ADAPTIVE DAMPING AND MEAN REMOVAL FOR
THE GENERALIZED APPROXIMATE MESSAGE PASSING ALGORITHM

Jeremy Vila[*]   *Philip Schniter*[*]   *Sundeep Rangan*[†]   *Florent Krzakala*[‡]   *Lenka Zdeborová*[°]

[*] Dept. of ECE, The Ohio State University, Columbus, OH 43202, USA.
[†] Dept. of ECE, Polytechnic Institute of New York University, Brooklyn, NY 11201, USA.
[‡] Sorbonne Universités, UPMC Univ Paris 06 and École Normale Supérieure, 75005 Paris, France.
[°] Institut de Physique Théorique, CEA Saclay, and CNRS URA 2306, 91191 Gif-sur-Yvette, France.

ABSTRACT

The generalized approximate message passing (GAMP) algorithm is an efficient method of MAP or approximate-MMSE estimation of x observed from a noisy version of the transform coefficients $z = Ax$. In fact, for large zero-mean i.i.d sub-Gaussian A, GAMP is characterized by a state evolution whose fixed points, when unique, are optimal. For generic A, however, GAMP may diverge. In this paper, we propose adaptive-damping and mean-removal strategies that aim to prevent divergence. Numerical results demonstrate significantly enhanced robustness to non-zero-mean, rank-deficient, column-correlated, and ill-conditioned A.

Index Terms— Approximate message passing, belief propagation, compressed sensing.

1. INTRODUCTION

Consider estimating a realization $x \in \mathbb{R}^N$ of a random vector x with statistically independent components $\mathsf{x}_n \sim p_{\mathsf{x}_n}$ from observations $y = [y_m] \in \mathbb{R}^M$ that are conditionally independent given the transform outputs

$$z = Ax \qquad (1)$$

for some known matrix $A = [a_{mn}] \in \mathbb{R}^{M \times N}$. Here, the likelihood function can be written as $p_{\mathsf{y}|\mathsf{z}}(y|Ax)$ with separable $p_{\mathsf{y}|\mathsf{z}}$, i.e., $p_{\mathsf{y}|\mathsf{z}}(y|z) = \prod_{m=1}^{M} p_{\mathsf{y}_m|\mathsf{z}_m}(y_m|z_m)$. Such problems arise in a range of applications including statistical regression, inverse problems, and compressive sensing. Note that, for clarity, we use sanserif fonts (e.g., x, x_n) to denote random quantities and serif fonts (e.g., x, x_n) to denote deterministic ones.

Assuming knowledge of the prior $p_{\mathsf{x}}(x) = \prod_{n=1}^{N} p_{\mathsf{x}_n}(x_n)$ and likelihood $p_{\mathsf{y}|\mathsf{z}}(y|z)$, typical estimation goals are to compute the minimum mean-squared error (MMSE) estimate $\widehat{x}_{\mathsf{MMSE}} \triangleq \int_{\mathbb{R}^N} x p_{\mathsf{x}|\mathsf{y}}(x|y) dx$ or the maximum a posteriori (MAP) estimate $\widehat{x}_{\mathsf{MAP}} \triangleq \arg\max_x p_{\mathsf{x}|\mathsf{y}}(x|y) = \arg\min_x J_{\mathsf{MAP}}(x)$ for the MAP cost

$$J_{\mathsf{MAP}}(\widehat{x}) \triangleq -\ln p_{\mathsf{y}|\mathsf{z}}(y|A\widehat{x}) - \ln p_{\mathsf{x}}(\widehat{x}). \qquad (2)$$

Recently, the *generalized approximate message passing* (GAMP) algorithm [1] has been proposed as a means of tackling these two problems in the case that M and N are large. Essentially, GAMP uses a high-dimensional approximation of loopy belief propagation to convert the MMSE or MAP inference problems into a sequence of tractable scalar inference problems.

GAMP is well motivated in the case that A is a realization of a large random matrix with i.i.d zero-mean sub-Gaussian entries. For such A, in the large-system limit (i.e., $M, N \to \infty$ for fixed $M/N \in \mathbb{R}_+$), GAMP is characterized by a state evolution whose fixed points, when unique, are MMSE or MAP optimal [1–3]. Furthermore, for *generic* A, it has been shown [4] that MAP-GAMP's fixed points coincide with the critical points of the cost function (2) and that MMSE-GAMP's fixed points coincide with those of a Bethe free entropy [5], as discussed in detail in Section 2.2.

For generic A, however, GAMP may not reach its fixed points, i.e., it may diverge (e.g., [6]). The convergence of GAMP has been fully characterized in [7] for the simple case that p_{x_n} and $p_{\mathsf{y}_m|\mathsf{z}_m}$ are Gaussian. There, it was shown that Gaussian-GAMP converges if and only if the peak-to-average ratio of the squared singular values of A is sufficiently small. A *damping* modification was then proposed in [7] that guarantees the convergence of Gaussian-GAMP with arbitrary A, at the expense of a slower convergence rate. For strictly logconcave p_{x_n} and $p_{\mathsf{y}_m|\mathsf{z}_m}$, the *local* convergence of GAMP was also characterized in [7]. However, the global convergence of GAMP under generic A, p_{x_n}, and $p_{\mathsf{y}_m|\mathsf{z}_m}$ is not yet understood.

Because of its practical importance, prior work has attempted to robustify the convergence of GAMP in the face of "difficult" A (e.g., high peak-to-average singular values) for generic p_{x_n} and $p_{\mathsf{y}_m|\mathsf{z}_m}$. For example, "swept" GAMP (SwAMP) [8] updates the estimates of $\{\mathsf{x}_n\}_{n=1}^N$ and $\{\mathsf{z}_m\}_{m=1}^M$ sequentially, in contrast to GAMP, which updates them in parallel. Relative to GAMP, experiments in [8] show that SwAMP is much more robust to difficult A, but it is slower and cannot facilitate fast implementations of A like an FFT. As another example, the public-domain GAMPmatlab implementation [9] has included "adaptive damping" and "mean removal" mechanisms for some time, but they have never been described in the literature.

In this paper, we detail the most recent versions of GAMPmatlab's adaptive damping and mean-removal mechanisms, and we experimentally characterize their performance on non-zero-mean, rank-deficient, column-correlated, and ill-conditioned A matrices. Our results show improved robustness relative to SwAMP and enhanced convergence speed.

2. ADAPTIVELY DAMPED GAMP

Damping is commonly used in loopy belief propagation to "slow down" the updates in an effort to promote convergence. (See, e.g., [10] for damping applied to the sum-product algorithm and [7, 9, 11] for damping applied to GAMP.) However, since not enough damping

Table 1. The adaptively damped GAMP algorithm. In lines (R1) and (R8), g_{z_m}' and g_{x_n}' denote the derivatives of g_{z_m} and g_{x_n} w.r.t their first arguments.

allows divergence while too much damping unnecessarily slows convergence, we are motivated to develop an *adaptive* damping scheme that applies just the right amount of damping at each iteration.

Table 1 details the proposed *adaptively damped GAMP* (AD-GAMP) algorithm. Lines (R3)-(R6) and (R10) use an iteration-t-dependent damping parameter $\beta(t) \in (0,1]$ to slow the updates,[1] and lines (R12)-(R18) adapt the parameter $\beta(t)$. When $\beta(t) = 1 \ \forall t$, AD-GAMP reduces to the original GAMP from [1]. Due to lack of space, we refer readers to [1,4] for further details on GAMP.

2.1. Damping Adaptation

The damping adaptation mechanism in AD-GAMP works as follows. Line (R12) computes the current cost $J(t+1)$, as described in the sequel. Line (R13) then checks evaluates whether the current iteration "passes" or "fails": it passes if the current cost is at least as good as the worst cost over the last $T_\beta \geq 0$ iterations or if $\beta(t)$

---

[1]The GAMPmatlab implementation [9] allows one to disable damping in (R6) and/or (R10).

is already at its minimum allowed value $\beta_{\min}$, else it fails. If the iteration passes, (R14)-(R15) implement a stopping condition, (R16) increases $\beta(t)$ by a factor $G_{\text{pass}} \geq 1$ (up to the maximum value $\beta_{\max}$), and (R17) increments the counter t. If the iteration fails, (R18) decreases $\beta(t)$ by a factor $G_{\text{fail}} < 1$ (down to the minimum value $\beta_{\min}$) and the counter t is *not* advanced, causing AD-GAMP to re-try the tth iteration with the new value of $\beta(t)$.

In the MAP case, line (R12) simply computes the cost $J(t+1) = J_{\text{MAP}}(\widehat{x}(t+1))$ for J_{MAP} from (2). The MMSE case, which is more involved, will be described next.

2.2. MMSE-GAMP Cost Evaluation

As proven in [4] and interpreted in the context of Bethe free entropy in [5], the fixed points of MMSE-GAMP are critical points of the optimization problem

$$(f_{\mathbf{x}}, f_{\mathbf{z}}) = \arg\min_{b_{\mathbf{x}}, b_{\mathbf{z}}} J_{\text{Bethe}}(b_{\mathbf{x}}, b_{\mathbf{z}}) \ \text{s.t. } \mathrm{E}\{\mathbf{z}|b_{\mathbf{z}}\} = \mathbf{A}\, \mathrm{E}\{\mathbf{x}|b_{\mathbf{x}}\} \quad (3)$$

$$J_{\text{Bethe}}(b_{\mathbf{x}}, b_{\mathbf{z}}) \triangleq D\big(b_{\mathbf{x}}\|p_{\mathbf{x}}\big) + D\big(b_{\mathbf{z}}\|p_{\mathbf{y}|\mathbf{z}}Z^{-1}\big) + H\big(b_{\mathbf{z}}, \boldsymbol{\nu}^p\big) \quad (4)$$

$$H\big(b_{\mathbf{z}}, \boldsymbol{\nu}^p\big) \triangleq \frac{1}{2}\sum_{m=1}^M \left(\frac{\text{var}\{\mathsf{z}_m|b_{\mathsf{z}_m}\}}{\nu_m^p} + \ln 2\pi\nu_m^p\right), \quad (5)$$

where $b_{\mathbf{x}}$ and $b_{\mathbf{z}}$ are separable pdfs, $Z^{-1} \triangleq \int p_{\mathbf{y}|\mathbf{z}}(\mathbf{y}|\mathbf{z})d\mathbf{z}$ is the scaling factor that renders $p_{\mathbf{y}|\mathbf{z}}(\mathbf{y}|\mathbf{z})Z^{-1}$ a valid pdf over $\mathbf{z} \in \mathbb{R}^M$, $D(\cdot\|\cdot)$ denotes Kullback-Leibler (KL) divergence, and $H(b_{\mathbf{z}})$ is an upper bound on the entropy of $b_{\mathbf{z}}$ that is tight when $b_{\mathbf{z}}$ is independent Gaussian with variances in $\boldsymbol{\nu}^p$. In other words, the pdfs $f_{\mathbf{x}}(\mathbf{x}; \widehat{\mathbf{r}}, \boldsymbol{\nu}^r) = \prod_{n=1}^N f_{\mathsf{x}_n}(x_n; \widehat{r}_n, \nu_n^r)$ and $f_{\mathbf{z}}(\mathbf{z}; \widehat{\mathbf{p}}, \boldsymbol{\nu}^p) = \prod_{m=1}^M f_{\mathsf{z}_m}(z_m; \widehat{p}_m, \nu_m^p)$ given in lines (D1) and (D2) of Table 1 are critical points of (3) for fixed-point versions of $\widehat{\mathbf{r}}, \boldsymbol{\nu}^r, \widehat{\mathbf{p}}, \boldsymbol{\nu}^p$.

Since $f_{\mathbf{x}}$ and $f_{\mathbf{z}}$ are functions of $\widehat{\mathbf{r}}, \boldsymbol{\nu}^r, \widehat{\mathbf{p}}, \boldsymbol{\nu}^p$, the cost J_{Bethe} can be written in terms of these quantities as well. For this, we first note

$$D\big(f_{\mathsf{x}_n}\|p_{\mathsf{x}_n}\big) = \int f_{\mathsf{x}_n}(x; \widehat{r}_n, \nu_n^r) \ln \frac{p_{\mathsf{x}_n}(x)\mathcal{N}(x;\widehat{r}_n,\nu_n^r)}{p_{\mathsf{x}_n}(x)C_n(\widehat{r}_n,\nu_n^r)}dx \quad (6)$$

$$= -\ln C_n(\widehat{r}_n, \nu_n^r) - \frac{\ln 2\pi\nu_n^r}{2} - \int f_{\mathsf{x}_n}(x; \widehat{r}_n, \nu_n^r)\frac{|x - \widehat{r}_n|^2}{2\nu_n^r}dx \quad (7)$$

$$= -\ln C_n(\widehat{r}_n, \nu_n^r) - \frac{\ln 2\pi\nu_n^r}{2} - \frac{|\widehat{x}_n - \widehat{r}_n|^2 + \nu_n^x}{2\nu_n^r}, \quad (8)$$

where $\widehat{x}_n$ and ν_n^x are the mean and variance of $f_{\mathsf{x}_n}(\cdot; \widehat{r}_n, \nu_n^r)$ from (R9) and (R8). Following a similar procedure,

$$D\big(f_{\mathsf{z}_m}\|p_{\mathsf{y}_m|\mathsf{z}_m}Z_m^{-1}\big)$$
$$= -\ln \frac{B_m(\widehat{p}_m, \nu_m^p)}{Z_m} - \frac{\ln 2\pi\nu_m^p}{2} - \frac{|\widehat{z}_m - \widehat{p}_m|^2 + \nu_m^z}{2\nu_m^p}, \quad (9)$$

where $\widehat{z}_m$ and ν_m^z are the mean and variance of $f_{\mathsf{z}_m}(\cdot; \widehat{p}_m, \nu_m^p)$ from (R2) and (R1). Then, since $D(f_{\mathbf{x}}\|p_{\mathbf{x}}) = \sum_{n=1}^N D\big(f_{\mathsf{x}_n}\|p_{\mathsf{x}_n}\big)$ and $D(f_{\mathbf{z}}\|p_{\mathbf{y}|\mathbf{z}}Z^{-1}) = \sum_{m=1}^M D\big(f_{\mathsf{z}_m}\|p_{\mathsf{y}_m|\mathsf{z}_m}Z_m^{-1}\big)$, (4) and (5) imply

$$J_{\text{Bethe}}(\widehat{\mathbf{r}}, \boldsymbol{\nu}^r, \widehat{\mathbf{p}}, \boldsymbol{\nu}^p) = -\sum_{m=1}^M \left(\ln B_m(\widehat{p}_m, \nu_m^p) + \frac{|\widehat{z}_m - \widehat{p}_m|^2}{2\nu_m^p}\right)$$
$$- \sum_{n=1}^N \left(\ln C_n(\widehat{r}_n, \nu_n^r) + \frac{\ln \nu_n^r}{2} + \frac{\nu_n^x + |\widehat{x}_n - \widehat{r}_n|^2}{2\nu_n^r}\right) + \text{const}, (10)$$

where we have written $J_{\text{Bethe}}(f_{\mathbf{x}}, f_{\mathbf{z}})$ as "$J_{\text{Bethe}}(\widehat{\mathbf{r}}, \boldsymbol{\nu}^r, \widehat{\mathbf{p}}, \boldsymbol{\nu}^p)$" to make the $(\widehat{\mathbf{r}}, \boldsymbol{\nu}^r, \widehat{\mathbf{p}}, \boldsymbol{\nu}^p)$-dependence clear, and where const collects terms invariant to $(\widehat{\mathbf{r}}, \boldsymbol{\nu}^r, \widehat{\mathbf{p}}, \boldsymbol{\nu}^p)$.

```
inputs:
    g_{z_m}, [A\widehat{x}]_m, \nu_m^p, \widetilde{p}_m(1), I_{max} \geq 1, \epsilon_{inv} \geq 0, \alpha \in (0,1], \phi \geq 0
for i = 1 : I_{max},
    e_m(i) = [A\widehat{x}]_m - g_{z_m}(\widetilde{p}_m(i), \nu_m^p)                    (F1)
    if |e_m(i)/g_{z_m}(\widetilde{p}_m(i), \nu_m^p)| < \epsilon_{inv}, stop          (F2)
    \nabla_m(i) = g'_{z_m}(\widetilde{p}_m(i), \nu_m^p)                             (F3)
    \widetilde{p}_m(i+1) = \widetilde{p}_m(i) + \alpha \frac{e_m(i)\nabla_m(i)}{\nabla_m^2(i)+\phi}    (F4)
end
outputs: \widetilde{p}_m(i)
```

Table 2. A regularized Newton's method to find the value of $\widetilde{p}_m$ that solves $[A\widehat{x}]_m = g_{z_m}(\widetilde{p}_m, \nu_m^p)$ for a given $[A\widehat{x}]_m$ and ν_m^p.

Note that the iteration-t MMSE-GAMP cost is not obtained simply by plugging $(\widehat{r}(t), \boldsymbol{\nu^r}(t), \widehat{p}(t+1), \boldsymbol{\nu^p}(t+1))$ into (10), because the latter quantities do not necessarily yield (f_x, f_z) satisfying the moment-matching constraint $\mathrm{E}\{z|f_z\} = A\,\mathrm{E}\{x|f_x\}$ from (3). Thus, it was suggested in [5] to compute the cost as

$$J_{MSE}(\widehat{r}(t), \boldsymbol{\nu^r}(t)) = J_{Bethe}(\widehat{r}(t), \boldsymbol{\nu^r}(t), \widetilde{p}, \boldsymbol{\nu^p}(t+1)), \quad (11)$$

for $\widetilde{p}$ chosen to match the moment-matching constraint, i.e., for

$$[A\widehat{x}(t+1)]_m = g_{z_m}(\widetilde{p}_m, \nu_m^p(t+1)) \text{ for } m = 1, \dots, M \quad (12)$$

where $\widehat{x}_n(t+1) = g_{x_n}(\widehat{r}_n(t), \mu_n^r(t))$ for $n = 1, \dots, N$ from (R9). Note that, since $\boldsymbol{\nu^p}(t+1)$ can be computed from $(\widehat{r}(t), \boldsymbol{\nu^r}(t))$ via (R8) and (R10), the left side of (11) uses only $(\widehat{r}(t), \boldsymbol{\nu^r}(t))$.

In the case of an additive white Gaussian noise (AWGN), i.e., $p_{y_m|z_m}(y_m|z_m) = \mathcal{N}(z_m; y_m, \nu^w)$ with $\nu^w > 0$, the function $g_{z_m}(\widetilde{p}_m, \nu_m^p)$ is linear in $\widetilde{p}_m$. In this case, [5] showed that (12) can be solved in closed-form, yielding the solution

$$\widetilde{p}_m = \big((\nu_m^p(t+1) + \nu^w)[A\widehat{x}(t+1)]_m - \nu_m^p(t+1)y_m\big)/\nu^w. \quad (13)$$

For general $p_{y_m|z_m}$, however, the function $g_{z_m}(\widetilde{p}_m, \nu_m^p)$ is nonlinear in $\widetilde{p}_m$ and difficult to invert in closed-form. Thus, we propose to solve (12) numerically using the regularized Newton's method detailed in Table 2. There, $\alpha \in (0, 1]$ is a stepsize, $\phi \geq 0$ is a regularization parameter that keeps the update's denominator positive, and I_{max} is a maximum number of iterations, all of which should be tuned in accordance with $p_{y_m|z_m}$. Meanwhile, $\widetilde{p}_m(1)$ is an initialization that can be set at $\widehat{p}_m(t+1)$ or $[A\widehat{x}(t+1)]_m$ and ϵ_{inv} is a stopping tolerance. Note that the functions g_{z_m} and g'_{z_m} employed in Table 2 are readily available from Table 1.

2.3. Mean Removal

To mitigate the difficulties caused by A with non-zero mean entries, we propose to rewrite the linear system "$z = Ax$" in (1) as

$$\underbrace{\begin{bmatrix} z \\ z_{M+1} \\ z_{M+2} \end{bmatrix}}_{\triangleq \overline{z}} = \underbrace{\begin{bmatrix} \widetilde{A} & b_{12}\gamma & b_{13}\mathbf{1}_M \\ b_{21}\mathbf{1}_N^H & -b_{21}b_{12} & 0 \\ b_{31}c^H & 0 & -b_{31}b_{13} \end{bmatrix}}_{\triangleq \overline{A}} \underbrace{\begin{bmatrix} x \\ x_{N+1} \\ x_{N+2} \end{bmatrix}}_{\triangleq \overline{x}} \quad (14)$$

where $(\cdot)^H$ is conjugate transpose, $\mathbf{1}_P \triangleq [1, \dots, 1]^H \in \mathbb{R}^P$, and

$$\mu \triangleq \tfrac{1}{MN}\mathbf{1}_M^H A\mathbf{1}_N \quad (15)$$

$$\gamma \triangleq \tfrac{1}{N}A\mathbf{1}_N \quad (16)$$

$$c^H \triangleq \tfrac{1}{M}\mathbf{1}_M^H(A - \mu\mathbf{1}_M\mathbf{1}_N^H) \quad (17)$$

$$\widetilde{A} \triangleq A - \gamma\mathbf{1}_N^H - \mathbf{1}_M c^H. \quad (18)$$

The advantage of (14) is that the rows and columns of $\overline{A}$ are approximately zero-mean. This can be seen by first verifying, via the definitions above, that $c^H\mathbf{1}_N = 0$, $\widetilde{A}\mathbf{1}_N = \mathbf{0}$, and $\mathbf{1}_M^H\widetilde{A} = \mathbf{0}^H$, which implies that the elements in every row and column of $\widetilde{A}$ are zero-mean. Thus, for large N and M, the elements in all but a vanishing fraction of the rows and columns in $\overline{A}$ will also be zero-mean. The mean-square coefficient size in the last two rows and columns of $\overline{A}$ can be made to match that in $\widetilde{A}$ via choice of $b_{12}, b_{13}, b_{21}, b_{31}$.

To understand the construction of (14), note that (18) implies

$$z = Ax = \widetilde{A}x + b_{12}\gamma\underbrace{\mathbf{1}_N^H x/b_{12}}_{\triangleq x_{N+1}} + b_{13}\mathbf{1}_M\underbrace{c^H x/b_{13}}_{\triangleq x_{N+2}}, \quad (19)$$

which explains the first M rows of (14). To satisfy the definitions in (19), we then require that $z_{M+1} = 0$ and $z_{M+2} = 0$ in (14), which can be ensured through the Dirac-delta likelihood

$$p_{y_m|z_m}(y_m|z_m) \triangleq \delta(z_m) \text{ for } m \in \{M+1, M+2\}. \quad (20)$$

Meanwhile, we make no assumption about the newly added elements x_{N+1} and x_{N+2}, and thus adopt the improper uniform prior

$$p_{x_n}(x_n) \propto 1 \text{ for } n \in \{N+1, N+2\}. \quad (21)$$

In summary, the mean-removal approach suggested here runs GAMP or AD-GAMP (as in Table 1) with $\overline{A}$ in place of A and with the likelihoods and priors augmented by (20) and (21). It is important to note that, if multiplication by A and A^H can be implemented using a fast transform (e.g., FFT), then multiplication by $\overline{A}$ and $\overline{A}^H$ can too; for details, see the GAMPmatlab implementation [9].

3. NUMERICAL RESULTS

We numerically studied the recovery $\mathsf{NMSE} \triangleq \|\widehat{x} - x\|^2/\|x\|^2$ of SwAMP [8] and the MMSE version of the original GAMP from [1] relative to the proposed mean-removed (M-GAMP) and adaptively damped (AD-GAMP) modifications, as well as their combination (MAD-GAMP). In all experiments, the signal x was drawn Bernoulli-Gaussian (BG) with sparsity rate τ and length $N = 1000$, and performance was averaged over 100 realizations. Average NMSE was clipped to 0 dB for plotting purposes. The matrix A was drawn in one of four ways:

(a) **Non-zero mean**: i.i.d $\mathsf{a}_{mn} \sim \mathcal{N}(\mu, \frac{1}{N})$ for a specified $\mu \neq 0$.

(b) **Low-rank product**: $\mathsf{A} = \frac{1}{N}\mathsf{UV}$ with $\mathsf{U} \in \mathbb{R}^{M \times R}$, $\mathsf{V} \in \mathbb{R}^{R \times N}$, and i.i.d $\mathsf{u}_{mr}, \mathsf{v}_{rn} \sim \mathcal{N}(0, 1)$, for a specified R. Note A is rank deficient when $R < \min\{M, N\}$.

(c) **Column-correlated**: the rows of A are independent zero-mean stationary Gauss-Markov processes with a specified correlation coefficient $\rho = \mathrm{E}\{\mathsf{a}_{mn}\mathsf{a}_{m,n+1}^H\}/\mathrm{E}\{|\mathsf{a}_{mn}|^2\}$.

(d) **Ill-conditioned**: $\mathsf{A} = \mathsf{U\Sigma V}^H$ where U and V^H are the left and right singular vector matrices of an i.i.d $\mathcal{N}(0, 1)$ matrix and $\mathbf{\Sigma}$ is a singular value matrix such that $[\mathbf{\Sigma}]_{i,i}/[\mathbf{\Sigma}]_{i+1,i+1} = (\kappa)^{1/\min\{M,N\}}$ for $i = 1, \dots, \min\{M, N\} - 1$, with a specified condition number $\kappa > 1$.

For all algorithms, we used $T_{max} = 1000$ and $\epsilon = 10^{-5}$. Unless otherwise noted, for adaptive damping, we used $T_\beta = 0$, $G_{pass} = 1.1$, $G_{fail} = 0.5$, $\beta_{max} = 1$, and $\beta_{min} = 0.01$. For SwAMP, we used the authors' publicly available code [12].

First we experiment with compressive sensing (CS) in AWGN at $\mathsf{SNR} \triangleq \mathrm{E}\{\|z\|^2\}/\mathrm{E}\{\|y - z\|^2\} = 60$ dB. For this, we used
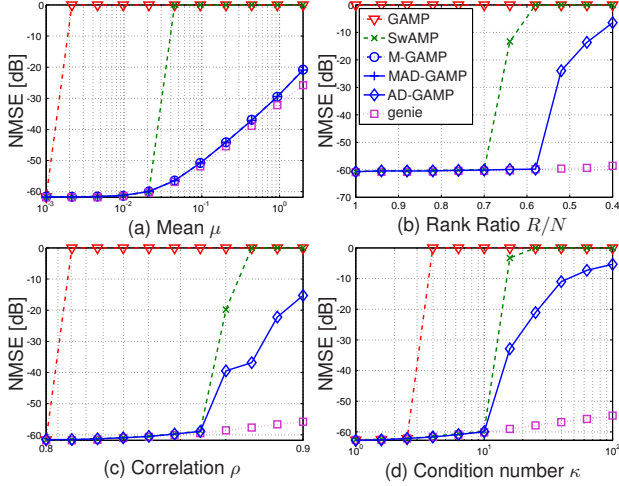
Fig. 1. AWGN compressive sensing under (a) non-zero-mean, (b) low-rank product, (c) column-correlated, and (d) ill-conditioned $\boldsymbol{A}$.
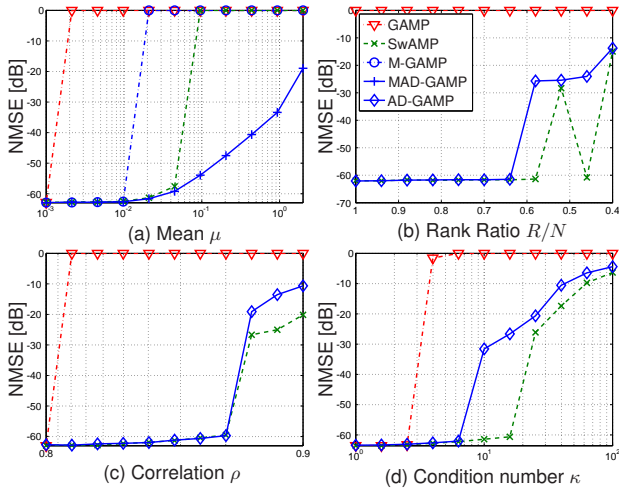
Fig. 2. "Robust" compressive sensing under (a) non-zero-mean, (b) low-rank product, (c) column-correlated, and (d) ill-conditioned $\boldsymbol{A}$.

$M = 500 = N/2$ measurements and sparsity rate $\tau = 0.2$. As a reference, we compute a lower-bound on the achievable NMSE using a genie who knows the support of $\boldsymbol{x}$. For non-zero-mean matrices, Fig. 1(a) shows that the proposed M-GAMP and MAD-GAMP provided near-genie performance for all tested means μ. In contrast, GAMP only worked with zero-mean $\boldsymbol{A}$ and SwAMP with small-mean $\boldsymbol{A}$. For low-rank product, correlated, and ill-conditioned matrices, Figs. 1(b)-(d) show that AD-GAMP is slightly more robust than SwAMP and significantly more robust than GAMP.

Next, we tried "robust" CS by repeating the previous experiment with sparsity rate $\tau = 0.15$ and with 10% of the observations (selected uniformly at random) replaced by "outliers" corrupted by AWGN at SNR $= 0$ dB. For (M)AD-GAMP, we set $\beta_{\max} = 0.1$ and $T_{\max} = 2000$. With non-zero-mean $\boldsymbol{A}$, Fig. 2(a) shows increasing performance as we move from GAMP to M-GAMP to SwAMP to MAD-GAMP. For low-rank product, correlated, and ill-conditioned matrices, Fig. 2(b)-(d) show that SwAMP was slightly more robust than AD-GAMP, and both where much more robust than GAMP.

Finally, we experimented with noiseless 1-bit CS [13], where $\boldsymbol{y} = \mathrm{sgn}(\boldsymbol{Ax})$, using $M = 3000$ measurements and sparsity ratio $\tau = 0.125$. In each realization, the empirical mean was subtracted
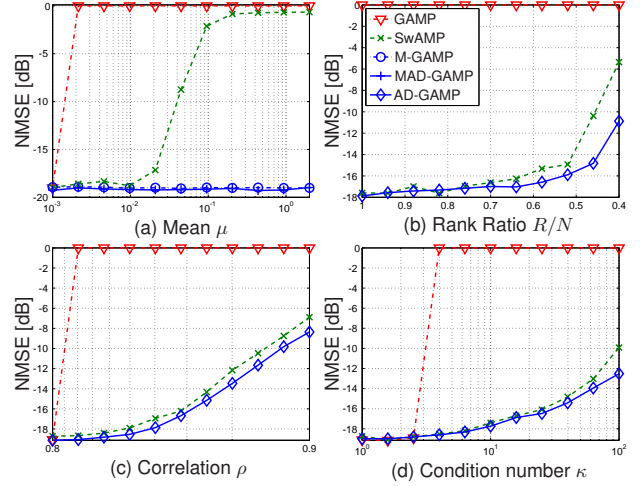

Fig. 3. 1-bit compressive sensing under (a) non-zero-mean, (b) low-rank product, (c) column-correlated, and (d) ill-conditioned $\boldsymbol{A}$.

		$\mu = 0.021$		$R/N = 0.64$		$\rho = 0.8$		$\log_{10} \kappa = 1$	
		MAD-GAMP	SwAMP	AD-GAMP	SwAMP	AD-GAMP	SwAMP	AD-GAMP	SwAMP
seconds	AWGN	**1.06**	1.90	**0.88**	2.74	**1.36**	3.84	**0.81**	1.49
	1-bit	**53.34**	83.21	**49.22**	137.46	**42.32**	149.40	**50.25**	117.62
	Robust	**3.47**	8.81	**2.66**	11.13	**3.33**	15.70	**2.38**	12.22
# iters	AWGN	42.9	**39.2**	130.0	**109.5**	221.9	**153.2**	121.4	**58.8**
	1-bit	947.8	**97.4**	942.7	**160.8**	866.2	**175.8**	927.3	**136.3**
	Robust	187.3	**42.2**	208.7	**56.1**	269.1	**79.2**	187.7	**61.7**

Table 3. Average runtime (in seconds) and # iterations of MAD-GAMP and SwAMP for various problem types and matrix types.

from the non-zero entries of $\boldsymbol{x}$ to prevent $y_m = 1 \,\forall m$. For (M)AD-GAMP, we used $\beta_{\max} = 0.5$. For SwAMP, we increased the stopping tolerance to $\epsilon = 5 \times 10^{-5}$, as it significantly improved runtime without degrading accuracy. For non-zero-mean $\boldsymbol{A}$, Fig. 3(a) shows that M-GAMP and MAD-GAMP were more robust than SwAMP, which was in turn much more robust than GAMP. For low-rank product, correlated, and ill-conditioned matrices, Figs. 3(b)-(d) show that MAD-GAMP and SwAMP gave similarly robust performance, while the original GAMP was very fragile.

Finally, we compare the convergence speed of MAD-GAMP to SwAMP. For each problem, we chose a setting that allowed MAD-GAMP and SwAMP to converge for each matrix type. Table 3 shows that, on the whole, MAD-GAMP ran several times faster than SwAMP but used more iterations. Thus, it may be possible to reduce SwAMP's runtime below that of MAD-GAMP using a more efficient (e.g., BLAS-based) implementation, at least for explicit $\boldsymbol{A}$. When $\boldsymbol{A}$ has a fast $O(N \log N)$ implementation (e.g., FFT), only (M)AD-GAMP will be able to exploit the reduced complexity.

4. CONCLUSIONS

We proposed adaptive damping and mean-removal modifications of GAMP that help prevent divergence in the case of "difficult" $\boldsymbol{A}$ matrices. We then numerically demonstrated that the resulting modifications significantly increase GAMP's robustness to non-zero-mean, low-rank product, column-correlated, and ill-conditioned $\boldsymbol{A}$ matrices. Moreover, they provide robustness similar to the recently proposed SwAMP algorithm, whilerunning faster than the current SwAMP implementation. For future work, we note that the sequential update of SwAMP could in principle be combined with the proposed mean-removal and/or adaptive damping to perhaps achieve a level robustness greater than either SwAMP or (M)AD-GAMP.

5. REFERENCES

[1] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," in *Proc. IEEE Int. Symp. Inform. Thy.*, Aug. 2011, pp. 2168–2172, (full version at *arXiv:1010.5141*).

[2] M. Bayati, M. Lelarge, and A. Montanari, "Universality in polytope phase transitions and iterative algorithms," in *Proc. IEEE Int. Symp. Inform. Thy.*, Boston, MA, June 2012, pp. 1643–1647, (full paper at *arXiv:1207.7321*).

[3] Adel Javanmard and Andrea Montanari, "State evolution for general approximate message passing algorithms, with applications to spatial coupling," *Inform. Inference*, vol. 2, no. 2, pp. 115–144, 2013.

[4] 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. Inform. Thy.*, July 2013, pp. 664–668, (full version at *arXiv:1301.6295*).

[5] F. Krzakala, A. Manoel, E. W. Tramel, and L. Zdeborová, "Variational free energies for compressed sensing," in *Proc. IEEE Int. Symp. Inform. Thy.*, July 2014, pp. 1499–1503, (see also *arXiv:1402.1384*).

[6] F. Caltagirone, F. Krzakala, and L. Zdeborová, "On convergence of approximate message passing," in *Proc. IEEE Int. Symp. Inform. Thy.*, July 2014, pp. 1812–1816, (see also *arXiv:1401.6384*).

[7] S. Rangan, P. Schniter, and A. Fletcher, "On the convergence of generalized approximate message passing with arbitrary matrices," in *Proc. IEEE Int. Symp. Inform. Thy.*, July 2014, pp. 236–240, (full version at *arXiv:1402.3210*).

[8] Andre Manoel, Florent Krzakala, Eric W. Tramel, and Lenka Zdeborová, "Sparse estimation with the swept approximated message-passing algorithm," *arXiv:1406.4311*, June 2014.

[9] S. Rangan, P. Schniter, J. T. Parker, J. Ziniel, J. Vila, M. Borgerding, and et al., "GAMPmatlab," https://sourceforge.net/projects/gampmatlab/.

[10] T. Heskes, "Stable fixed points of loopy belief propagation are minima of the Bethe free energy," in *Proc. Neural Inform. Process. Syst. Conf.*, Vancouver, B.C., Dec. 2002, pp. 343–350.

[11] P. Schniter and S. Rangan, "Compressive phase retrieval via generalized approximate message passing," in *Proc. Allerton Conf. Commun. Control Comput.*, Monticello, IL, Oct. 2012, pp. 815–822, (full version at *arXiv:1405.5618*).

[12] Andre Manoel, Florent Krzakala, Eric W. Tramel, and Lenka Zdeborová, "SwAMP demo user's manual," https://github.com/eric-tramel/SwAMP-Demo.

[13] U. S. Kamilov, V. K. Goyal, and S. Rangan, "Message-passing de-quantization with applications to compressed sensing," *IEEE Trans. Signal Process.*, vol. 60, no. 12, pp. 6270–6281, Dec. 2012.