing (BiG-AMP) for matrix completion," presented at the Asilomar Conf., Pacific Grove, CA, USA, Nov. 2012.
- [5] J. Cai, E. Candès, and Z. Shen, "A singular value thresholding algorithm for matrix completion," *SIAM J. Optim.*, vol. 20, no. 4, pp. 1956–1982, 2010 [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/080738970>
- [6] S. Ma, D. Goldfarb, and L. Chen, "Fixed point and bregman iterative methods for matrix rank minimization," *Math. Programm.*, vol. 128, no. 1-2, pp. 321–353, 2011 [Online]. Available: <http://dx.doi.org/10.1007/s10107-009-0306-5>
- [7] Z. Zhou, X. Li, J. Wright, E. Candès, and Y. Ma, "Stable principal component pursuit," in *Proc. IEEE Int. Symp. Inf. Theory Process. (ISIT)*, Jun. 2010, pp. 1518–1522.
- [8] B. Recht, M. Fazel, and P. Parrilo, "Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization," *SIAM Rev.*, vol. 52, no. 3, pp. 471–501, 2010 [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/070697835>
- [9] Z. Lin, M. Chen, L. Wu, and Y. Ma, "The augmented Lagrange multiplier method for exact recovery of corrupted low-rank matrices," arXiv preprint, arXiv:1009.5055 [Online]. Available: <http://arxiv.org/abs/1009.5055>
- [10] M. Tao and X. Yuan, "Recovering low-rank and sparse components of matrices from incomplete and noisy observations," *SIAM J. Optim.*, vol. 21, no. 1, pp. 57–81, 2011.
- [11] T. Zhou and D. Tao, L. Getoor and T. Scheffer, Eds., "Godec: Randomized low-rank and sparse matrix decomposition in noisy case," in *Proc. 28th Int. Conf. Mach. Learn. (ICML-11), Ser. ICML'11, ACM*, New York, NY, USA, Jun. 2011, pp. 33–40.
- [12] H. Ghasemi, M. Malek-Mohammadi, M. Babaei-Zadeh, and C. Jutten, "SRF: Matrix completion based on smoothed rank function," presented at the IEEE Int. Conf. Acoust. Speech, Signal Process., Prague, Czech Republic, May 2011.
- [13] Z. Wen, W. Yin, and Y. Zhang, "Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm," *Math. Programm. Comput.*, vol. 4, pp. 333–361, 2012 [Online]. Available: <http://dx.doi.org/10.1007/s12532-012-0044-1>
- [14] A. Kyrillidis and V. Cevher, "Matrix recipes for hard thresholding methods," *J. Math. Imag. Vis.*, vol. 48, pp. 235–265, 2014.
- [15] G. Marjanovic and V. Solo, "On optimization and matrix completion," *IEEE Trans. Signal Process.*, vol. 60, no. 11, pp. 5714–5724, 2012.
- [16] J. Haldar and D. Hernando, "Rank-constrained solutions to linear matrix equations using Power Factorization," *IEEE Signal Process. Lett.*, vol. 16, no. 7, pp. 584–587, Jul. 2009.
- [17] L. Balzano, R. Nowak, and B. Recht, "Online identification and tracking of subspaces from highly incomplete information," in *Proc. Allerton Conf. Commun. Control Comput.*, Sep. 2010, pp. 704–711.
- [18] R. Keshavan, A. Montanari, and S. Oh, "Matrix completion from a few entries," *IEEE Trans. Inf. Theory*, vol. 56, no. 6, pp. 2980–2998, Jun. 2010.
- [19] W. Dai and O. Milenkovic, "SET: An algorithm for consistent matrix completion," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Processing (ICASSP)*, 2010, pp. 3646–3649.
- [20] J. He, L. Balzano, and J. Lui, "Online robust subspace tracking from partial information" arXiv:1109.3827 2011.
- [21] R. Salakhutdinov and A. Mnih, "Probabilistic matrix factorization," *Adv. Neural Inf. Process. Syst.*, vol. 20, pp. 1257–1264, 2008.
- [22] X. Ding, L. He, and L. Carin, "Bayesian robust principal component analysis," *IEEE Trans. Image Process.*, vol. 20, no. 12, pp. 3419–3430, Dec. 2011.
- [23] Y. Lim and Y. Teh, "Variational Bayesian approach to movie rating prediction," in *Proc. KDD Cup Workshop*, 2007, pp. 15–21, Citeseer.
- [24] S. D. Babacan, M. Luessi, R. Molina, and A. K. Katsaggelos, "Sparse Bayesian methods for low-rank matrix estimation," *IEEE Trans. Signal Process.*, vol. 60, no. 8, pp. 3964–3977, Aug. 2012.
- [25] M. Tipping and C. Bishop, "Probabilistic principal component analysis," *J. Roy. Statist. Soc., Series B (Statist. Methodol.)*, vol. 61, no. 3, pp. 611–622, 1999.
- [26] F. Léger, G. Yu, and G. Sapiro, "Efficient matrix completion with Gaussian models," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, 2011, pp. 1113–1116.
- [27] N. Wang, T. Yao, J. Wang, and D. Yeung, A. Fitzgibbon, Ed., "A probabilistic approach to robust matrix factorization," in *Proc. Eur. Conf. Comput. Vis.*, 2012, vol. VII, pp. 126–139.
- [28] B.-H. Kim, A. Yedla, and H. Pfister, "Imp: A message-passing algorithm for matrix completion," in *Proc. 6th Int. Symp. Turbo Codes Iterative Inf. Process. (ISTC)*, 2010, pp. 462–466.
- [29] B. J. Frey and D. J. C. MacKay, "A revolution: Belief propagation in graphs with cycles," in *Proc. Neural Inf. Process. Syst. Conf.*, Denver, CO, USA, 1997, pp. 479–485.
- [30] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing," *Proc. Nat. Acad. Sci.*, vol. 106, no. 45, pp. 18914–18919, Nov. 2009.
- [31] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing: I. motivation and construction," in *Proc. Inf. Theory Workshop*, Cairo, Egypt, Jan. 2010, pp. 1–5.

- [32] A. Montanari, “Graphical models concepts in compressed sensing,” in *Compressed Sensing: Theory and Application*, Y. C. Eldar and G. Kutyniok, Eds. Cambridge, U.K.: Cambridge Univ. Press, 2012.
- [33] S. Rangan, “Generalized approximate message passing for estimation with random linear mixing,” in *Proc. IEEE Int. Symp. Inf. Theory*, Saint Petersburg, Russia, Aug. 2011, pp. 2168–2172.
- [34] M. Bayati and A. Montanari, “The dynamics of message passing on dense graphs, with applications to compressed sensing,” *IEEE Trans. Inf. Theory*, vol. 57, no. 2, pp. 764–785, Feb. 2011.
- [35] A. Javanmard and A. Montanari, “State evolution for general approximate message passing algorithms, with applications to spatial coupling,” *Inf. Inference*, vol. 2, no. 2, pp. 115–144, 2013.
- [36] S. Rangan, P. Schniter, and A. Fletcher, “On the convergence of generalized approximate message passing with arbitrary matrices,” presented at the IEEE Int. Symp. Inf. Theory, Honolulu, HI, USA, Jul. 2014.
- [37] A. Dempster, N. M. Laird, and D. B. Rubin, “Maximum-likelihood from incomplete data via the EM algorithm,” *J. Roy. Statist. Soc.*, vol. 39, pp. 1–17, 1977.
- [38] J. P. Vila and P. Schniter, “Expectation-maximization Gaussian-mixture approximate message passing,” *IEEE Trans. Signal Process.*, vol. 61, no. 19, pp. 4658–4672, Oct. 2013.
- [39] P. Schniter, “Turbo reconstruction of structured sparse signals,” in *Proc. Conf. Inf. Sci. Syst.*, Princeton, NJ, USA, Mar. 2010, pp. 1–6.
- [40] J. Vila, P. Schniter, and J. Meola, “Hyperspectral image unmixing via bilinear generalized approximate message passing,” *Proc. SPIE*, vol. 8743, no. 87430Y, p. 9, 2013.
- [41] S. Rangan and A. K. Fletcher, “Iterative estimation of constrained rank-one matrices in noise,” in *Proc. IEEE Int. Symp. Inf. Theory*, 2012, pp. 1246–1250.
- [42] F. Krzakala, M. Mézard, and L. Zdeborová, “Phase diagram and approximate message passing for blind calibration and dictionary learning,” in *Proc. IEEE Int. Symp. Inf. Theory*, 2013, pp. 659–663.
- [43] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA, USA: Morgan Kaufmann, 1988.
- [44] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, “Factor graphs and the sum-product algorithm,” *IEEE Trans. Inf. Theory*, vol. 47, pp. 498–519, Feb. 2001.
- [45] G. F. Cooper, “The computational complexity of probabilistic inference using Bayesian belief networks,” *Artif. Intell.*, vol. 42, pp. 393–405, 1990.
- [46] W. T. Freeman, E. C. Pasztor, and O. T. Carmichael, “Learning low-level vision,” *Intl. J. Comput. Vis.*, vol. 40, no. 1, pp. 25–47, Oct. 2000.
- [47] R. J. McEliece, D. J. C. MacKay, and J.-F. Cheng, “Turbo decoding as an instance of pearl’s ‘belief propagation’ algorithm,” *IEEE J. Sel. Areas Commun.*, vol. 16, no. 2, pp. 140–152, Feb. 1998.
- [48] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms*. New York, NY, USA: Cambridge Univ. Press, 2003.
- [49] J. Boutros and G. Caire, “Iterative multiuser joint decoding: Unified framework and asymptotic analysis,” *IEEE Trans. Inf. Theory*, vol. 48, no. 7, pp. 1772–1793, Jul. 2002.
- [50] J. T. Parker, V. Cevher, and P. Schniter, “Compressive sensing under matrix uncertainties: An approximate message passing approach,” in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Pacific Grove, CA, USA, Nov. 2011, pp. 804–808.
- [51] P. Schniter and S. Rangan, “Compressive phase retrieval via generalized approximate message passing,” presented at the Allerton Conf. Commun., Control, Comput., Monticello, IL, USA, Oct. 2012.
- [52] S. Rangan *et al.*, GAMPmatlab [Online]. Available: <https://sourceforge.net/projects/gampmatlab/>
- [53] S. Wright, R. Nowak, and M. Figueiredo, “Sparse reconstruction by separable approximation,” *IEEE Trans. Signal Process.*, vol. 57, no. 7, pp. 2479–2493, Jul. 2009.
- [54] S. Rangan, P. Schniter, E. Riegler, A. Fletcher, and V. Cevher, “Fixed points of generalized approximate message passing with arbitrary matrices,” in *Proc. IEEE Int. Symp. Inf. Theory*, Istanbul, Turkey, Jul. 2013, pp. 664–668.
- [55] R. Neal and G. Hinton, “A view of the EM algorithm that justifies incremental, sparse, and other variants,” in *Learning in Graphical Models*, M. I. Jordan, Ed. Cambridge, MA, USA: MIT Press, 1999, pp. 355–368.
- [56] P. Stoica and Y. Selen, “Model-order selection: a review of information criterion rules,” *IEEE Signal Process. Mag.*, vol. 21, no. 4, pp. 36–47, Jul. 2004.
- [57] S. Rangan, “Generalized approximate message passing for estimation with random linear mixing,” Aug. 2012, arXiv:1010.5141v2.



Jason T. Parker (M’06) received the B.S., M.S., and Ph.D. degrees in electrical and computer engineering from The Ohio State University, Columbus, in 2004, 2006, and 2014, respectively. Since 2006, he has been a research engineer with the U.S. Air Force Research Laboratory Sensors Directorate. His research interests currently include compressive sensing, adaptive signal processing, and inverse problems, with applications to radar target detection and imaging.



Philip Schniter (F’14) received the B.S. and M.S. degrees in Electrical Engineering from the University of Illinois at Urbana-Champaign in 1992 and 1993, respectively, and the Ph.D. degree in Electrical Engineering from Cornell University in Ithaca, NY, in 2000.

From 1993 to 1996 he was employed by Tektronix Inc. in Beaverton, OR as a systems engineer. After receiving the Ph.D. degree, he joined the Department of Electrical and Computer Engineering at The Ohio State University, Columbus, where he is currently a 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 Supélec, Gif-sur-Yvette, France.

In 2003, Dr. Schniter received the National Science Foundation CAREER Award. His areas of interest currently include statistical signal processing, wireless communications and networks, and machine learning.



Volkan Cevher (SM’10) received the B.S. (valedictorian) degree in electrical engineering in 1999 from Bilkent University in Ankara, Turkey, and he received the Ph.D. degree in Electrical and Computer Engineering in 2005 from the Georgia Institute of Technology in Atlanta. He held research scientist positions at the University of Maryland, College Park from 2006 to 2007 and at Rice University in Houston, Texas, from 2008 to 2009. Currently, he is an Assistant Professor at the Swiss Federal Institute of Technology Lausanne with a complimentary

appointment at the Electrical and Computer Engineering Department at Rice University. His research interests include signal processing theory, machine learning, graphical models, and information theory. He received a Best Paper Award at SPARS in 2009 and an ERC StG in 2011.