
Swept Approximate Message Passing for Sparse Estimation

Andre Manoel

Institute of Physics, University of São Paulo, R. do Matão 187, São Paulo, SP 05508-090, Brazil

AMANOEL@IF.USP.BR

Florent Krzakala

Université Pierre et Marie Curie and École Normale Supérieure, 24 rue Lhomond, 75005 Paris, France

KRZAKALA@ENS.FR

Eric W. Tramel

École Normale Supérieure, 24 rue Lhomond, 75005 Paris, France

ERIC.TRAMEL@LPS.ENS.FR

Lenka Zdeborová

Institut de Physique Théorique, CEA Saclay, and CNRS URA 2306, 91191 Gif-sur-Yvette, France

LENKA.ZDEBOROVA@GMAIL.COM

Abstract

Approximate Message Passing (AMP) has been shown to be a superior method for inference problems, such as the recovery of signals from sets of noisy, lower-dimensionality measurements, both in terms of reconstruction accuracy and in computational efficiency. However, AMP suffers from serious convergence issues in contexts that do not exactly match its assumptions. We propose a new approach to stabilizing AMP in these contexts by applying AMP updates to individual coefficients rather than in parallel. Our results show that this change to the AMP iteration can provide expected, but hitherto unobtainable, performance for problems on which the standard AMP iteration diverges. Additionally, we find that the computational costs of this *swept* coefficient update scheme is not unduly burdensome, allowing it to be applied efficiently to signals of large dimensionality.

1. Introduction

Belief Propagation (BP) is a powerful iterative message passing algorithm for graphical models (Pearl, 1988; Mézard & Montanari, 2009; Opper & Saad, 2001). However, it presents two main drawbacks when applied to highly connected continuous variable problems: first, the need to work with continuous probability distributions; and second, the necessity to iterate over one such probability distribution for each pair of variables.

The first problem can be addressed by projecting the distributions onto a finite number of moments (Sudderth et al., 2010) and the second by utilizing the Thouless–Anderson–Palmer (TAP) approach (Mézard & Montanari, 2009; Opper & Saad, 2001) where only single variable marginals are required. Approximate message passing (AMP), first introduced in (Donoho et al., 2009), is one such relaxation of BP that utilizes both of the aforementioned approximations in order to solve sparse estimation problems. In AMP’s more general setting, as is considered in Generalized AMP (GAMP) (Rangan, 2011a), the goal of the algorithm is the reconstruction of an N -dimensional sparse vector \mathbf{x} given the knowledge of an M -dimensional vector \mathbf{y} obtained via a possibly non-linear and/or probabilistic output function $h(z)$ performed on a set of linear projections. Specifically,

$$y_\mu = h(z_\mu), \quad \text{where} \quad z_\mu = \sum_{i=1}^N \Phi_{\mu i} x_i. \quad (1)$$

For example, if $h(z) = z + \xi$ where ξ is a zero-mean *i.i.d.* Gaussian random variable, then $h(z)$ represents an additive white Gaussian noise (AWGN) channel. With this output function, in the setting $M \ll N$, (1) is simply the application of Compressed Sensing (CS) (Candès & Romberg, 2005) under noise. AMP is currently acknowledged as one of the foremost algorithms for such problems in terms of both its computational efficiency and in the number of measurements required for exact reconstruction of \mathbf{x} . In fact, with properly chosen measurement matrices (Krzakala et al., 2012b;a; Donoho et al., 2012), one

¹In the present work, we use subscript notation to denote the individual coefficients of vectors, i.e. y_μ refers to the μ^{th} coefficient of \mathbf{y} where $\mu \in \{1, 2, \dots, M\}$, and the double-subscript notation to refer to individual matrix elements in row-column order.

can achieve information-theoretically optimal reconstruction performance for CS, a hitherto unachievable bound with standard convex optimization approaches.

Just as with any iterative algorithm, the convergence properties of AMP are of chief analytical concern. Many rigorous results have been obtained on the performance of AMP in the case of *i.i.d.* and block *i.i.d.* matrices (Bayati & Montanari, 2011; Donoho et al., 2012). Unfortunately, while AMP performs well for zero-mean *i.i.d.* projections, performance tends to drastically decline if one moves away from these simple scenarios. In fact, even for *i.i.d.* matrices with a small positive mean, the algorithm may violently diverge, leading to poor reconstruction results (Caltagirone et al., 2014). This instability to slight variations from these strict assumptions on the projections is a serious problem for many practical applications of AMP.

The main theoretical reason for these convergence issues has been identified in (Caltagirone et al., 2014) using a rigorous tool called State Evolution. The issue stems from AMP’s use of a parallel update, instead of a sequential one, on the BP variables at each iteration. This result motivates the present work.

Three strategies have been proposed in recent literature to avoid this problem. First, one can highly damp the AMP iterations, as in (Krzakala et al., 2012b; Vila & Schniter, 2012). However, this often requires a damping factor so large that the cost, in terms of the number of iterations until convergence, is prohibitive. Additionally, it is not entirely clear how to determine an optimal damping factor to ensure convergence in general.

Second, one can modify the problem *a posteriori* in order to come back to a more favorable situation. For instance, one might remove the mean of the matrix and of the measurements (Caltagirone et al., 2014), or one might modify the algorithm according to the theoretical spectrum of the operator Φ (Rangan et al., 2014; Çakmak et al., 2014), if it is known. This knowledge about the operator may be prohibitive and could therefore present a strong limitation in practice. This procedure may work for non-zero mean projectors, but for more complicated operators it is not clear what one should do.

A third solution was proposed in (Caltagirone et al., 2014) where it was shown that one can take one step backward in approximation from AMP to a come back to a BP-style iteration (Caltagirone et al., 2014), using BP with a sequential update rather than the parallel one. It is indeed rather natural when using BP to work with a sequential update, sweeping through all variables, and this usually gives better convergence performance empirically. Going back to BP, however, amounts to a huge cost in terms of both memory and computational efficiency as there are $O(N^2)$ variables

to update per-iteration with BP as opposed to the $O(N)$ utilized in AMP.

In this contribution, we solve these problems by deriving a slightly modified and efficient AMP algorithm with greatly improved convergence properties while preserving the $O(N)$ iteration and memory cost of AMP. We accomplish this by a careful analysis of the relaxation leading from BP to AMP where we preserve the sequential, or swept, variable update pattern of BP in our AMP approach, and by paying great attention to the “time indices”. This leads to a subtly modified set of update rules for the AMP and GAMP algorithms without affecting the fixed point in any way. The resulting algorithm, which we denote as Swept AMP (SwAMP), possesses impressive empirical convergence properties.

2. Belief-Propagation for Signal Recovery

2.1. Signal Recovery as Statistical Estimation

To describe AMP, we focus on the CS signal recovery problem with real valued signals in terms of statistical inference. Given an unknown signal $\mathbf{x} \in \mathbb{R}^N$, a linear projection $\Phi \in \mathbb{R}^{M \times N}$, and a set of observations $\mathbf{y} \in \mathbb{R}^M$ generated from \mathbf{x} and Φ , we write the posterior distribution for the unknown signal according to Bayes’ rule,

$$P(\mathbf{x} | \Phi, \mathbf{y}) \propto P(\mathbf{y} | \Phi, \mathbf{x}) P_0(\mathbf{x}), \quad (2)$$

where we write \propto as we neglect the normalization constant. The likelihood $P(\mathbf{y} | \Phi, \mathbf{x})$ is determined according to the constraints one wishes to enforce, which we consider to be of form $\mathbf{y} = h(\Phi \mathbf{x})$, with h being, in general, any stochastic function. We consider h to be an AWGN channel ²,

$$y_\mu = h(\Phi_\mu \mathbf{x}) = \Phi_\mu \mathbf{x} + \mathcal{N}(0, \Delta), \quad (3)$$

where Δ is the variance of the AWGN and Φ_μ is the μ^{th} row-vector of Φ . Hence,

$$P(\mathbf{y} | \Phi, \mathbf{x}) = \prod_{\mu=1}^M \frac{1}{\sqrt{2\pi\Delta}} \exp \left[-\frac{(y_\mu - \sum_i \Phi_{\mu i} x_i)^2}{2\Delta} \right]. \quad (4)$$

The prior $P_0(\mathbf{x})$ is determined from the information we have on the structure of \mathbf{x} . For CS, we are concerned with the recovery of *sparse* signals, i.e. ones with few non-zero values. Unstructured sparse signals can be modeled well

²One can generalize h to be a more complicated output function. This generalization constitutes the change of AMP to GAMP (Rangan, 2011a). For example, we examine the case of 1-bit CS in Sec. 4.3 where h is a non-linear sign function.

by an *i.i.d.* Bernoulli sparse prior,

$$P_0(\mathbf{x}) \propto \prod_{i=1}^N P_0(x_i), \quad \text{where} \quad (5)$$

$$P_0(x_i) = \rho\psi(x_i) + (1 - \rho)\delta(x_i), \quad (6)$$

where $\psi(x_i)$ can be any distribution, e.g. $\psi(x_i) = \mathcal{N}(x_i; \bar{x}, \sigma^2)$, $\delta(x_i)$ is the Dirac delta function and the degree of sparsity is controlled by the value $\rho \in [0, 1]$. Notice that, in this usual setting, both distributions are factorized, that is, the likelihood is in M terms relative to the constraint over each y_μ , and the prior is in N terms relative to what is expected of each x_i . Factorized distributions such as these are well represented by graphical models (Wainwright & Jordan, 2008), specifically, bipartite graphs in which the $M + N$ factors are represented by one type of node and the N variables x_i by another. Once the posterior distribution is written down, the estimate $\hat{\mathbf{x}}$ may be assigned in different ways, according to what loss function one wishes to minimize. In this work, we are chiefly concerned with the minimum mean-squared error (MMSE) estimate, which can be shown to be the average of x_i with respect to the posterior $P(\mathbf{x} | \Phi, \mathbf{y})$; if one were able to compute the posterior's marginals, the MMSE estimate would read

$$\hat{x}_i^{\text{MMSE}} = \int dx_i x_i P(x_i | \Phi, \mathbf{y}), \quad \forall i. \quad (7)$$

The strategy employed by AMP is to infer the marginals of the posterior by using a relaxed version of the BP algorithm (Pearl, 1988; Mézard & Montanari, 2009), and thus to arrive at the MMSE estimate of the unknown signal \mathbf{x} .

2.2. Relaxed Belief-Propagation

BP implements a message-passing scheme between nodes in a graphical model, ultimately allowing one to compute approximations of the posterior marginals. Messages $m_{i \rightarrow \mu}$ are sent from the variable nodes to the factor nodes and subsequent messages $m_{\mu \rightarrow i}$ are sent from factor nodes back to variable nodes that corresponds to algorithm's current "beliefs" about the probabilistic distribution of the variables x_i . Since these distributions are continuous, the first relaxation step is to move to a projected version of these distributions, as described in (Rangan, 2011a; Krzakala et al., 2012a). Here, we shall follow the notation of reference and use the following parametrization

$$a_{i \rightarrow \mu} \triangleq \int dx_i x_i m_{i \rightarrow \mu}(x_i), \quad (8)$$

$$v_{i \rightarrow \mu} \triangleq \int dx_i x_i^2 m_{i \rightarrow \mu}(x_i) - a_{i \rightarrow \mu}^2, \quad (9)$$

$$m_{\mu \rightarrow i}(x_i) \propto e^{-\frac{x_i^2}{2} A_{\mu \rightarrow i} + B_{\mu \rightarrow i} x_i}. \quad (10)$$

This leads (see (Krzakala et al., 2012a)) to the following closed recursion sometimes called relaxed BP (r-BP),

$$A_{\mu \rightarrow i} = \frac{\Phi_{\mu i}^2}{\Delta + \sum_{j \neq i} \Phi_{\mu j}^2 v_{j \rightarrow \mu}}, \quad (11)$$

$$B_{\mu \rightarrow i} = \frac{\Phi_{\mu i}(y_\mu - \sum_{j \neq i} \Phi_{\mu j} a_{j \rightarrow \mu})}{\Delta + \sum_{j \neq i} \Phi_{\mu j}^2 v_{j \rightarrow \mu}}, \quad (12)$$

$$a_{i \rightarrow \mu} = f_1 \left(\frac{1}{\sum_{\nu \neq \mu} A_{\nu \rightarrow i}}, \frac{\sum_{\nu \neq \mu} B_{\nu \rightarrow i}}{\sum_{\nu \neq \mu} A_{\nu \rightarrow i}} \right), \quad (13)$$

$$v_{i \rightarrow \mu} = f_2 \left(\frac{1}{\sum_{\nu \neq \mu} A_{\nu \rightarrow i}}, \frac{\sum_{\nu \neq \mu} B_{\nu \rightarrow i}}{\sum_{\nu \neq \mu} A_{\nu \rightarrow i}} \right), \quad (14)$$

where the functions f are defined by the following prior-dependent integrals

$$f_1(\Sigma^2, R) \triangleq \int dx x P_0(x) \frac{1}{\sqrt{2\pi\Sigma}} e^{-\frac{(x-R)^2}{2\Sigma^2}}, \quad (15)$$

and

$$\begin{aligned} f_2(\Sigma^2, R) &\triangleq \int dx x^2 P_0(x) \frac{1}{\sqrt{2\pi\Sigma}} e^{-\frac{(x-R)^2}{2\Sigma^2}} - f_1^2(\Sigma^2, R) \\ &= \Sigma^2 \frac{df_1}{dR}(\Sigma^2, R). \end{aligned} \quad (16)$$

After convergence, the single point marginals are given by

$$a_i = f_1 \left(\frac{1}{\sum_{\nu} A_{\nu \rightarrow i}}, \frac{\sum_{\nu} B_{\nu \rightarrow i}}{\sum_{\nu} A_{\nu \rightarrow i}} \right), \quad (17)$$

$$v_i = f_2 \left(\frac{1}{\sum_{\nu} A_{\nu \rightarrow i}}, \frac{\sum_{\nu} B_{\nu \rightarrow i}}{\sum_{\nu} A_{\nu \rightarrow i}} \right). \quad (18)$$

We intentionally write r-BP without specifying time indices since the updates can be performed in one of two ways. The first approach is to update in parallel, where *all* variables are updated at time t given the state at time $t - 1$. The second is the random sequential update where one picks a *single* index i and updates all messages corresponding to it. A time-step is completed once all indices have been visited and updated once.

2.3. The AMP algorithm

We consider the AMP algorithm in the form that was derived in (Donoho et al., 2010; Rangan, 2011b; Krzakala et al., 2012a). The main steps are a) going from belief propagation (BP) to a relaxed BP (r-BP) where only the two first moments of all messages are kept and b) using N sites marginals instead of $N \times M$ messages and adding the compensating Onsager terms (Thouless et al., 1977). Fi-

nally, AMP reads

$$V_\mu^{t+1} = \sum_i \Phi_{\mu i}^2 v_i^t, \quad (19)$$

$$\omega_\mu^{t+1} = \sum_i \Phi_{\mu i} a_i^t - \frac{(y_\mu - \omega_\mu^t)}{\Delta + V_\mu^t} \sum_i \Phi_{\mu i}^2 v_i^t, \quad (20)$$

$$(\Sigma_i^{t+1})^2 = \left[\sum_\mu \frac{\Phi_{\mu i}^2}{\Delta + V_\mu^{t+1}} \right]^{-1}, \quad (21)$$

$$R_i^{t+1} = a_i^t + \frac{\sum_\mu \Phi_{\mu i} \frac{(y_\mu - \omega_\mu^{t+1})}{\Delta + V_\mu^{t+1}}}{\sum_\mu \frac{\Phi_{\mu i}^2}{\Delta + V_\mu^{t+1}}}, \quad (22)$$

$$a_i^{t+1} = f_1((\Sigma_i^{t+1})^2, R_i^{t+1}), \quad (23)$$

$$v_i^{t+1} = f_2((\Sigma_i^{t+1})^2, R_i^{t+1}). \quad (24)$$

where $f_k(\Sigma^2, R)$, here and in what follows, are the k -th connected cumulants w.r.t. the probability measure

$$\mathcal{Q}(x) = \frac{1}{Z(\Sigma^2, R)} P(x) \frac{e^{-\frac{(x-R)^2}{2\Sigma^2}}}{\sqrt{2\pi\Sigma^2}}, \quad (25)$$

with $Z(\Sigma^2, R)$ as the normalization constant.

The variables a_i and v_i are the AMP estimators for the mean and variance of the component i of the signal. The quality of the reconstruction can be evaluated by computing the mean squared error (MSE) and the average variance

$$E^t = \frac{1}{N} \sum_{i=1}^N (s_i - a_i^t)^2, \quad V^t = \frac{1}{N} \sum_{i=1}^N v_i^t. \quad (26)$$

where s_i is the original signal component. When $\gamma = 0$, the performance of the AMP algorithm was analyzed rigorously in the limit of large system size via the state evolution $(E^{t+1}, V^{t+1}) = G(E^t, V^t)$, where G is a function specified in (Bayati & Montanari, 2011; Donoho et al., 2010; Rangan, 2011b; Krzakala et al., 2012a). An important property of the Bayes optimal inference (i.e. when the signal was indeed generated from the assumed prior distribution) is that the two parameters are equal in the large size limit, $E^t = V^t$, and the state evolution hence reduces to an iterative equation of a single real number, which is amenable to rigorous analysis (Bayati & Montanari, 2011). In statistical physics $E^t = V^t$ is called the Nishimori condition and is discussed in the context of compressed sensing in detail in (Krzakala et al., 2012a). In general, when $\gamma = 0$ we observed by analyzing the state evolution equations that even when at initial times $E^{t=0} \neq V^{t=0}$ the equality $E^t = V^t$ is restored after a sufficient number of iterations.

3. Swept Approximate Message Passing

3.1. Rederiving the Time Indices

In the message-passing described in the previous section, $2(M \times N)$ messages are sent, one between each variable component and each measurement at each iteration. This creates a very large computational and memory burden for applications with large N, M . It is possible to rewrite the BP equations in terms of only $N + M$ messages by making the assumption that Φ is dense and that its elements are of magnitude $O(1/\sqrt{N})$. In statistical physics, this assumption leads to the TAP equations (Thouless et al., 1977) used in the study of spin glasses. For graphical models, such strategies have been discussed in (Oppor & Saad, 2001). The use of TAP with r-BP provides the standard AMP iteration. Now let us investigate the expansion of the factor ω_μ from Eq. (20) of the AMP iteration as we include the time, or iteration, indices t . First one has

$$\begin{aligned} a_{i \rightarrow \mu}^{t+1} &= f_1 \left(\frac{1}{\sum_\nu A_{\nu \rightarrow i}^t - A_{\mu \rightarrow i}^t}, \frac{\sum_\nu B_{\nu \rightarrow i}^t - B_{\mu \rightarrow i}^t}{\sum_\nu A_{\nu \rightarrow i}^t - A_{\mu \rightarrow i}^t} \right) \\ &= a_i^{t+1} - B_{\mu \rightarrow i}^t (\Sigma_i^t)^t \frac{\partial f_1}{\partial R} ((\Sigma_i^t)^t, R_i^t), \\ &= a_i^{t+1} - B_{\mu \rightarrow i}^t v_i^{t+1}, \end{aligned} \quad (27)$$

making the expansion for ω_μ

$$\begin{aligned} \omega_\mu^{t+1} &= \sum_i \Phi_{\mu i} a_i^{t+1} - \frac{(y_\mu - \omega_\mu^t)}{\Delta_\mu + V_\mu^t} \sum_i \Phi_{\mu i}^2 v_i^{t+1} \\ &= \sum_i \Phi_{\mu i} a_i^{t+1} - \frac{(y_\mu - \omega_\mu^t)}{\Delta_\mu + V_\mu^t} V_\mu^{t+1}, \end{aligned} \quad (28)$$

which allows us to close the equations on the set of a, v, R, Σ, V and ω . Iterating all relations in parallel (i.e. updating all R, Σ 's, then a, v 's and then the ω, V 's) provides the AMP iteration.

The implementation of the sequential update is not a straightforward task as many otherwise intuitive attempts lead to non-convergent algorithms. The key observation in the derivation of SwAMP is that (28) mixes different time indices: while the “ a ” and “ V ” are the “new ones”, the expression in the fraction is the “old” one, i.e. the one *before* the last iteration. The implication of this is that while $\sum_i \Phi_{\mu i} a_i$ and V_μ should be recalculated as the updates sweep over i at a single time-step, the term $(y_\mu - \omega_\mu)/(\Delta_\mu + V_\mu)$ (which we denote as g_μ later on) should not. A corresponding bookkeeping then leads to the SwAMP algorithm for the evolution of $\omega_\mu, \Sigma_i^2, V_\mu$ and R_i described in Alg. 1. At this point, the difference between AMP and SwAMP appears minimal, but, as we shall see, the differences in convergence properties turn out to be spectacular.

Algorithm 1 Swept AMP

Input: $y, \Phi, \Delta, \theta_{\text{prior}}, t_{\text{max}}, \varepsilon$
 $t \leftarrow 0$
 Initialize $\{\mathbf{a}^{(0)}, \mathbf{v}^{(0)}\}, \{\omega^{(0;N+1)}, \mathbf{V}^{(0;N+1)}\}$
while $t < t_{\text{max}}$ **and** $\|\mathbf{a}^{(t+1)} - \mathbf{a}^{(t)}\| > \varepsilon$ **do**
 for $\mu = 1$ **to** M **do**
 $g_{\mu}^{(t)} \leftarrow \frac{y_{\mu} - \omega_{\mu}^{(t;N+1)}}{\Delta + V_{\mu}^{(t;N+1)}}$
 $V_{\mu}^{(t+1;1)} \leftarrow \sum_i \Phi_{\mu i}^2 v_i^{(t)}$
 $\omega_{\mu}^{(t+1;1)} \leftarrow \sum_i \Phi_{\mu i} a_i^{(t)} - V_{\mu}^{(t+1;1)} g_{\mu}^{(t)}$
 end for
 $\mathbf{S} \leftarrow \text{Permute}([1, 2, \dots, N])$
 for $k = 1$ **to** N **do**
 $i \leftarrow S_k$
 $\Sigma_i^{2(t+1)} \leftarrow \left[\sum_{\mu} \frac{\Phi_{\mu i}^2}{\Delta + V_{\mu}^{(t+1;k)}} \right]^{-1}$
 $R_i^{(t+1)} \leftarrow a_i^{(t)} + \Sigma_i^{2(t+1)} \sum_{\mu} \Phi_{\mu i} \frac{y_{\mu} - \omega_{\mu}^{(t+1;k)}}{\Delta + V_{\mu}^{(t+1;k)}}$
 $a_i^{(t+1)} \leftarrow f_1(R_i^{(t+1)}, \Sigma_i^{2(t+1)}; \theta_{\text{prior}})$
 $v_i^{(t+1)} \leftarrow f_2(R_i^{(t+1)}, \Sigma_i^{2(t+1)}; \theta_{\text{prior}})$
 for $\mu = 1, m$ **do**
 $V_{\mu}^{(t+1;k+1)} \leftarrow V_{\mu}^{(t+1;k)} + \Phi_{\mu i}^2 (v_i^{(t+1)} - v_i^{(t)})$
 $\omega_{\mu}^{(t+1;k+1)} \leftarrow \omega_{\mu}^{(t+1;k)} + \Phi_{\mu i} (a_i^{(t+1)} - a_i^{(t)}) - g_{\mu}^{(t)} (V_{\mu}^{(t+1;k+1)} - V_{\mu}^{(t+1;k)})$
 end for
 end for
 $t \leftarrow t + 1$
end while

3.2. Generalized Swept AMP

We note that this procedure can also be generalized, a la GAMP, for output channels other than the AWGN. The required change is minimal (Rangan, 2011a): one should replace the term $(y_{\mu} - \omega_{\mu})/(\Delta_{\mu} + V_{\mu})$ in the R_i and ω_{μ} updates with $g_{\text{out}}(\omega_{\mu}, V_{\mu})$, a generic function which depends on the channel. Specifically, $g_{\text{out}}(\omega, V) = \int dz P(y|z) e^{-\frac{(z-\omega)^2}{2V}} \left(\frac{z-\omega}{V}\right)$. Additionally, the $\frac{1}{\Delta_{\mu} + V_{\mu}}$ term in the Σ_i^2 update should be replaced by $-\frac{\partial g_{\text{out}}}{\partial \omega}$. Notice that all AWGN specific terms are recovered for $P(y|z) \propto e^{-\frac{(y-z)^2}{2\Delta}}$.

4. Numerical Results

Here, we present a range of numerical results demonstrating the effectiveness of the SwAMP algorithm for problems on which both standard AMP and ℓ_1 minimization via convex optimization fail to provide desirable reconstruction performance. All experiments were conducted on a computer with an i7-3930K processor and run via Matlab. We have provided demonstrations of the SwAMP code on-line

³. For calculating ℓ_1 recoveries, we utilize an implementation of the SPGL1 (van den Berg & Friedlander, 2008) algorithm.

4.1. Compressed Sensing with Troublesome Projections

As discussed earlier, using projections of non-zero mean to sample \mathbf{x} is one of the simplest cases for which AMP can fail to converge. However, by using the proposed SwAMP approach, accurate estimates of \mathbf{x} can be obtained even when the mean of the projections is non-negligible. While it may be possible to use mean subtraction, our proposed approach does not require such preprocessing. Additionally, as we will show later, not all problems are amenable to such mean subtraction. To evaluate the effectiveness of SwAMP as compared to the standard parallel-update AMP iteration, we draw *i.i.d.* projections according to

$$\Phi_{\mu i} \sim \mathcal{N}\left(\frac{\gamma}{N}, \frac{1}{N}\right), \quad (29)$$

where the magnitude of the projector mean is controlled by the term γ . For a given signal \mathbf{x} and noise variance Δ , as γ increases from 0, we expect to see AMP failing to converge. This behavior can be observed in the numerical experiments presented in Fig. 1. Here, we observe that SwAMP is robust to values of γ over an order of magnitude larger than the standard AMP, converging to a low-MSE solution even for $\gamma \approx 140$ while AMP fails already at $\gamma = 2$. Additionally, for the tested parameters, ℓ_1 minimization fails to provide a meaningful reconstruction for any value of γ .

We also consider an even more troublesome case, namely, strongly correlated projections. Such problems are of interest as they arise naturally in machine learning and biomedical applications, where the practitioner does not design the projections for optimal signal recovery, but rather the projections represent sampled observations from which the practitioner desires to predict some response variables by finding an interpretable, i.e. sparse, set of regression coefficients. Such observations can be highly correlated and thus represent a significant impediment to the use of parallel AMP for regression tasks.

For these tests, we draw

$$\Phi = \frac{1}{N} P Q, \quad \text{where } P_{\mu k}, Q_{ki} \sim \mathcal{N}(0, 1) \quad (30)$$

with $P \in \mathbb{R}^{M \times R}$, $Q \in \mathbb{R}^{R \times N}$ and $R \triangleq \eta N$. That is, Φ is *rank-deficient* for $\eta < \alpha$, where $\alpha = \frac{M}{N}$. In our experiments, we use η to denote the level of independence of the rows of Φ , with lower values of η representing a more difficult problem. We observe that the elements of

³<https://github.com/eric-tramel/SwAMP-Demo>

Φ are neither normal nor *i.i.d.* for these experiments. In Fig. 2 we see that SwAMP is robust to even these extremely troublesome projections while AMP fails to converge.

As we cannot guarantee that SwAMP provides Bayes optimal estimation for problems exhibiting strong correlations in Φ , we provide a comparison to other well-known approaches in Fig. 2. Here, we compare the recovery performance of SwAMP to adaptive Lasso (Zou, 2006), basis pursuit denoising (BPDN) (Chen et al., 2001), as well as the ℓ_p generalization of Lasso (Tibshirani, 1996) for the non-convex case $p < 1$ (Chartrand, 2007). In order to address the free-parameter of the regularization strength in both ℓ_p and adaptive Lasso, we compute the entire solution path and report only the most accurate solution, providing best-case oracle results for these approaches. Additionally, for adaptive Lasso we choose a weight exponent of 0.1. The tests are conducted over 500 independent realizations of the sparse reconstruction problem for $N = 1024$, $\alpha = 0.6$, and $\rho = 0.2$ with a noise variance $\Delta = 10^{-8}$. We assume that the value of Δ is known to all the tested approaches *a posteriori*. For the implementation of the ℓ_p regression we utilized the SparseReg Matlab toolbox (Zhou, 2013), while we use the SpaSM Matlab toolbox (Sjöstrand et al., 2012) for the implementation of adaptive Lasso.

We can see that SwAMP provides more accurate estimates across all η while also possessing a more robust transition between successful and unsuccessful reconstructions in terms of η . Over the tested η , SwAMP averaged 2 seconds of computation time, BPDN averaged 8 seconds, and ℓ_p and adaptive Lasso required on average 40 and 14.88 seconds of compute time respectively. We point out these run times to demonstrate that despite the potential complexity of the SwAMP fixed-point iteration, its computational burden is on par with, or an improvement on, other well known techniques.

These two experiments demonstrate how the proposed SwAMP iteration allows for AMP-like performance while remaining robust to conditions outside of the TAP assumptions about the projector.

4.2. Group Testing

Group testing, also known as *pooling* in molecular biology, is an approach to designing experiments so as to reduce the number of tests required to identify rare events or faulty items. In the most naive approach to this problem, the number of tests is equal to the number of items, as each item is tested individually. However, since only a small fraction of the items may be faulty, the number of tests can be significantly reduced via pooling, i.e. testing many items simultaneously and allowing items to be included within multiple different tests. The nature of this linear combination of tests allows for a CS-type approach




Figure 1. AMP, SwAMP, and ℓ_1 solvers compared for CS signal reconstruction for sensing matrices with positive mean on sparse signals of size $N = 10^4$ and sparsity $\rho = 0.2$ with noise variance $\Delta = 10^{-8}$. The projections for have been created following (29) using $M = \alpha N$ measurements with $\alpha = 0.5$. Finally, a comparison between the reconstruction error obtained by SwAMP and ℓ_1 -minimization is given at the bottom for the same experimental settings.

to faulty item detection, but with a few important caveats. First, the operator is extremely sparse since the number of pools, and the number of items in them, may be limited due to physical testing constraints. Second, the elements of this operator are commonly 0/1. Group testing is therefore a very challenging application for AMP since the properties of the group testing operator do not match AMP’s assumptions.

In one recent work (Zhang et al., 2013), the authors use both BP and AMP for group testing and found that while basic AMP would not converge, very good results—optimal ones, in fact—could be obtained by using a BP approach. This came at a large computational cost, however. Here, we have repeated the experiment of (Zhang et al., 2013) using the SwAMP approach instead of AMP and BP. In fact, for SwAMP, a sparse operator is a very advantageous situation in terms of computational efficiency. Since the projector is extremely sparse by construction, we may explicitly ignore operations involving null elements, thus considerably improving the algorithm’s speed, as seen in Fig. 3(b). Here, we also see that SwAMP’s computational complexity is on the order of $O(N^2)$, as is AMP’s. Group testing experiments are shown in Fig. 3(a) where we use random 0/1 projections, under the constraint that each projection should sum to 7, to sample sparse 0/1 signals with $K \ll N$ non-zero elements, where N is the signal dimensionality. While AMP diverges when attempting to

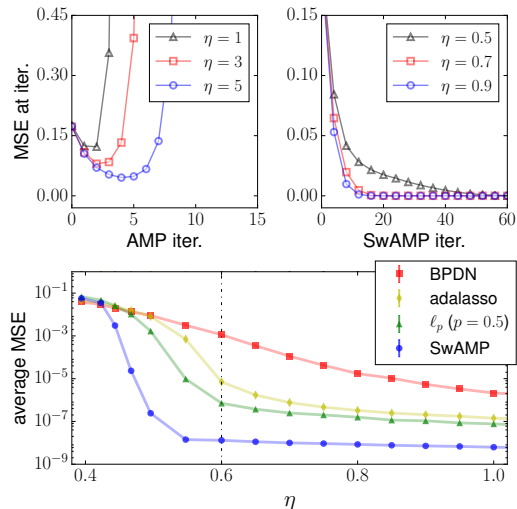


Figure 2. At the top, convergence behavior of AMP and SwAMP are compared for CS signal reconstruction for correlated sensing matrices on sparse signals of size $N = 10^4$ and sparsity $\rho = 0.2$ with noise variance $\Delta = 10^{-8}$. The projectors have been created according to (30) and are rank-deficient for $\eta < \alpha = 0.6$. At the bottom, a comparison between log-scale average reconstruction MSE obtained by SwAMP, BPDN, adaptive Lasso, and ℓ_p regularization is given for signals of size $N = 1024$ for $\Delta = 10^{-8}$, $\rho = 0.2$, and $\alpha = 0.6$.

recover these signals, SwAMP converges to the correct solution in few iterations. Additionally, SwAMP very closely matches the BP transition, thus providing recovery performance better than convex optimization, just as BP does, but with much less computational complexity.

4.3. 1-bit Compressed Sensing

One of the confounding factors regarding the practical implementation of CS in hardware devices is the treatment of measurement quantization. The original CS analysis provides recovery bounds based upon the assumption of real-valued measurements. However, in practice, hardware devices cannot capture such values with infinite precision, and so some kind of quantization on the measurements must be implemented. Specifically, if $Q(\cdot)$ is a uniform scalar quantizer, then $\mathbf{y} = Q(\Phi \mathbf{x}, B)$, where B is the number of bits used to represent the measurement. If signal recoverability is significantly impacted by small B , then the dimensionality reduction provided by CS may be lost by the requirement for many bits to encode each measurement.

Thankfully, recent works have shown CS recovery to be robust to quantization and the non-linear error it introduces. In fact, CS has been shown (Boufounos & Baraniuk, 2008; Jacques et al., 2012) to be robust even in the extreme case $B = 1$ known as 1-bit CS. In this case, the quantized mea-




Figure 3. (a) Group testing phase transition diagram between successful and unsuccessful signal recovery over M , the number of pools, and K , the number of non-zero signal elements. Successful recovery means the correct identification of all signal elements. The top-left of the diagram represents the easiest problems while the bottom-right the most difficult. The transition lines are drawn along the contour of 50% of recoveries succeeding for many trials. (b) Execution times for both SwAMP and AMP using a sparse matrix with 25% of its elements having non-zero value. The reported times are measured for 500 iterations of the algorithms for each value of N for the parameters $\rho = 0.25$ and $\alpha = 0.75$.

asurements are given by

$$\mathbf{y} = \text{sign}(\Phi \mathbf{x}). \quad (31)$$

The non-linearity and severity of 1-bit CS requires special treatment from the CS recovery procedure. In (Boufounos & Baraniuk, 2008), a renormalized fixed-point continuation (RFPC) algorithm was proposed. Later, (Jacques et al., 2012) analyzed the sensitivity of 1-bit CS to sign flips and proposed a noise-robust recovery algorithm, binary iterative hard thresholding (BIHT).

Recognizing the capability of GAMP to handle non-linear output channels, (Kamilov et al., 2012) proposed the use of GAMP for signal recovery from quantized CS measurements. Further analysis of message-passing approaches to the 1-bit CS problem from the perspective of statistical me-

chanics was given in (Xu & Kabashima, 2013) where a modified fixed-point iteration was derived via the cavity method which provided both improved recovery accuracy and reconstruction time as compared to the RFPC. Additionally, the authors used replica analysis to estimate the optimal MSE performance of ℓ_1 -minimization based 1-bit CS reconstruction. Finally, this analysis is extended in (Xu et al., 2014) to include the theoretical Bayesian optimal performance, which we will use as a baseline of comparison in Fig. 4(a).

Both methods (Kamilov et al., 2012) and (Xu & Kabashima, 2013) show the effectiveness of algorithms grounded in statistical mechanics for quantized CS reconstruction. However, both assume an amenable set of projections. Even projections possessing small mean can cause large degradations in performance. While mean removal is occasionally effective in the usual CS setting, it cannot be used for 1-bit CS due to the nature of the sign operation in (31). An algorithm that can handle troublesome projectors can therefore be of great use. In Sec. 2.3, we show how the SwAMP can be modified to the general-channel setting, as was done in GAMP. This generalization allows for 1-bit CS recovery with SwAMP under much more relaxed requirements for Φ .

In Fig. 4(a), we see Generalized SwAMP (G-SwAMP) results for $\Phi_{\mu_i} \sim \mathcal{N}(\frac{20}{N}, \frac{1}{N})$. We observe that G-SwAMP performs admirably even for this non-negligible mean on the projectors. In terms of recovery performance, it does not quite meet the theoretical Bayes optimal performance (Xu et al., 2014), however, this is expected as the Bayes optimal performance is calculated for $\gamma = 0$. Additionally, we see that even for this non-zero mean, G-SwAMP outperforms both the BIHT's empirical performance for the same mean, as well as the best-case theoretical ℓ_1 performance for zero mean (Xu & Kabashima, 2013). Finally, in Fig. 4(b), we see that GAMP fails to provide any meaningful signal recovery for γ small, while G-SwAMP continues to converge to low-MSE even for large values of γ .

5. Conclusion

While the AMP algorithm has been shown to be a very desirable approach for signal recovery and statistical inference problems in terms of both computational efficiency and accuracy, it is also very sensitive to problems which deviate from its fundamental assumptions. In this work, we propose the SwAMP algorithm which matches AMP's accuracy while remaining robust to such variations, all without unduly increasing computation or memory requirements. We also demonstrate how SwAMP can be used to solve practical problems for which AMP and GAMP cannot be applied, namely, group testing and 1-bit CS with




Figure 4. Results for 1-bit CS. (a) Top: Comparison between the Bayes optimal MSE for zero-mean projectors (Xu et al., 2014) (dashed lines) and that obtained by SwAMP for projectors with $\gamma = 20$ (markers) for three different levels of signal sparsity. The reported empirical results were obtained by averaging over 200 instances of size $N = 512$. Bottom: Comparison of SwAMP and BIHT for $\rho = 1/8$ for experiment conditions identical to the figure above; theoretical results for zero-mean projectors are also presented for completeness, including theoretical ℓ_1 performance (Xu & Kabashima, 2013). (b) Single instance comparison between GAMP and G-SwAMP for 1-bit CS with $N = 2048$, $\rho = 1/8$, and $\alpha = 3$.

troublesome projections. In all cases, SwAMP provides superior accuracy as compared to ℓ_1 -minimization, as well as convergence properties superior to AMP and GAMP, and all with less computational and memory burden than BP or r-BP.

Exact analysis of the asymptotic state evolution of SwAMP, as well as a thorough analytical proof of its convergence, remains a challenging open problem for future work. However, the results are promising for many difficult applications such as sparse logistic regression for identifying significant biological markers from among many thousands of potential features when very few test samples are available. Such applications would be impossible for GAMP due to the sensitivity of its convergence, but are a quite natural application for G-SwAMP.

References

- Bayati, M. and Montanari, A. The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Transactions on Information Theory*, 57(2):764, 2011.
- Boufounos, P. T. and Baraniuk, R. G. 1-bit compressive sensing. In *Proceedings of the 42nd Annual Conference on Information Sciences and Systems*, pp. 16–21, Princeton, NJ, 2008.
- Caltagirone, F., Krzakala, F., and Zdeborová, L. On convergence of approximate message passing. In *Information Theory Proceedings (ISIT), 2014 IEEE International Symposium on*, 2014. URL <http://arxiv.org/abs/1401.6384>.
- Candès, E. J. and Romberg, J. Signal recovery from random projections. In *Computational Imaging III*, pp. 76–86, San Jose, CA, 2005. Proc. SPIE 5674.
- Çakmak, B., Winther, O., and Fleury, B. H. S-amp: Approximate message passing for general matrix ensembles. *arXiv preprint 1405.2767*, 2014.
- Chartrand, Rick. Exact reconstruction of sparse signals via nonconvex minimization. *IEEE Signal Processing Letters*, 14(10):707–710, 2007.
- Chen, Scott Shaobing, Donoho, David L., and Saunders, Michael A. Atomic decomposition by basis pursuit. *SIAM Review*, 43(1):129–159, March 2001.
- Donoho, D. L., Maleki, A., and Montanari, A. Message-passing algorithms for compressed sensing. *Proc. National Academy of Sciences of the United States of America*, 106(45):18914, 2009.
- Donoho, D. L., Maleki, A., and Montanari, A. Message passing algorithms for compressed sensing: I. motivation and construction. In *Proc. IEEE Information Theory Workshop*, pp. 1–5, 2010.
- Donoho, D. L., Javanmard, A., and Montanari, A. Information-theoretically optimal compressed sensing via spatial coupling and approximate message passing. In *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*, pp. 1231. IEEE, 2012.
- Jacques, L., Laska, J. N., Boufounos, P. T., and Baraniuk, R. G. Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *arXiv preprint 1104.3160v3*, 2012.
- Kamilov, U. S., Goyal, V. K., and Rangan, S. Message-passing de-quantization with applications to compressed sensing. *IEEE Transactions on Image Processing*, 60(12):6270, 2012.
- Krzakala, F., Mézard, M., Sausset, F., Sun, Y. F., and Zdeborová, L. Probabilistic reconstruction in compressed sensing: Algorithms, phase diagrams, and threshold achieving matrices. *J. Stat. Mech.: Th. and Exp.*, (8):P08009, 2012a.
- Krzakala, F., Mézard, M., Sausset, F., Sun, Y. F., and Zdeborová, L. Statistical-physics-based reconstruction in compressed sensing. *Physical Review X*, 2(2):021005, 2012b.
- Mézard, M. and Montanari, A. *Information, Physics, and Computation*. OUP, 2009. ISBN 9780198570837.
- Opper, M. and Saad, D. *Advanced Mean Field Methods: Theory and Practice*. MIT Press, 2001. ISBN 0585417725. NIPS workshop series.
- Pearl, J. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, 1988. ISBN 9781558604797.
- Rangan, S. Generalized approximate message passing for estimation with random linear mixing. In *Information Theory Proceedings, IEEE International Symposium on*, pp. 2168, 2011a.
- Rangan, S. Generalized approximate message passing for estimation with random linear mixing. In *Proc. IEEE International Symposium on Information Theory*, pp. 2168–2172, 2011b.
- Rangan, S., Schniter, P., and Fletcher, A. K. On the convergence of approximate message passing with arbitrary matrices. *arXiv preprint 1402.3210*, 2014. URL <http://arxiv.org/abs/1402.3210>.
- Sjöstrand, Karl, Clemmensen, Line Harder, Larsen, Rasmus, and Ersbøll, Bjarne. Spasm: A matlab toolbox for sparse statistical modeling. *Journal of Statistical Software*, pp. 1–24, 2012.
- Sudderth, E. B., Ihler, A. T., Isard, M., Freeman, W. T., and Willsky, A. S. Nonparametric belief propagation. *Communications of the ACM*, 53(10):95, 2010.
- Thouless, D. J., Anderson, P. W., and Palmer, R. G. Solution of ‘solvable model of a spin glass’. *Philosophical Magazine*, 35(3):593, 1977.
- Tibshirani, Robert. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58(1):267–288, 1996.
- van den Berg, E. and Friedlander, M. P. Probing the pareto frontier for basis pursuit solutions. *SIAM Journal on Scientific Computing*, 31(2):890–912, 2008. doi: 10.1137/080714488. URL <http://link.aip.org/link/?SCE/31/890>.

- Vila, J. P. and Schniter, P. Expectation-maximization gaussian-mixture approximate message passing. In *Proc. 46th Annual Conference on Information Sciences and Systems*, pp. 1, 2012.
- Wainwright, M. J. and Jordan, M. I. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1, 2008. ISSN 1935-8237. doi: 10.1561/2200000001.
- Xu, Y. and Kabashima, Y. Statistical mechanics approach to 1-bit compressed sensing. *Journal of Statistical Mechanics: Theory and Experiment*, (2):P02041, 2013.
- Xu, Y., Kabashima, Y., and Zdeborová, L. Bayesian signal reconstruction for 1-bit compressed sensing. arXiv to appear, 2014.
- Zhang, P., Krzakala, F., Mézard, M., and Zdeborová, L. Non-adaptive pooling strategies for detection of rare faulty items. In *Communications Workshops, Proc. IEEE International Conference on*, pp. 1409, Budapest, Hungary, 2013.
- Zhou, Hua. Matlab SparseReg toolbox version 0.0.1. Available online: <http://hua-zhou.github.io/software/sparsereg>, July 2013.
- Zou, Hui. The adaptive lasso and its oracle properties. *Journal of the American Statistical Association*, 101 (476):1418–1429, 2006.