Massachusetts Institute of Technology

The Overlap Gap Property and Approximate Message Passing Algorithms for p-spin models

David Gamarnik[*]and Aukosh Jagannath[†]

November 27, 2019

Abstract

We consider the algorithmic problem of finding a near ground state (near optimal solution) of a p-spin model. We show that for a class of algorithms broadly defined as Approximate Message Passing (AMP), the presence of the Overlap Gap Property (OGP), appropriately defined, is a barrier. We conjecture that when $p \geq 4$ the model does indeed exhibits OGP (and prove it for the space of binary solutions). Assuming the validity of this conjecture, as an implication, the AMP fails to find near ground states in these models, per our result. We extend our result to the problem of finding pure states by means of Thouless, Anderson and Palmer (TAP) based iterations, which is yet another example of AMP type algorithms. We show that such iterations fail to find pure states approximately, subject to the conjecture that the space of pure states exhibits the OGP, appropriately stated, when $p \geq 4$.

1   Introduction

Given an N-tensor $A = (A_{i_1,\dots,i_p}, 1 \leq i_1, \dots, i_p \leq N) \in (\mathbb{R}^N)^{\otimes p}$ of order p and an N-vector $u \in \mathbb{R}^N$, define the usual inner tensor product by

$$A(u) \triangleq \sum_{1 \leq i_1, \dots, i_p \leq N} A_{i_1,\dots,i_p} u_{i_1} \cdots u_{i_p}. \tag{1}$$

Consider the associated normalized variational problem over the binary cube $\mathcal{B}_N \triangleq \{-1, 1\}^N$:

$$\eta_N \triangleq \frac{1}{N} \min_{\sigma \in \mathcal{B}_N} A(\sigma). \tag{2}$$

The case when A consists of i.i.d. zero mean Gaussian random entries with variance $1/N^{p-1}$ that is $\mathcal{N}\left(0, \frac{1}{N^{p-1}}\right)$ corresponds to the problem of finding a ground state of a p-spin model with Gaussian couplings and the (unique) vector u^* achieving the minimization value is called the ground state [Pan13]. The choice of variance $1/N^{p-1}$ and the normalization $1/N$ is dictated by the associated Gibbs distribution defined by assigning probability weight proportional to $\exp(-\beta A(\sigma))$ to each $\sigma \in \mathcal{B}_N$ for some fixed inverse temperature parameter $\beta \in \mathbb{R}_+$. In this case the partition function

$$Z \triangleq \sum_{\sigma \in \mathcal{B}_N} \exp(-\beta A(\sigma))$$

---

[*]Operations research Center and Sloan School of Management, MIT. Email: gamarnik@mit.edu

[†]Department of Statistics and Actuarial Sciences and Department of Applied Mathematics, University of Waterloo. Email:a.jagannath@uwaterloo.ca

is well approximated by $N\eta_N$ as β increases and η_N is known to converge to a strictly negative limiting value $\eta^* < 0$ with high probability (w.h.p.) as $N \to \infty$. For us though the details of the choice of scaling are immaterial and the variational problem above is equivalent to the case when A consists of i.i.d. standard normal entries and the normalization $1/N$ is skipped. Another standard assumption in the literature is to assume a symmetry of A, for example assuming that entries are fully determined by i.i.d. entries corresponding to $i_1 \le \cdots \le i_p$. This difference is again immaterial. Indeed, consider the tensor $\bar{A}$ defined by

$$\bar{A}_{i_1,\ldots,i_p} = \frac{1}{p!} A^\pi,$$

where A^π is defined by

$$A^\pi_{i_1,\ldots,i_p} = A_{\pi(i_1,\ldots,i_p)}, \qquad 1 \le i_1,\ldots,i_p \le N,$$

for any permutation π of $1,\ldots,p$. Note that $\bar{A}$ is symmetric and satisfies $\bar{A}(u) = A_\pi(u)$ for every π.

In the present paper we focus on the algorithmic question of solving the minimization problem (2) approximately and efficiently (in polynomial time). That is, the question is one of existence of a polynomial time algorithm which for every $\epsilon > 0$ produces a sequence of solution $\sigma_N \in \mathcal{B}_N$ satisfying

$$A(\sigma_N)/N \le (1 - \epsilon)\eta^*,$$

as $N \to \infty$, ideally w.h.p. as $N \to \infty$. This problem was successfully solved recently by Montanari [Mon18] in the case of the Sherrington-Kirkpatrick model, which is the special case corresponding to $p = 2$. The result though assumes the validity of a (widely believed) conjecture that the overlap distribution function is strictly increasing. In particular, it assumes the absence of an interval $[\nu_1, \nu_2]$ inside the support of the overlap distribution with zero mass, namely that the overlap distribution does not exhibit the Overlap Gap Property (OGP). The algorithm is based on a variant of Approximate Message Passing (AMP) type algorithms, which in the context of spin glasses is well motivated by the so-called Thouless, Anderson and Palmer (TAP) equation describing the magnetization of spins in spin glass models. AMP, as a class of algorithms was also found to be one of the most effective classes of algorithm in many models of signal processing [Kab03, DMM09, Bol14, BM11, JM13, BLM+15, BMN17], specifically models involving a "planted signal" (which does not apply to our p-spin model). The algorithmic result of [Mon18] in its order was inspired by a similar result by Subag [Sub18] regarding the problem of finding a near ground states in a spherical "mixed p-spin" model. Here one considers a linear combination of objectives of the form (1) as one varies p, with the coefficients being fixed, and optimizing over the unit sphere $\{u : \|u\|_2 \le 1\}$ instead of $\mathcal{B}_N$. Here a polynomial time constrictions of near optimal solutions is provided, under the assumption that the model does not exhibit OGP (see part (2) of Proposition 1 in [Sub18]). For the case of spherical models, the necessary and sufficient conditions for the OGP are known [CS17, JT17a, Tal06]. Both p-spin and spherical p-spin models are related to the Random Energy Model (REM) considered from the algorithmic perspective by Addario-Berry and Maillard [ABM18], where in contrast to [Mon18] and [Sub18] algorithmic hardness is established away from the ground state value. One should note, however that REM is an oracle based optimization problem and thus does not fit classical input size based algorithmic complexity questions arising in the context of p-spin and spherical p-spin models.

At the same time it is known that the OGP does take place in p-spin models when $p \ge 4$, as was established in [CGP+19, Theorem 3]. In particular, it was shown that for every even $p \ge 4$, there exists $\mu > 0, 0 < \nu_1 < \nu_2 < 1$ such that w.h.p. for every pair of solutions σ_1, σ_2 satisfying $A(\sigma_j)/N \le \eta^* + \mu$ for $j = 1, 2$, the associated normalized overlap, defined simply as the normalized absolute value of the inner product $(1/N)|\langle \sigma_1, \sigma_2 \rangle|$ is either at most ν_1 or at least some ν_2. This naturally raises the question as to

whether the OGP creates a barrier to the success of AMP when $p \geq 4$. The main result of our paper, Theorem 3.3, is to establish precisely this fact under the assumption that a certain relaxed version of the OGP takes place when $p \geq 4$. The relaxed version concerns the optimization problem min $A(u)$ when u is relaxed to be in Hilbert cube $u \in [-1,1]^N \triangleq \mathcal{H}_N$, and otherwise is defined in the same way as for $\mathcal{B}_N$. This relaxed version would be a rather straightforward implication of the OGP for binary solutions if one could show that every nearly optimal solution in $\mathcal{H}_N$ is nearly binary. Unfortunately, even this fact is not known, and we leave it as an interesting, though, as we believe, an approachable open problem. As a consequence of our main result, we show that extension of the AMP result of [Mon18] to the case $p \geq 4$ is not possible. As another implication, we show that a natural iterative scheme of computing the fixed point of the TAP equations fails as well in the case $p \geq 4$. We note, that this iterative scheme is known to succeed in the high temperature regime due to the result of [Bol14]. Another important class of algorithms ruled out by our negative result is gradient descent type algorithms. Since the gradient of $A(u)$ is a linear combination of the vectors of the form $A(\cdot, u)$ defined below in (3), then a discrete implementation of the gradient descent algorithm in the form $u^t = u^{t-1} + \eta_{t-1} \nabla A(u^{t-1})$ for some step choices η_t is also a special case of our class of AMP algorithms.

One challenge in establishing our result formally is the formalization of the class of AMP algorithms to begin with. Unfortunately, there is no one formal definition for it, but rather there is a vaguely proposed scheme for a class of iterations inspired by the Belief Propagation type algorithms. The iterations take the form $u^{t+1} = F_t(G_t(u^t), G_{t-1}(u^{t-1}), \ldots, G_0(u^0)), t \geq 0$ and are performed for some constant number of rounds $t = 0, 1, \ldots, T$, where F^t is an in general t-dependent function involving vector $A(\cdot, u) \in \mathbb{R}^N$ defined by

$$A(\cdot, u) \triangleq \left(\sum_{i_2, \ldots, i_p} A_{i, i_2, \ldots, i_p} u_{i_2} \cdots u_{i_p}, 1 \leq i \leq N \right), \tag{3}$$

for any $u \in \mathbb{R}^N$, as well other non-linear operators $G_t : \mathbb{R}^N \to \mathbb{R}^N$ defined typically through some kind of univariate or t variant non-linear functions $g_t : \mathbb{R}^t \to \mathbb{R}$ applied coordinate-wise in some way. We note that (3) is simply a matrix vector product Au when $p = 2$.

Thus, as a first step, we introduce a precise class of such iterative algorithms (functions) F^t, and associate it with a precise set of assumptions. We show separately that the algorithm of [Mon18] is a special case. We assume that the results $U^t, 0 \leq t \leq T$ of each iterations are truncated so that the resulting vector always belongs to $\| \cdot \|_\infty$ bounded region of the form $[-M, M]^N$ for some constant M. The rational for the truncation is as follows. In the implementation of the AMP the iterations $F^t, 0 \leq t \leq T$ produce a real valued vector $U^T \in \mathbb{R}^N$ which is then projected to a vector in $\mathcal{B}_N$ in a way discussed below. The idea here is that U^T is a vector which is "close enough" to some vector $\sigma \in \mathcal{B}_N$ which is a near-ground state. In particular, the typical entries of U^t are "not too far" from interval $[-1, 1]$, and in particular are bounded by M. We restrict every vector U^t to be in $[-M, M]^N$ for technical convenience. The rounding scheme $[-M, M]^N \to \mathcal{B}_N$ assumed to be adopted by our class of AMP is similar the one that was used in [Mon18]: first U^T is projected to vector $V \in \mathcal{H}_N$ via a natural truncation $x \to \min(\max(x, -1), 1)$, and then some rounding scheme $\Pi : \mathcal{H}_N \to \mathcal{B}_N$ is adopted by the algorithm designer, which is guaranteed asymptotically to never lower the quality of the solution, that is, it guarantees that $A(\Pi(V))/N \leq A(V)/N + o(1)$. Our main result, stated more precisely, says that for *any* AMP algorithm thus defined, the vector V is w.h.p. sub-optimal, namely $A(V)/N$ exceeds η^* by some fixed constant $\mu > 0$, w.h.p. as $N \to \infty$. Thus we establish that the vector V obtained in the pen-ultimate (before Π) step of the AMP is sub-optimal. We note that it is precisely this vector which is shown to be nearly optimal in the case $p = 2$ in the argument of [Mon18]. The last step of converting a real vector $V \in \mathcal{H}_N$ to $\Pi(V) \in \mathcal{B}_N$ is just used there in order to obtain a genuinely binary vector. We do not establish that the ultimate vector $\Pi(V) \in \mathcal{B}_N$ is sub-optimal, and this is a limitation of our

technique. We note however that showing near optimality of $\pi(V)$ without showing near optimality of V would amount to believing that the rounding Π is somehow mysteriously capable of producing near optimal binary solution $\Pi(V)$ from a presumably far from optimal fractional solution $V \in \mathcal{H}_N$, which is something which does not seem to be plausible, and something which not established in [Mon18]. Nevertheless, it would be admittedly a more complete result to show directly that $\Pi(V)$ is far from optimal, without assuming the same for V, but we are currently unable to make this argument rigorous, and leave it for further investigation.

Proof of the main result. Outline

We now describe the main ingredients of our proof. First, as a consequence of a result established in [CGP$^+$19], we show that the OGP holds w.h.p. not just for one instance of A, but for a continuous family of sets of the form $\mathcal{A} = \left(\sqrt{1-\tau}A + \sqrt{\tau}\hat{A} \right), \tau \in [0,1]$ where $\hat{A}$ is an independent instance of A. Note that for each fixed τ, the corresponding tensor has the same distribution as A. An easy consequence of the OGP result in [CGP$^+$19] and the chaos result of [CHL18], is the fact, which we prove in this paper (Theorem 3.4), that the OGP holds for $\mathcal{A}$ as well in the sense that for any two $A_{\tau_1}, A_{\tau_2} \in \mathcal{A}$ and any $\sigma_1, \sigma_2 \in \mathcal{B}_N$ satisfying $A_{\tau_j}(\sigma_j)/N \leq \eta^* + \mu$, it is again the case that $N^{-1}|\langle \sigma_1, \sigma_2 \rangle| \in [0, \nu_1] \cup [\nu_2, 1]$, for the same values ν_1, ν_2, μ. The chaos property roughly speaking says for any fixed $\tau_1 \neq \tau_2$ near optimal solutions of A_{τ_1} and A_{τ_2} are nearly orthogonal, see Theorem 3.5 below. Our main conjecture regarding OGP (Conjecture 3.2), which we use as an assumption of the main result, is the conjecture that OGP holds in fact for near optimal solutions in $\mathcal{H}_N$ as opposed to those in $\mathcal{B}_N$, for the same family of instances. Establishing this conjecture is an interesting open question.

Our main ingredient of the proof is then to show that the iterations $U^T = U^T(A)$ as function of A are sufficiently "continuous" to perturbation of the entries of A. Specifically, we obtain an upper bound on $N^{-1}\|U^T(A_\tau) - U^T(A_0)\|_2$ for the interpolation scheme $\mathcal{A}$ which is sufficiently continuous in τ. This result is the subject of Theorem 5.1. A straightforward implication is that the same bound applies to $N^{-1}\|V(A_\tau) - V(A_0)\|_2$, where, as we recall $V(A)$ is projection of $U^T(A)$ through the truncation $x \to \min(\max(x, -1), 1)$. Separately, we use the independence of A and $\hat{A}$ to argue the near orthogonality of $V(A)$ and $V(\hat{A})$. The continuity result above then is used to show that for an appropriate choice of τ, it will hold that $N^{-1}\langle V(A_\tau), V(A) \rangle \in (\nu_1, \nu_2)$. The (conjectured) OGP property implies that this choice of τ corresponds to a sufficiently sub-optimal solution $V(A_\tau)$, which contradicts concentration property of $A(V(A_\tau))$, which we establish separately using standard techniques, including Gaussian concentration of measure and Kirszbraun's Theorem.

Prior results on OGP and algorithmic implications

The concept of OGP originates in the study of spin glass models, specifically the study of overlap distribution of replicas generated according to some associated Gibbs distribution, such as the one described above. Understanding the limiting distribution of overlaps is of an utmost importance to spin glass theory and has recieved significant attention [AC15, ACZ17, JT17b, JT18]. The first connections between to study of the overlaps and algorithms were made in the context of random constraint satisfaction problems, such as random K-SAT problem and many other similar problems. These problems exhibit an "infamous" gap between the range of parameters for which a satisfying assignment exists vs those for which solutions can be found in polynomial time. The apparent hardness was linked conjecturally to the clustering (shattering) property of these models which were discovered to appear roughly in the regime where known polynomial time algorithms fail [ACO08, ACORT11, MMZ05, MPZ02, COE11]. The clustering property says, roughly speaking, that a large part of the set of satisfying assignments can be partitioned into clusters separated by Hamming distance, which is of the order of the size of the model itself. It is notable that the proof technique used to establish such a clustering property

actually shows something more: the overlaps between pairs of typical (random in some appropriate sense) satisfying assignments lie in a disconnected union of intervals $[0, \nu_1] \cup [\nu_2, 1]$. Thus the set of solutions is disconnected not only with respect to its ambient metric space, but also with respect to its one-dimensional projection onto the set of possible overlap values. The proof technique relies on fairly standard application of the moment method. The disconnectivity of overlaps (that is the presence of the overlap gaps of the form (ν_1, ν_2)) was later used as an obstruction to a class of local algorithms defined as so-called Factors of IID in [GS17a, RV$^+$17, GS17b, CGP$^+$19], and for random walk type algorithms in [COHH17]. It is this line of work which the closest in spirit to the present one, as one can think of AMP as a natural definition for "local" algorithms defined on dense instances – instances not defined on sparse graphs and hypergraphs.

It is important to note that while OGP implies the clustering property, the converse in general is not true. Indeed if the OGP takes place then one can partition the set of all solutions of interest into those which have overlap at least ν_2 with some arbitrarily marked solution σ^o (thus marked "Cluster 1"), vs solutions with overlap at most ν_1 with σ^o (thus marked "other clusters"), leading to a set of at least two clusters separated by a significant distance. On the other hand, one can easily create a subset of $\mathcal{B}_N$ for which the set of all overlaps spans the entire interval $[0, 1]$, though at the same time admits clustering partition.

The OGP was further established for some other models, some involving planted signals [GL18, DI17, GJS19]. It was shown in [GJS19] to be an obstruction to Glauber Dynamics type algorithms by showing that OGP implies the existence of a free energy well, a property which was shown to be a barrier for Markov chain type algorithms in problems involving planted signals [BAGJ18a]. A related notion of free energy barriers associated with these gaps were also shown to be obstructions for local Markov chain type algorithms for problems of the class considered here in [BAJ18], where it was also shown that these free energy barriers occur in a broad class of models including both the p-spin and spherical p-spin models. It can be shown that OGP implies the existence of a free energy barrier at sufficiently low temperatures. It is of interest to establish the broadest class of algorithms for which OGP is a provable barrier.

The remainder of the paper is structured as follows. In the next section we introduce the formalism of the AMP algorithms. In Section 3 we give the definition of the OGP, state the corresponding conjecture and state our main result. The validity of OGP for binary solutions is proven in the same section. Some preliminary technical results are established in Section 4. Our main technical result is Theorem 5.1 which is stated and proven in Section 5. We note that it is a purely deterministic result showing that the output of the AMP depends on the values of the tensor A sufficiently continuously. In Section 6 we establish the concentration property of the solution V produced by the AMP around its expectation. Our main theorem is proven in Section 7. In Section 8 we consider TAP solutions and show that a natural class of iterations suggested by TAP fails to find the fixed point of TAP, modulo the same Conjecture 3.2, since the iterations are a special case of the class of AMP algorithms we define. This result is a direct implication of our main result, Theorem 3.3. It contrasts with the positive result of Bolthausen [Bol14], which establishes that these iterations do converge to the solution of TAP equations in the high-temperature setting. In Section 9 we verify that the AMP algorithm constructed in [Mon18] also fit the general definition of AMP introduced in this paper. Finally, we conclude in Section 10 where we state some open questions.

2 Approximate Message Passing iterations formalism

$\langle x, y \rangle = \sum_{1 \le i \le N} x_i y_i$ denotes inner product of vectors $x, y \in R^N$. For any tensor $B \in (\mathbb{R}^N)^{\otimes p}$ $\|B\|_2$ denotes the Frobenius norm $\sqrt{\sum_{1 \le i_1, \ldots, i_p \le N} B_{i_1, \ldots, i_p}^2}$, and $\|B\|_{\mathrm{op}}$ denotes the operator norm

$$\max_{u_1, \ldots, u_p} B(u_1, \ldots, u_p)$$

where the maximum is over all $u_1, \ldots, u_p \in \mathbb{R}^N, \|u_j\|_2 = 1, 1 \le j \le p$. By Cauchy-Schwartz inequality $\|B\|_{\mathrm{op}} \le \|B\|_2$.

Throughout the paper $A \in (\mathbb{R}^N)^{\otimes p}$ denotes N-size order p tensor consisting of $\mathcal{N}(0, N^{-(p-1)})$ i.i.d. entries. For any $u_1, \ldots, u_{p-1} \in \mathbb{R}^N$ let

$$A(u_1, \ldots, u_p) = \sum_{1 \le i_1, \ldots, i_p \le N} A_{i_1, i_2, \ldots, i_{p-1}} u_{i_1}^1 \cdots u_{i_p}^p,$$

so that for any $u \in \mathbb{R}^N$, $A(u) = A(u, \cdots, u)$ as in (1). Here $u_r = (u_1^r, \ldots, u_N^r)$. For any $u_1, \ldots, u_{p-1} \in \mathbb{R}^N$ we also introduce

$$y = A(\cdot, u_1, \ldots, u_{p-1}) \in \mathbb{R}^N \tag{4}$$

defined by

$$y_i = \sum_{1 \le i_1, \ldots, i_{p-1} \le N} A_{i, i_1, \ldots, i_{p-1}} u_{i_1}^1 \cdots u_{i_{p-1}}^{p-1} \qquad 1 \le i \le N.$$

Similarly, for any $u \in \mathbb{R}^N$ we write $A(\cdot, u)$ instead of $A(\cdot, u, u, \ldots, u)$ for short. We recall the definition of η_N from (2). Observe that we may view $A(u)$ as a centered Gaussian process indexed by $\mathcal{H}_N$, which has covariance

$$\mathbb{E}[A(u)A(v)] = N \Big(\frac{\langle u, v \rangle}{N} \Big)^p.$$

In particular, $|\mathbb{E}[A(u)A(v)]| \le N$ for any $u, v \in \mathbb{R}^N$ with $\|u\|_2, \|v\|_2 \le 1$. The following concentration result is then an immediate consequence of the Borell-TIS inequality, Theorem 2.1.11 of [AT09].

Theorem 2.1. *For every $\delta > 0$*

$$\mathbb{P}\left(|\eta_N - \mathbb{E}[\eta_N]| \ge \delta \right) \le \exp(-(1/4)\delta^2 N),$$

for all sufficiently large N.

A major consequence of the development in spin glass theory is the existence of the limit

$$\lim_{N \to \infty} \mathbb{E}[\eta_N] = \eta^* < 0, \tag{5}$$

which by Theorem 2.1 also implies that the limit $\eta_N \to \eta^*$ holds w.h.p. as $N \to \infty$.

We now introduce a set of assumptions which are used to define a class of AMP algorithms. Fix a positive integer T and an $M > 0$. Consider two sequences of functions $f_t : [-M, M]^t \to \mathbb{R}$ and $F_t : \mathbb{R} \times [-M, M]^t \to \mathbb{R}, 1 \le t \le T$.

Assumption 2.2. $f_t(0) = 0$. *Furthermore, functions f_t, F_t are Lipschitz continuous on their respective domains. More precisely, there exists $\zeta \in \mathbb{R}_+$ such that for all $1 \le t \le T$,*

$$\sup_{u, v \in [-M, M]^t} |f_t(u) - f_t(v)| \le \zeta \|u - v\|_2, \tag{6}$$

$$\sup_{u, v \in \mathbb{R} \times [-M, M]^t} |F_t(u) - F_t(v)| \le \zeta \|u - v\|_2. \tag{7}$$

The assumption (7) says that the function F_t is Lipschitz on an infinite rectangle $\mathbb{R} \times [-M, M]^t$ This will be required due to the special role played by the first variable of F_t.

Fix a positive constant $M > 1$. Let $x_M = \max(-M, \min(x, M))$ denote an M-truncation for any $x \in \mathbb{R}$. When x is a vector, x_M is assumed to be applied coordinate-wise. We now define the iterations forming the basis of AMP. Fix $U^0 \in [-M, M]^N$ and define the sequence $U^t \in \mathbb{R}^N, 1 \leq t \leq T$ as follows

$$U^t = \left[F_t(A(\cdot, f_t(U^0, \ldots, U^{t-1})), U^0, \ldots, U^{t-1}) \right]_M \in [-M, M]^N, \tag{8}$$

where F_t, f_t and M are applied component-wise. In other words, in step t, first a vector $f_t(U^0, \ldots, U^{t-1}) \in \mathbb{R}^N$ is formed by applying f_t coordinate-wise (recall that the domain of f_t is $\mathbb{R}^t$). Then this vector is used to define vector $A(\cdot, f_t(U^0, \ldots, U^{t-1}))$ via (3). This vector is concatinated with prior vectors $U^0, \ldots, U^{t-1}$ to form an $N \times (t+1)$ matrix $\left(A(\cdot, f_t(U^0, \ldots, U^{t-1})), U^0, \ldots, U^{t-1} \right) \in \mathbb{R}^{N \times (t+1)}$. Then function F_t is applied coordinate-wise. Finally the M-truncation is applied to each of the N coordinates of the vector thus obtained, resulting in U^t.

We now describe an algorithm which uses AMP to generate a solution in $\mathcal{B}_N$. For this purposes we assume that the algorithm designer has access to some (computable) projection function $\Pi_N : \mathcal{H}_N \to \mathcal{B}_N$. We discuss this further below.

Algorithm 2.3 (AMP Algorithm). *The algorithm is parametrized by $U^0, M, T, (f_t, 1 \leq t \leq T), (F_t, 1 \leq t \leq T), \Pi_N$.*

Input $A \in (\mathbb{R}^N)^{\otimes p}$.

Step 1 *Compute U^T using (8).*

Step 2 *Project U^T to $\mathcal{H}_N$ by applying transformation $x \to [x]_1 = \max(\min(x, 1), -1)$, coordinate-wise. Denote the resulting vector by $V \in \mathcal{H}_N$.*

Step 3 *Output $\sigma = \Pi(V) \in \mathcal{B}_N$.*

In some sense the details of the projection Π_N are immaterial to us since our negative result will be concerned with the quality of the solution V itself and not its projection. Nevertheless, for completeness we describe now the projection used in ([Mon18]), which we denote by Π_N^{sign}. The projection was defined only for $p = 2$ which was the case of interest. But it is straightforward to extend the idea to the case of general p. Set $z^{(0)} = V$. For $j = 1, \ldots, N$, construct $z^{(j)}$ by making all coordinates $\ell \neq j$ of $z^{(j)}$ to be the same as of $z^{(j-1)}$, and setting the j-th coordinate of $z^{(j)}$ to be the sign opposite of

$$\sum_{j \neq i_1 \neq i_2 \neq \cdots \neq i_{p-1}} A_{j, i_1, i_2, \ldots, i_{p-1}} z^{(j-1)}_{i_1, \ldots, i_{p-1}}. \tag{9}$$

In particular, the first j coordinates of $z^{(j)}$ are ± 1, but the remaining coordinates are real valued in general. Set $\Pi_N^{\text{sign}}(V) = z^{(N)}$.

Let us comment on the meaning and motivation behind the steps of the AMP algorithm above and also the motivation behind the projection Π_N described above and used in ([Mon18]). The idea is that when the AMP algorithm succeeds, the vector V, while not being an element of the binary cube $\mathcal{B}_N$, should be nearly optimal in the sense that

$$A(V) \approx \inf_{w \in \mathcal{B}_N} A(w),$$

and should not be too far from $\mathcal{H}_N$, so that the projecting U^T to $V \in \mathcal{H}_N$ does not change the objective value significantly. That is $A(V) \approx A(U^T)$. Next one observes that Π_N^{sign} effectively rounds V to a

7

vector $z^{(N)}$ in $\mathcal{B}_N$ in such a way that the objective value only decreases asymptotically. This is verified by observing that for each coordinate j, the dependence of $A(V)$ on variable V_j is linear in V_j, except for terms $A_{i_1,\dots,i_p}$ with repeating coordinates (i.e. such that $i_\ell = i_r$ for some $1 \le \ell \ne r \le p$). Since $V \in \mathcal{H}_N$ and thus $|V_j| \le 1$, the linearity allows to round V_j to -1 or 1 while only decreasing the objective value. This is done trivially by setting V_j to be the sign opposite of the one of the multiplier of V_j, which is (9). This is done iteratively over all N coordinates. The terms corresponding to repeating coordinates are easily shown to be of lower order of magnitude than the objective value. As a result one obtains a vector $z = z^{(N)} \in \mathcal{B}_N$ satisfying

$$A(z) \lesssim A(V) \approx \inf_{\sigma \in \mathcal{B}_N} A(\sigma).$$

But since z belongs to the solution space itself (the binary cube $\mathcal{B}_N$), it must be the case that in fact

$$A(z) \approx A(V) \approx \inf_{\sigma \in \mathcal{B}_N} A(\sigma),$$

and thus the success of AMP is validated. Importantly, the near optimality of z is argued from the near optimality of V itself. This discussion is of key essence to the main result of our work, which is stated in the next section.

3   The OGP conjecture and the main result

Consider an arbitrary set $\mathcal{A}$ of tensors $A \in (\mathbb{R}^N)^{\otimes p}$.

Definition 3.1. *The set $\mathcal{A}$ satisfies the Overlap Gap Property (OGP) with domain $\mathcal{S}_N \subset \mathbb{R}^N$, and parameters $\mu > 0, 0 < \nu_1 < \nu_2 < 1$ if for every pair $A_j \in \mathcal{A}, j = 1, 2$ and every $u_j, j = 1, 2$ satisfying*

$$\frac{1}{N} A_j(u_j) \le \frac{1}{N} \inf_{w \in \mathcal{S}_N} A_j(w) + \mu, \qquad j = 1, 2,$$

it holds

$$\frac{|\langle u_1, u_2 \rangle|}{\|u_1\|_2 \|u_2\|_2} \in [0, \nu_1] \cup [\nu_2, 1]. \tag{10}$$

Namely, every pair of nearly (μ-close) optimal solutions with respect to any two members of $\mathcal{A}$ cannot have normalized inner product in the interval (ν_1, ν_2).

Consider two independent random tensors A and $\hat{A}$ in $(\mathbb{R}^N)^{\otimes p}$ both with i.i.d. $\mathcal{N}\left(0, 1/N^{p-1}\right)$ entries. Introduce the interpolated set of tensors $A_\tau \triangleq \sqrt{1-\tau} A + \sqrt{\tau} \hat{A}$ with τ varying in $[0, 1]$. Note that for each fixed τ, A_τ is distributed as A. Our main conjecture regarding the OGP concerns the set $\mathcal{A} \triangleq (A_\tau, 0 \le \tau \le 1)$.

Conjecture 3.2. *For every even $p \ge 4$ here exists $\mu > 0, 0 < \nu_1 < \nu_2 < 1$ such that $\mathcal{A}$ described above satisfies the OGP with domain $\mathcal{S}_N = \mathcal{H}_N$ and parameters μ, ν_1, ν_2, with probability at least $1 - \exp(-cN)$, for some $c > 0$ for all sufficiently large N. Furthermore, for every $\delta > 0$ and every $v_1, v_2 \in \mathcal{H}_N$ satisfying $A(v_1)/N \le (1-\delta)\mathbb{E}[\eta_N], \hat{A}(v_2)/N \le (1-\delta)\mathbb{E}[\eta_N]$, it holds $|\langle v_1, v_2 \rangle| \le \delta N$ with probability at least $1 - \exp(-cN)$ for some $c > 0$ and all large N.*

Our main result stated below assumes the validity of this conjecture. To state this result, let us introduce the following. Let $\mathcal{M}_1([-M, M]^N)$ denote the space of probability measures on $[-M, M]^N$. Let $V(A, T, U^0)$ denote the output of the first two steps of Algorithm 2.3 after T steps with coefficient matrix A and initial data U^0, where the entires of $A \in (\mathbb{R}^N)^{\otimes p}$ are i.i.d. $\mathcal{N}(0, N^{-(p-1)})$. Then the following holds.

Theorem 3.3. *Let $p \geq 4$ be even. Let $M \geq 1$ and $\zeta > 0$. Assume that $(f_t), (F_t)$ satisfy Assumption 2.2 with Lipschitz constant ζ and that Conjecture 3.2 holds. Then there exists $\bar{\mu} > 0$ and $c > 0$, such that for N sufficiently large and any $\nu \in \mathcal{M}_1([-M, M]^N)$, if $U^0 \sim \nu$, then $V = V(A, T, U^0)$ satisfies*

$$\mathbb{P}\left(\frac{A(V)}{N} \geq \frac{\min_{\sigma \in \mathcal{B}_N} A(\sigma)}{N} + \bar{\mu}\right) \geq 1 - \exp(-cN).$$

Thus we argue the failure of the AMP to find a vector $V \in \mathcal{H}_N$ which is a near optimizer of A. As discussed earlier, this is a negative result regarding the performance of AMP, since finding such near optimal V is a key step towards finding a near optimal member z of the binary cube $\mathcal{B}_N$. Ideally, one would establish that the vector z obtained from V via any projection scheme, such as the one described above is also μ-away from optimality. Unfortunately, our proof technique stops short of that due to the potential sensitivity of the sign function used on obtaining z to perturbation of A, thus potentially violating stability used crucially in the proof of our main result. We leave this as an interesting open question.

A partial support to the validity of Conjecture 3.2 above is its validity for the domain $\mathcal{S}_N = \mathcal{B}_N$ as we now establish.

Theorem 3.4. *For every even $p \geq 4$ here exists $\mu > 0, 0 < \nu_1 < \nu_2 < 1$ such that $\mathcal{A}$ described above satisfies the OGP with domain $\mathcal{S} = \mathcal{B}_N$ and parameters μ, ν_1, ν_2, with probability at least $1 - \exp(-cN)$, for some $c > 0$ for all sufficiently large N. Furthermore, for every $\delta > 0$ and every $\sigma_1, \sigma_2 \in \mathcal{B}_N$ satisfying $A(\sigma_1)/N \leq (1-\delta)\mathbb{E}[\eta_N], \hat{A}(\sigma_2)/N \leq (1-\delta)\mathbb{E}[\eta_N]$, it holds $|\langle \sigma_1, \sigma_2 \rangle| \leq \delta N$ with probability at least $1 - \exp(-cN)$ for some $c > 0$ and all large N.*

Proof. We note that in the case $\mathcal{S}_N = \mathcal{B}_N$, since $\|\sigma\|_2 = N$ for each $\sigma \in \mathcal{B}_N$, the requirement (10) in definition of OGP simplifies to

$$\frac{|\langle \sigma_1, \sigma_2 \rangle|}{N} \in [0, \nu_1] \cup [\nu_2, 1].$$

It was established in [CGP$^+$19], Theorem 3 that the OGP holds for a single instance of a tensor A, i.e. $\mathcal{A} = \{A\}$, with probability at least $1 - \exp(-cN)$ for some $c > 0$ and all large N. At the same time the following chaos property was established in [CHL18]:

Theorem 3.5 ([CHL18], Theorem 2). *For every $\epsilon > 0$ and $\tau \in (0,1)$ there exists $C, \tilde{\mu} > 0$ such that with probability $1 - \exp(-CN)/C$, for every $\sigma_1, \sigma_2 \in \mathcal{B}_N$ satisfying $A(\sigma_1)/N \leq \mathbb{E}[\eta_N] + \tilde{\mu}, A_\tau(\sigma_2)/N \leq \mathbb{E}[\eta_N] + \tilde{\mu}$ it holds $|\langle \sigma_1, \sigma_2 \rangle| \leq \epsilon N$.*

We now combine these two results. We first claim that it suffices to establish OGP for a discrete finite subsets. Fix $\delta > 0$ such that $1/\delta$ is an integer and consider A_τ for $\tau = 0, \delta, 2\delta, \ldots, \delta(1/\delta)$. We assume OGP holds for this set for some μ, ν_1, ν_2 for every sufficiently small such δ. Now for any $\sigma \in \mathcal{B}_N$

$$A_\tau(\sigma) - A(\sigma) = (\sqrt{1-\tau} - 1)A(\sigma) + \sqrt{\tau}A(\sigma).$$

In light of concentration bound of Theorem 2.1, for any $\epsilon > 0$ we can find small enough δ so that

$$\max_{0 \leq k \leq (1/\delta)-1} \sup_{0 \leq \tau \leq \delta} \max_{\sigma \in \mathcal{B}_N} |A_{k\delta+\tau}(\sigma) - A_{k\delta}(\sigma)| \leq \epsilon N$$

with probability at least $1 - \exp(-cN)$, for some $c > 0$ and large N. This means that modulo exponentially small probability, every σ satisfying $A_{k\delta+\tau}(\sigma)/N \leq \mathbb{E}[\eta_N] + \epsilon$ also satisfies $A_{k\delta}(\sigma)/N \leq \mathbb{E}[\eta_N] + 2\epsilon$. Thus if $\epsilon < \mu - 2\epsilon$ then the set $(A_\tau, \tau \in [0,1])$ satisfies OGP with $\hat{\mu} = \mu - 2\epsilon > 0$ and the same ν_1, ν_2,

provided that the discrete set $(A_{k\delta}, 0 \leq k \leq 1/\delta)$ satisfies OGP with μ, ν_1, ν_2. Thus we now prove OGP for this discrete set.

Let μ, ν_1, ν_2 be OGP parameters for a single instance A. By the union bounds over $k = 0, 1, \ldots, 1/\delta$, the OGP holds for each $A_{k\delta}$ modulo exponentially small probability. Fix $\delta > 0$. Applying Theorem 3.5, we find $\tilde{\mu}$ so that the theorem claim holds for $\epsilon = \nu_1$ and $\tau = \delta$. By union bounds this also holds for all pairs $A_{k_1\delta}, A_{k_2\delta}, k_1 \neq k_2$ modulo exponentially small probability. Then OGP holds for $\bar{\mu} \triangleq \min(\tilde{\mu}, \mu), \nu_1, \nu_2$ by considering separately the cases $k_1 = k_2$ and $k_1 \neq k_2$, where in the latter case for every $\sigma_j, j = 1, 2$ satisfying $A_{k_j\delta}(\sigma_j)/N \leq \mathbb{E}[\eta_N] + \bar{\mu}, j = 1, 2$, we simply have $|\langle \sigma_1, \sigma_2 \rangle| \leq \nu_1 N$.

The second part of the theorem follows immediately from the chaos property of Theorem 3.5 in the special case $\tau = 1$. $\square$

Conjecture 3.2 would follow from Theorem 3.4 if we could establish that every nearly optimal solution in $\mathcal{H}_N$ is actually close to a point in $\mathcal{B}_N$. This is quite plausible as one does not expect nearly optimal solutions to exist "deep" inside the Hilbert cube $\mathcal{H}_N$. Unfortunately, we are not able to show this and thus state it as an interesting open problem.

Conjecture 3.6. *Suppose the entries of $A \in (\mathbb{R}^N)^{\otimes p}$ are generated i.i.d. according to $\mathcal{N}(0, N^{-(p-1)})$. For every $\epsilon > 0$ there exists $\delta > 0$ such that with probability at least $1 - \exp(-cN)$ for some $c > 0$ and large enough N, every $u \in \mathcal{H}_N$ satisfying $A(u)/N \leq (1-\delta)\mathbb{E}[\eta_N]$ also satisfies $\min_{v \in \mathcal{B}_N} \|u - v\|_2 \leq \epsilon\sqrt{N}$.*

4 Preliminary technical results

In this section we establish several preliminary results. We begin by recalling the following operator norm bound

Lemma 4.1. *There exist constants $C, c > 0$, such that*

$$\mathbb{P}(\|A\|_{\mathrm{op}} > CN^{1-p/2}) \leq e^{-cN},$$

for all sufficiently large N.

Proof. The proof of this result is verbatim that from [BAGJ18b, Lemma 4.7]. We include this for completeness.

Let $\mathbb{S}^N = \{x : \|x\|_2 = 1\}$ denote the unit ℓ_2- ball. We may then view A as a centered Gaussian process on $(\mathbb{S}^N)^{\times p}$, with covariance

$$\mathbb{E}[A(x_1, .., x_p)A(y_1, ..., y_p)] = \frac{1}{N^{(p-1)}} \prod_{1 \leq i \leq p} \langle x_i, y_i \rangle.$$

This process is rotationally invariant. Fix an $\epsilon > 0$, let Σ_ϵ denote an ϵ-net for $\mathbb{S}^N$ with respect to $\|\cdot\|_2$ norm, and let Σ_ϵ^p denote is p-fold cartesian product. By multilinearity of A,

$$\|A\|_{\mathrm{op}} \leq \sup_{(x_1, \ldots, x_k) \in \Sigma_\epsilon^p} A(x_1, ..., x_p) + \epsilon p \|A\|_{\mathrm{op}}.$$

If we choose ϵ so that $2p\epsilon \leq 1$, we have

$$\mathbb{P}(\|A\|_{\mathrm{op}} > \lambda) \leq \mathbb{P}\left(\cup_{\mathbf{x} \in \Sigma_\epsilon^p} \{A(x_1, ..., x_p) \geq \lambda/2\}\right).$$

To bound the right hand side, note that for any fixed $(x_1, .., x_p) \in (\mathbb{S}^N)^p$, $A(x_1, ..., x_p)$ is a centered Gaussian with variance N^{-p+1}, so that

$$\mathbb{P}(A(x_1, ..., x_p) \geq \lambda N^{1-p/2}) \leq e^{-N\lambda^2/2}.$$

10

where in the second line, $(x_1, ..., x_p)$ is any point in Σ^p, $|\Sigma^k|$ denotes its cardinality, and the final inequality comes from a Gaussian tail bound since $A(x_1, .., x_p)$ is a centered Gaussian with variance N^{-p+1} . Note furthermore that we may choose this net so that $|\Sigma_\epsilon| \leq (4/\epsilon)^N$, [Ver18, Lemma 5.1]. Thus by rotation invariance and a union bound, we see that

$$\mathbb{P}(\|A\|_{\mathrm{op}} \geq \lambda N^{1-p/2}) \leq \left(\frac{4}{\epsilon}\right)^{pN} e^{-N\lambda^2/2}.$$

Choosing λ sufficiently large yields the result. $\qquad\square$

Proposition 4.2. *There exists $c_2, c > 0$ which depend on M such that*

$$\mathbb{P}(\max_{u,v \in [-M,M]^N} \frac{\|A(\cdot, u) - A(\cdot, v)\|_2}{\|u - v\|_2} \geq c_2) \leq \exp(-cN),$$

for all sufficiently large N.

Proof. By multilinearity of A, the triangle inequality, and the definition of the operator norm,

$$\|A(\cdot, u, \ldots, u) - A(\cdot, v, \ldots, v)\|_2 \leq \|A(\cdot, u - v, \ldots, u)\|_2 + \ldots + \|A(\cdot, v, \ldots, v, u - v)\|_2$$
$$\leq \|A\|_{\mathrm{op}} \max\{\|v\|_2^{p-2}, \|u\|_2^{p-2}\}\|u - v\|_2.$$

Since $u, v \in [-M, M]^N$, it follows that $\|v\|_2, \|u\|_2 \leq M\sqrt{N}$. Applying the bound from Lemma 4.1, we obtain the result. $\qquad\square$

Lemma 4.3. *There exist $c, C > 0$ such that with probability for every at least $1 - \exp(-cN)$ for all large N the following holds: for every $\eta > 0$, every $u \in H_N$ satisfying $A(u) \leq -\eta N$ also satisfies $\|u\|_2 \geq C\eta^{1/p}\sqrt{N}$.*

Proof. Note that by Lemma 4.1, with probability at least $1 - \exp(-cN)$

$$|\max_{\|u\|_2 \leq \delta\sqrt{N}} A(u)| \leq \delta^p N^{\frac{p}{2}} \cdot \|A\|_{\mathrm{op}} \leq C\delta^p N.$$

Thus if $A(u) \leq -\eta N$, it must be that $\|u\|_2^p \|A\|_{\mathrm{op}} \geq \eta N$, so that

$$\|u\|_2 \geq \frac{1}{C}\eta^{1/p}N^{1/2},$$

where C is as in Lemma 4.1. $\qquad\square$

5 Continuous dependence

When we view Algorithm 2.3 as a discrete time dynamical system, it is natural to expect that this admits a similar dependence on the tensor A as a time-inhomogenous differential equation of the same form. Thus our proof of continuous dependence of iterations on the tensor A can be viewed as a discrete analogue of similar standard result for differential equations, see, e.g., [Tes12, Section 2.4].

Given any tensor $B \in (\mathbb{R}^N)^{\otimes p}$, let

$$c_2(B) \triangleq \sup_{u \neq v \in [-M,M]^N} \frac{\|B(\cdot, u) - B(\cdot, v)\|_2}{\|u - v\|_2}. \tag{11}$$

We now state the main result of this section.

Theorem 5.1. *Let* $B, \hat{B} \in (\mathbb{R}^N)^{\otimes p}$, *and let* $\hat{V}^t, V^t$ *denote the corresponding sequences output by Step 2 of Algorithm 2.3, with the same initial vector* $U^0 = \hat{U}^0$. *Under Assumption 2.2, there is some constant* K *which depends only on* ζ *and* $c_2(\hat{B})$, *such that for every* $T \geq 1$ *and* U^0,

$$\sup_{1 \leq t \leq T} \|\hat{V}^t - V^t\| \leq K^T \|\hat{B} - B\|_{\mathrm{op}} \left(\zeta M \sqrt{NT}\right)^{p-1}.$$

Proof. Define U^t and $\hat{U}^t$ as in Step 1 of Algorithm 2.3, and let $\mathbf{U}_t = (U^s)_{0 \leq s \leq t}$ and $\hat{\mathbf{U}}_t = \left(\hat{U}^s\right)_{0 \leq s \leq t}$. Since the map $f(x) = [x]_M$ is 1-Lipschitz for any M, we see that the claim of the theorem follows, provided that

$$\beta_N(t) \triangleq \|\mathbf{U}_t - \hat{\mathbf{U}}_t\|_2 = \sqrt{\sum_{s \leq t} \|U^s - \hat{U}^s\|_2^2}$$

satisfies

$$\beta_N(t) \leq (K(c_2(\hat{B}) + 1))^t \|\hat{B} - B\|_{\mathrm{op}} \left(\zeta M \sqrt{Nt}\right)^p, \tag{12}$$

as trivially $\|\hat{U}^t - U^t\|_2 \leq \beta_N(t)$.

Thus we establish (12). By 1-Lipschitz continuity of $[\cdot]_M$ we have

$$\|\hat{U}^{t+1} - U^{t+1}\|_2 \leq$$
$$\|F_{t+1}(\hat{B}(\cdot, f_{t+1}(\hat{U}^t, \hat{U}^{t-1}, \ldots, \hat{U}_0)), \hat{U}^0, \ldots, \hat{U}^t) - F_{t+1}(B(\cdot, f_{t+1}(U^0, \ldots, U^t)), U_0, \ldots, U^t)\|_2.$$

Applying the part of Assumption 2.2 regarding F_t, we see that

$$\|\hat{U}^{t+1} - U^{t+1}\|_2 \leq \zeta \sqrt{\beta_N^2(t) + \|\hat{B}(\cdot, f_{t+1}(\hat{\mathbf{U}}_t)) - B(\cdot, f_{t+1}(\mathbf{U}_t))\|_2^2},$$

so that

$$\beta_N(t+1) \leq (1 + \zeta^2) \sqrt{\beta_N^2(t) + \|\hat{B}(\cdot, f_{t+1}(\hat{\mathbf{U}}_t)) - B(\cdot, f_{t+1}(\mathbf{U}_t))\|_2^2}.$$

By the triangle inequality,

$$\|\hat{B}(\cdot, f_{t+1}(\hat{\mathbf{U}}_t)) - B(\cdot, f_{t+1}(\mathbf{U}_t))\|_2$$
$$\leq \|\hat{B}(\cdot, f_{t+1}(\hat{\mathbf{U}}_t)) - \hat{B}(\cdot, f_{t+1}(\mathbf{U}_t))\|_2 + \|\hat{B}(\cdot, f_{t+1}(\mathbf{U}_t)) - B(\cdot, f_{t+1}(\mathbf{U}_t))\|_2$$
$$= I + II.$$

By definition of $c_2(\hat{B})$,

$$I \leq c_2(\hat{B}) \|f_{t+1}(\hat{\mathbf{U}}_t) - f_{t+1}(\mathbf{U}_t)\|_2 \leq \zeta c_2(\hat{B}) \beta_N(t).$$

We now analyze II. Note that by Assumption 2.2

$$\|f_{t+1}(\mathbf{U}_t)\|_2^2 \leq \zeta^2 \sum_{0 \leq i \leq t} \|U^i\|_2^2 \leq \zeta^2 (t+1) M^2 N.$$

Then

$$II \leq \|\hat{B} - B\|_{\mathrm{op}} (M\zeta\sqrt{N(t+1)})^{p-1}.$$

Combining these bounds, we obtain,

$$\|B(\cdot, f_{t+1}(\mathbf{U}_t)) - \hat{B}(\cdot, f_{t+1}(\hat{\mathbf{U}}_t))\|_2 \leq \zeta c_2(\hat{B}) \beta_N(t) + \|\hat{B} - B\|_{\mathrm{op}} \left(\zeta M \sqrt{N(t+1)}\right)^{p-1}.$$

12

Plugging this in to the above, yields

$$\beta_N(t+1) \le (1+\zeta^2)\sqrt{\beta_N^2(t) + \left(\zeta c_2(\hat{B})\beta_N(t) + \|\hat{B} - B\|_{\mathrm{op}}\left(\zeta M\sqrt{Nt}\right)^{p-1}\right)^2}.$$

We can write the inequality above in the form

$$\beta_N(t+1) \le K\beta_N(t) + b(t),$$

where $b(t)$ is non-decreasing, and $K > 1$ which depends only on $c_2(\hat{B})$ and ζ. The inequality above is a discrete version of Gronwall's inequality and using $\beta_N(0) = 0$, easily leads to a bound

$$\beta_N(t) \le K^t b(t) = K^t \|\hat{B} - B\|_{\mathrm{op}}\left(\zeta M\sqrt{Nt}\right)^{p-1}. \qquad \square$$

6 Concentration property of the AMP solution

In this section we establish that the value associated with the solution V produced by the AMP is concentrated around its expectation.

Theorem 6.1. *Suppose that Assumption 2.2 holds. For any ϵ, M, T, ζ, there exists $c > 0$ such that such that the value $A(V)$ associated with the solution V produced in Step 2 of Algorithm 2.3 satisfies*

$$\sup_{U^0 \in [-M,M]^N} \mathbb{P}\left(|A(V) - \mathbb{E}[A(V)|U^0]| \ge \epsilon N|U^0\right) \le \exp(-cN),$$

for all sufficiently large N.

Proof. Fix U^0. Our approach is based on Gaussian concentration combined with Kirszbraun's Theorem. Let $Z \in (\mathbb{R}^N)^{\otimes p}$ denote a tensor consisting of i.i.d. standard normal entries, so that $A = Z/N^{\frac{p-1}{2}}$ in distribution. We let $f(Z) = A(V(A)) = Z(V(Z/N^{\frac{p-1}{2}}))/N^{\frac{p-1}{2}}$, where $V = V(Z)$ is again the solution produced by AMP viewed as a function of Z, and thus introduce $f : (\mathbb{R}^N)^{\otimes p} \to \mathbb{R}$ defined by $f(z) = z(V(z/N^{\frac{p-1}{2}}))/N^{\frac{p-1}{2}}$. We first establish that this function is Lipschitz with an appropriate constant on an appropriate subspace of $(\mathbb{R}^N)^{\otimes p}$. Recall the constant c_2 introduced in Proposition 4.2. Let

$$K_{2,N} = \left\{ z \in (\mathbb{R}^N)^{\otimes p} : c_2\left(\frac{z}{N^{\frac{p-1}{2}}}\right) \le c_2 \right\}.$$

In particular, a random Z with i.i.d. standard normal entries satisfies

$$\mathbb{P}\left(Z \in K_{2,N}\right) \ge 1 - \exp(-cN), \tag{13}$$

for all large enough N, where c is as in the proposition.

Lemma 6.2. *There exists a constant $c = c(M, c_2, \zeta, T)$ such that for every $z_1, z_2 \in K_{2,N}$*

$$|f(z_2) - f(z_1)| \le c\sqrt{N}\|z_2 - z_1\|_2.$$

Proof. Applying Theorem 5.1, for any $z_1, z_2 \in K_{2,N}$ we have

$$\|V(z_2/N^{\frac{p-1}{2}}) - V(z_1/N^{\frac{p-1}{2}})\|_2 \le cN^{\frac{p-1}{2}}\|N^{-\frac{p-1}{2}}(z_2 - z_1)\|_{\text{op}}$$
$$= c\|z_2 - z_1\|_{\text{op}}, \tag{14}$$

where $c = c(M, c_2, \zeta, T)$ is an appropriate constant.

Next,

$$|f(z_1) - f(z_2)| = N^{-\frac{p-1}{2}}|z_2(V(z_2/N^{\frac{p-1}{2}})) - z_1(V(z_1/N^{\frac{p-1}{2}}))|$$
$$\le N^{-\frac{p-1}{2}}|z_2(V(z_2/N^{\frac{p-1}{2}})) - z_1(V(z_2/N^{\frac{p-1}{2}}))|$$
$$+ N^{-\frac{p-1}{2}}|z_1(V(z_2/N^{\frac{p-1}{2}})) - z_1(V(z_1/N^{\frac{p-1}{2}}))| \tag{15}$$

We first analyze the second summand above. For simplicity we use v_1, v_2 in place of $V(z_1/N^{\frac{p-1}{2}})$ and $V(z_2/N^{\frac{p-1}{2}})$. Note $z_1(u) = \langle u, z(\cdot, u)\rangle$. Thus

$$|z_1(v_2) - z_1(v_1)| = |\langle v_2, z_1(\cdot, v_2)\rangle - \langle v_1, z_1(\cdot, v_1)\rangle|$$
$$\le |\langle v_2, z_1(\cdot, v_2)\rangle - \langle v_2, z_1(\cdot, v_1)\rangle| + |\langle v_2, z_1(\cdot, v_1)\rangle - \langle v_1, z_1(\cdot, v_1)\rangle|.$$

Then

$$|\langle v_2, z_1(\cdot, v_2)\rangle - \langle v_2, z_1(\cdot, v_1)\rangle| = |\langle v_2, z_1(\cdot, v_2) - z_1(\cdot, v_1)\rangle|$$
$$\le \|v_2\|_2 \|z_1(\cdot, v_2) - z_1(\cdot, v_1)\|_2$$
$$\le M\sqrt{N}N^{\frac{p-1}{2}}c_2(z_1/N^{\frac{p-1}{2}})\|v_2 - v_1\|_2.$$

Since $z_1 \in K_{2,N}$, we obtain instead a bound

$$M\sqrt{N}N^{\frac{p-1}{2}}c_2\|v_2 - v_1\|_2 \le M\sqrt{N}N^{\frac{p-1}{2}}c_2 c\|z_2 - z_1\|_{\text{op}},$$

where the inequality follows from (14).

For the second term we have

$$|\langle v_2, z_1(\cdot, v_1)\rangle - \langle v_1, z_1(\cdot, v_1)\rangle| \le \|v_2 - v_1\|_2 \|z_1(\cdot, v_1)\|_2$$

Since $z_1 \in K_{2,N}$ and $z_1(\cdot, 0) = 0$, then

$$\|z_1(\cdot, v_1)\|_2 \le N^{\frac{p-1}{2}}c_2\|v_1\|_2 \le N^{\frac{p-1}{2}}c_2 M\sqrt{N}.$$

Using (14) to $\|v_2 - v_1\|_2$, we obtain a bound $c\|z_2 - z_1\|_{\text{op}}N^{\frac{p-1}{2}}c_2 M\sqrt{N}$.

Applying both bounds to (15) and using $\|\cdot\|_{\text{op}} \le \|\cdot\|_2$ we complete the proof. $\square$

We now complete the proof of the theorem. For every $z \in (\mathbb{R}^N)^{\otimes p}$, define

$$g(z) = \inf_{\hat{z} \in K_{2,N}} \left(f(\hat{z}) + c\sqrt{N}\|\hat{z} - z\|_2 \right),$$

where c is as in Lemma 6.2. Kirszbraun's Theorem says that g is a Lipschitz continuous function with constant $c\sqrt{N}$ and $g = f$ on $K_{2,N}$. This is easy to verify. Indeed, fix any $z_1, z_2 \in (\mathbb{R}^N)^{\otimes p}$ and $\epsilon > 0$. Find $\hat{z}_1 \in K_{2,N}$ such that

$$|g(z_1) - \left(f(\hat{z}_1) + c\sqrt{N}\|\hat{z}_1 - z_1\|_2 \right)| \le \epsilon.$$

14

Then

$$g(z_2) - g(z_1) \leq f(\hat{z}_1) + c\sqrt{N}\|\hat{z}_1 - z_2\|_2 - \left(f(\hat{z}_1) + c\sqrt{N}\|\hat{z}_1 - z_1\|_2\right) + \epsilon$$
$$= c\sqrt{N}\left(\|\hat{z}_1 - z_2\|_2 - \|\hat{z}_1 - z_1\|_2\right) + \epsilon$$
$$\leq c\sqrt{N}\|z_2 - z_1\|_2.$$

Using a similar reversed inequality, the Lipschitz continuity of g is established. Now if $z \in K_{2,N}$ then by Lemma 6.2 for every $\hat{z} \in K_{2,N}$

$$f(\hat{z}) + c\sqrt{N}\|\hat{z} - z\|_2 \geq f(z),$$

implying that the infimum is achieved by $\hat{z} = z$, establishing the Kirszbraun's Theorem.

In conclusion, g is a Lipschitz continuous function with constant $c\sqrt{N}$. Thus by Gaussian concentration (see, e.g., [Ver18]), for every $t \geq 0$

$$\mathbb{P}\left(|g(Z) - \mathbb{E}\big[g(Z)|U^0\big]| \geq tN|U^0\right) \leq \exp\left(-\frac{t^2 N}{4c^2}\right).$$

We now use the fact that $f = g$ on the high probability set $K_{2,N}$. Specifically

$$\mathbb{E}\big[g(Z)|U^0\big] = \mathbb{E}\big[f(Z)|U^0\big] - \mathbb{E}\big[f(Z)\mathbf{1}\left(Z \notin K_{2,N}\right)|U^0\big] + \mathbb{E}\big[g(Z)\mathbf{1}\left(Z \notin K_{2,N}\right)|U^0\big].$$

Using $g(Z) \leq f(0) + c\sqrt{N}\|Z\|_2 = c\sqrt{N}\|Z\|_2$ we have

$$\mathbb{E}\big[g(Z)\mathbf{1}\left(Z \notin K_{2,N}\right)|U^0\big] \leq c\sqrt{N}\mathbb{E}[\|Z\|_2\mathbf{1}\left(Z \notin K_{2,N}\right)]$$
$$\leq c\sqrt{N}\left(\mathbb{E}\big[\|Z\|_2^2\big]\right)^{\frac{1}{2}}\mathbb{P}^{\frac{1}{2}}\left(Z \notin K_{2,N}\right)$$
$$\leq \exp(-c_4 N),$$

for some appropriately chosen $c > 0$ and all sufficiently large N, where in the second line, we used that U^0 and A are independent, and the last inequality follows from (13) and from $\mathbb{E}\big[\|Z\|_2^2\big] = N^{O(1)}$. Similarly, since $f(Z) \leq N^{-\frac{p-1}{2}}\sum|Z_{i_1,\dots,i_p}|$, we also have $\mathbb{E}\big[f^2(Z)|U^0\big] = N^{O(1)}$ and thus $\mathbb{E}\big[f(Z)\mathbf{1}\left(Z \notin K_{2,N}\right)|U^0\big]$ is at most $\exp(-c_4 N)$ for all large enough N, where we used the same notation for constant c_4 as above for convenience. We conclude

$$|\mathbb{E}\big[g(Z)|U^0\big] - \mathbb{E}\big[f(Z)|U^0\big]| \leq \exp(-c_5 N),$$

for some $c_5 > 0$ and all large N.

Thus for any $t > 0$

$$\mathbb{P}\left(|f(Z) - \mathbb{E}\big[f(Z)|U^0\big]| \geq tN\right)$$
$$\leq \mathbb{P}\left(|g(Z) - \mathbb{E}\big[f(Z)|U^0\big]| \geq tN, \mathbf{1}\left(Z \in K_{2,N}\right)\right) + \mathbb{P}(Z \notin K_{2,N})$$
$$\leq \mathbb{P}\left(|g(Z) - \mathbb{E}\big[g(Z)|U^0\big]| \geq tN - \left(\mathbb{E}\big[g(Z)|U^0\big] - \mathbb{E}\big[f(Z)|U^0\big]\right)\right) + \exp(-CN)$$
$$\leq \mathbb{P}\left(|g(Z) - \mathbb{E}\big[g(Z)|U^0\big]| \geq (t/2)N\right) + \exp(-CN),$$
$$\leq \exp(-c_6 N),$$

for all large enough N and appropriately chosen $c_6 > 0$. As U^0 was arbitrary the result then follows. $\square$

7 OGP is an obstruction to AMP. Proof of the main result

In this section we complete the proof of the main result, Theorem 3.3. Let us begin by first conditioning on the value of U^0. Let $A \in (\mathbb{R}^N)^{\otimes p}$ be a tensor with i.i.d. $\mathcal{N}(0, 1/N^{p-1})$ entries. Recall that by assumption A and U^0 are independent. Let $V = V(A)$ be the result of the Step 2 of Algorithm 2.3 after T steps. Applying the concentration properties given by Theorems 2.1 and 6.1, it suffices to show that for every $\epsilon > 0$,

$$\frac{\mathbb{E}[A(V)|U^0]}{N} \geq \mathbb{E}[\eta_N] + \mu - \epsilon,$$

for all large enough N, where μ is as in Conjecture 3.2, as in this case the main result would be established for $\bar{\mu} = \mu - 2\epsilon$ for every $\epsilon > 0$.

Thus for the purposes of contradiction, assume

$$\mathbb{E}[A(V)|U^0] / N \leq \mathbb{E}[\eta_N] + \mu_2. \tag{16}$$

for some $\mu_2 < \mu$ for infinitely many N.

Generate a tensor $\hat{A} \in (\mathbb{R}^N)^{\otimes p}$ distributed as A and independent from A and U^0. Consider the interpolated set $\mathcal{A} = (A_\tau, \tau \in [0,1])$ described in Section 3. Denote by $\mathcal{E}_{\text{OGP}}$ the high probability OGP event defined in Conjecture 3.2 with parameters μ, ν_1, ν_2. Let V_τ be the vector produced by AMP when run on tensor $A_\tau, \tau \in [0,1]$. For any $\tau \in [0,1]$ we have

$$c_2(A_\tau) = \sup_{u \neq v \in [-M,M]^N} \frac{\|\sqrt{1-\tau}A(\cdot, u) + \sqrt{\tau}\hat{A}(\cdot, u) - (\sqrt{1-\tau}A(\cdot, v) + \sqrt{\tau}\hat{A}(\cdot, v))\|_2}{\|u - v\|_2}$$
$$\leq \sqrt{1-\tau}c_2(A) + \sqrt{\tau}c_2(\hat{A})$$
$$\leq c_2(A) + c_2(\hat{A}).$$

Applying Proposition 4.2 we have $c_2(A) + c_2(\hat{A}) \leq 2c_2$ modulo exponentially small in N probability. We conclude that $\sup_{\tau \in [0,1]} c_2(A_\tau) \leq 2c_2$ modulo exponentially small probability.

By Theorem 5.1 then modulo exponentially small probability, using $\|\cdot\|_{\text{op}} \leq \|\cdot\|_2$, we have that for any $\tau_1, \tau_2 \in [0,1]$

$$\|V^{\tau_1} - V^{\tau_2}\|_2 \leq C^T N^{\frac{p-1}{2}} \|A_{\tau_1} - A_{\tau_2}\|_2,$$

for some constant $C > 0$ which incorporates c_2, ζ, and M (and which may change from line to line). We have

$$\|A_{\tau_1} - A_{\tau_2}\|_2 = \|\sqrt{1-\tau_1}A - \sqrt{1-\tau_2}A + \sqrt{\tau_1}\hat{A} - \sqrt{\tau_2}\hat{A}\|_2$$
$$\leq \left(|\sqrt{1-\tau_1} - \sqrt{1-\tau_2}| + |\sqrt{\tau_1} - \sqrt{\tau_2}|\right) \left(\|A\|_2 + \|\hat{A}\|_2\right).$$

Since $\|A\|_2^2$ is distributed as $N^{-p-1} \sum_{1 \leq i \leq N^p} Z_i^2$, which is N in expectation, then by standard large deviations estimates, $\mathbb{P}(\|A\|_2 \geq cN)$ is exponentially small for any $c > 1$. In particular, $\|A\|_2 + \|\hat{A}\|_2 \leq 4\sqrt{N}$, modulo exponentially small probability.

Combining and letting $h(\tau_1, \tau_2) = |\sqrt{1-\tau_1} - \sqrt{1-\tau_2}| + |\sqrt{\tau_1} - \sqrt{\tau_2}|$ we obtain

$$\|\|V^{\tau_1}\|_2 - \|V^{\tau_2}\|_2\| \leq \|V^{\tau_1} - V^{\tau_2}\|_2 \leq C^T N^{\frac{p}{2}} h(\tau_1, \tau_2), \tag{17}$$

for all τ_1, τ_2 modulo exponentially small in N probability.

Next, for any $\tau_1, \tau_2 \in [0,1]$

$$\frac{\langle V^0, V^{\tau_2} \rangle}{\|V^0\|_2 \|V^{\tau_2}\|_2} - \frac{\langle V^0, V^{\tau_1} \rangle}{\|V^0\|_2 \|V^{\tau_1}\|_2} = \frac{\|V^{\tau_1}\|_2 \langle V^0, V^{\tau_2} \rangle - \|V^{\tau_2}\|_2 \langle V^0, V^{\tau_1} \rangle}{\|V^0\|_2 \|V^{\tau_1}\|_2 \|V^{\tau_2}\|_2}.$$

For the numerator, applying the Cauchy-Scwhartz inequality

$$|\|V^{\tau_2}\|_2 \langle V^0, V^{\tau_1} \rangle - \|V^{\tau_1}\|_2 \langle V^0, V^{\tau_2} \rangle|$$
$$= |\|V^{\tau_2}\|_2 \langle V^0, V^{\tau_1} \rangle - \|V^{\tau_1}\|_2 \langle V^0, V^{\tau_1} \rangle + \|V^{\tau_1}\|_2 \langle V^0, V^{\tau_1} \rangle - \|V^{\tau_1}\|_2 \langle V^0, V^{\tau_2} \rangle|$$
$$\leq |\|V^{\tau_2}\|_2 - \|V^{\tau_1}\|_2| \; \|V^0\|_2 \|V^{\tau_1}\|_2 + \|V^{\tau_1}\|_2 \|V^0\|_2 \|V^{\tau_2} - V^{\tau_1}\|_2$$

Provided (17) holds, we obtain

$$\left| \frac{\langle V^0, V^{\tau_2} \rangle}{\|V^0\|_2 \|V^{\tau_2}\|_2} - \frac{\langle V^0, V^{\tau_1} \rangle}{\|V^0\|_2 \|V^{\tau_1}\|_2} \right| \leq \frac{2C^T N^{\frac{p}{2}} h(\tau_1, \tau_2)}{\|V^{\tau_2}\|_2}.$$

Next we fix $\alpha > 0$, to be specified later, let $\delta = N^{-\alpha}$. We assume for convenience that $1/\delta = N^\alpha$ is an integer. Introduce a discrete sequence $\tau_n = n\delta, 0 \leq n \leq 1/\delta$. By Lemma 4.3, applying union in N^α terms bound, we have that $\|V^{\tau_n}\| \geq C_2 \sqrt{N}$ for some $C_2 > 0$ for all sufficiently large N, modulo exponentially small probability. Provided this holds, the bound above can be replaced by

$$\frac{2C^T N^{\frac{p}{2}} h(\tau_1, \tau_2)}{C_2 \sqrt{N}} = C^T h(\tau_1, \tau_2) N^{\frac{p-1}{2}}.$$

Now we have

$$h(\tau_{n_1}, \tau_{n_2}) \leq C^T N^{\frac{\alpha}{2}} (\tau_{n_2} - \tau_{n_1}) = C^T N^{\frac{\alpha}{2}} N^{-\alpha} (n_2 - n_1) = C^T N^{-\frac{\alpha}{2}} (n_2 - n_1).$$

for all n_1, n_2 and some $C > 0$.

Combining, we conclude that modulo exponentially small in N probability, for all $n = 0, \ldots, N^\alpha$,

$$\left| \frac{\langle V^0, V^{\tau_{n+1}} \rangle}{\|V^0\|_2 \|V^{\tau_{n+1}}\|_2} - \frac{\langle V^0, V^{\tau_n} \rangle}{\|V^0\|_2 \|V^{\tau_n}\|_2} \right| \leq C^T N^{\frac{p-1}{2}} N^{-\frac{\alpha}{2}}$$

and provided $\alpha > p-1$ the bound above is $o(1)$, and in particular is smaller than $\nu_2 - \nu_1$ for N sufficiently large.

Next we examine $\frac{\langle V^0, V^{\tau_n} \rangle}{\|V^0\|_2 \|V^{\tau_n}\|_2}$ in the extreme case $n = 0$ and $n = N^\alpha$. The value is clearly 1 when $n = 0$. Applying the second part of Conjecture 3.2 we have that for every ϵ, this value is at most ϵ/C^2 modulo exponentially small probability, where C is the constant from Lemma 4.3. In particular at $n = N^\alpha$ this value is at most ν_1. It follows that there must exist an index n^*, (which is random in general) such that

$$\left| \frac{\langle V^0, V^{\tau_{n^*}} \rangle}{\|V^0\|_2 \|V^{\tau_{n^*}}\|_2} \right| \in (\nu_1, \nu_2).$$

Now in the event that OGP holds, which by Conjecture 3.2 holds modulo exponentially small probability, this implies $A(V^{\tau_{n^*}})/N \geq \mathbb{E}[\eta_N] + \mu$ and therefore the larger event

$$\max_{0 \leq n \leq N^\alpha} A(V^{\tau_n})/N \geq \mathbb{E}[\eta_N] + \mu.$$

However, this contradicts assumption (16) and the concentration bound of Theorem 6.1 applied in the union over $n = 0, \ldots, N^\alpha$ bound. This yields the result conditionally on U^0. Since the lower bound we obtain does not depend on U^0, we can take the expectation in U^0, and obtain the main result.

8 TAP-type iteration schemes

One motivations for the AMP algorithm discussed in the introduction is the prediction that the minimizers of $A(u)$ satisfy a self-consistent equation, called a "mean-field" equation. In this setting, equations of this type are called Thouless–Anderson–Palmer (TAP) equations, after the work of those three authors in [TAP77], on mean-field equations in the case $p = 2$ on $\mathcal{B}_N$, in a certain physically motivated relaxation. For a discussion of these and related results see also [MPV87]. In this section, we show that, as an implication of Theorem 3.3, the iterative methods designed to produce solutions to TAP-like equations fail, modulo Conjecture3.2.

More precisely, consider the following modification of the objective. Recall the Bernoulli entropy, $S : [-1, 1] \to \mathbb{R}_+$

$$S(x) = \frac{1}{2}(1 + x)\log(1 + x) + \frac{1}{2}(1 - x)\log(1 - x).$$

For any $\beta > 0$, let $f_\beta : [-1, 1] \to \mathbb{R}$

$$f_\beta(x) = \frac{\beta^2}{2}(1 - x^p - px^{p-1}(1 - x)),$$

Finally, define the one-parameter family of functions $F_\beta : \mathcal{H}_N \to \mathbb{R}^N$ given by

$$F_\beta(x) = \beta A(x) - S(x) + f_\beta(\frac{\|x\|^2}{N}).$$

Observe that as $\beta \to \infty$,

$$\frac{F_\beta(u)}{\beta} \to \begin{cases} A(x) & x \in \mathcal{B}_N, \\ \infty & x \in \mathcal{H}_N \setminus \mathcal{B}_N. \end{cases}$$

Thus one expects that for β very large, minimizers of F_β are near minimizers for $A(u)$. In particular, one approach to computing near minimizers for (2) would be to produce minimizers of F.

Differentiating F_β, we see that the critical points of F satisfy the fixed point equation

$$x = \tanh\left[\beta\nabla A(x) + 2f'\left(\frac{\|x\|^2}{N}\right)x\right]. \tag{18}$$

Thus one approach to produce these minimizers is to construct solutions of these fixed point equations. The AMP algorithm is one such method, based off of a deep intuition in the physics literature that suggests that in the case $p = 2$.

Another approach, would be a more naive approach, in the spirit of standard AMP iterations would be to simply iteratively construct solutions to (18) as in [Bol14]. It is expected that the critical points of this equation satisfy,

$$\frac{\|x\|^2}{N} = q_*(\beta),$$

where $q_*(\beta)$ is an explicit constant, called the Edwards-Anderson order parameter and is given by as in [Mon18]. Motivated by this, consider the following class of AMP iterations,

$$U^t = \tanh\left(\beta A(\cdot, U^{t-1}) + a_{t-2}U^{t-2}\right), \quad U_0 = 1_N q, \ U_{-1} = 0 \quad 1 \le t \le T.$$

where $q > 0$ is a fixed constant and a_t is any bounded sequence. For instance, we may take $a_t = 2f'(q_*)$ and $q = q_*$. One might make the replacement $a_t \mapsto f'(\|U^t\|_2^2/N)$, however, one can show that this will not change the performance if the original sequence was chosen appropriately. See Section 9 for a similar argument in the more detailed case of the AMP iteration from [Mon18].

As a consequence of OGP, we see that the above iteration will fail to produced fixed points which are also near optimizers of $A(u)$ for β large. More precisely we have the following.

Corollary 8.1. *Suppose that the entries of $A \in (\mathbb{R}^N)^{\otimes p}$ are i.i.d. $\mathcal{N}(0, N^{-p+1})$ and $p \geq 4$ is even. Assuming the validity of Conjecture 3.2, there exists a $\bar{\mu} > 0$ such that for any $M, T > 0$ and β sufficiently large, if $V = V(A)$ is the result of the firs two steps of the AMP algorithm after T iterations, then*

$$\frac{1}{\beta}F_\beta(V) \geq \frac{\min_{x \in \mathcal{H}_N} F(x)}{N} + \bar{\mu}.$$

Proof. First note that this iteration is of the form (8), for some functions F_t, f_t satisfying Assumption 2.2. Indeed, let

$$f_t(u_1, \ldots, u_t) = u_t$$
$$F_t(u_0, \ldots, u_t) = \tanh(\beta u_0 + a_{t-2} u_{t-2}).$$

These functions are Lipschitz on the relevant domains as $\tanh(x)$ is smooth with bounded derivatives. Thus by Theorem 3.3,

$$A(V) \geq \min_{x \in \mathcal{B}_N} A(x)/N + \bar{\mu}.$$

Now observe that $F_\beta(x)$ satisfies $F_\beta(x) \geq \beta A(x) - \log(2)$ on $\mathcal{H}_N$. In particular, this is an equality on $\mathcal{B}_N$. As a result,

$$\frac{1}{N\beta}F_\beta(V) \geq \frac{A(V)}{N} - \frac{\log 2}{\beta N} \geq \min_{x \in \mathcal{B}_N} \frac{A(x)}{N} + \bar{\mu}$$
$$\geq \min_{x \in \mathcal{B}_N} \frac{F_\beta(x)}{\beta N} + \bar{\mu} + \frac{\log(2)}{2\beta} \geq \min_{x \in \mathcal{H}_N} \frac{F_\beta(x)}{\beta N} + \bar{\mu}/2,$$

where in the last line we take $\beta > \frac{\log(2)}{2\bar{\mu}}$ by assumption. $\qquad\square$

9 Verification for AMP for p-spin models

In this section we show that the AMP algorithm defined in [Mon18] is a special case of the AMP defined in Section 2, modulo some truncation and averaging steps which we discuss below. Here $p = 2$ so $A \in \mathbb{R}^{N \times N}$ is a matrix. The algorithm constructed in [Mon18] is as follows. A one-dimensional measure μ is fixed which is a solution of the minimization problem of the Parisi functional. A function $\Phi : [0,1] \times \mathbb{R} \to \mathbb{R}$ is a solution of the associated PDF. It is known that $\partial_x \Phi(t, x)$ and $\partial_{xx} \Phi(t, x)$ of this function are Lipschitz continuous. A certain value $q_* \in [0,1]$ is fixed (it is Edward-Anderson parameter). Let $u^{-1} = 0 \in \mathbb{R}^N, u^0$ be i.i.d. standard normal vector in $\mathbb{R}^N$: $u^0 \stackrel{d}{=} \mathcal{N}(0, I_N)$, $g^{-2} = 0 \in \mathbb{R}^N, g^{-1} = 1_N \in \mathbb{R}^N, b^0 = 0 \in \mathbb{R}^N$. Given $a, b \in \mathbb{R}^N$, $a \cdot b \in \mathbb{R}^N$ denotes a coordinate-wise product of a and b. Then for $t = 0, 1, \ldots, \lfloor q_*/\delta \rfloor \triangleq T$,

$$u^{t+1} = A(g^{t-1} \cdot u^t) - b^t g^{t-2} \cdot u^{t-1}, \tag{19}$$
$$x^t = x^{t-1} + \beta^2 \mu(t\delta)\partial_x \Phi(t\delta, x^{t-1})\delta + \beta\sqrt{\delta}u^t, \tag{20}$$
$$g^t = \sqrt{N}\partial_{xx}\Phi(t\delta, x^t)/\|\partial_{xx}\Phi(t\delta, x^t)\|_2, \tag{21}$$
$$b^t = N^{-1} \sum_{1 \leq i \leq N} g_i^t, \tag{22}$$

where everywhere the functions are applied coordinate-wise.

We first consider modifications of these iterations and justify them. First set M to be a large constant. Since $u^0 \stackrel{d}{=} \mathcal{N}(0, I_N)$, then the fraction of coordinates of u^0 with absolute values larger than N decreases to zero as a function of M. Replace (19) by

$$u^{t+1} = [A(g^{t-1} \cdot u^t) - b_t g^{t-2} \cdot u^{t-1}]_M. \tag{23}$$

In the final step of algorithm in [Mon18], the resulting vectors $u^1, \ldots, u^t$ are used to construct

$$z = \sqrt{\delta} \sum_{1 \leq t \leq \lfloor q_* / \delta \rfloor} g^{t-1} \cdot u^t,$$

and then z is rounded to a vector in $\mathcal{H}_N$ via $\max(-1, \min(1, \cdot))$ operator. It is thus expected that the truncation of u^t in (23) by a value M does not affect the result significantly provided M is large, though we do not prove this fact.

Next, as an implication of the analysis in [Mon18], as $N \to \infty$, the norm $\|\Phi(t\delta, x^t)\|_2$ is concentrated around a deterministic function of t, which which has value $\Theta(\sqrt{N})$. In particular, there it is argued that the empirical measure $\mathcal{E}_N^t = \frac{1}{N} \sum \delta_{x_i^t}$ converges to some deterministic limit $\mathcal{E}^t$ weakly almost surely. Since $\partial_{xx}\Phi$ is smooth and bounded [JT16, Theorem 4], it then follows that

$$\|\partial_{xx}\Phi(t\delta, x^t)\|/\sqrt{N} = \sqrt{\int \partial_{xx}\Phi(t\delta, y) d\mathcal{E}_N^t(y)} \to h(t).$$

Denoting this function by $\sqrt{N}h(t), t = 0, 1, \ldots, T$, we thus rewrite (21) as

$$g^t = h^{-1}(t)\partial_{xx}\Phi(t\delta, x^t). \tag{24}$$

Similarly, b^t, which per (22) is defined as coordinate-wise average of g^t, as $N \to \infty$ is concentrated around a deterministic function of t, which we denote by $\eta(t), t = 0, 1, \ldots, T$. Thus we replace (22) by

$$b^t = \eta^t.$$

We now fit these iterations into our framework defined in Section 2. We begin by defining f_t. In light of (24) replacing (21) we may define $f_t : \mathbb{R}^{t+1} \to \mathbb{R}$ as follows

$$f_t(u^0, \ldots, u^t) = g^{t-1}(u^0, \ldots, u^{t-1})u^t, \tag{25}$$

where the function $g_t : \mathbb{R}^t \to \mathbb{R}$ is a one-dimensional version of g^t, namely

$$g_t(u^0, \ldots, u^t) = h^{-1}(t)\partial_{xx}\Phi(t\delta, x^t),$$

where x^t is defined through a one-dimensional version of (20):

$$x^t = x^{t-1} + \beta^2 \mu(t\delta)\partial_x\Phi(t\delta, x^{t-1})\delta + \beta\sqrt{\delta}u^t.$$

Since $\partial_{xx}\Phi$ and $\partial_{xx}\Phi$ are Lipschitz continuous in the second argument for each fixed first argument, it is then immediate to verify that $g_t : \mathbb{R}^t \to \mathbb{R}$ is Lipschitz continuous as a function of $u^0, \ldots, u^t$, wrt $\|\cdot\|_\infty$ norm, say with constant C_t and satisfies $g_t(0) = 0$. This implies $\|g_t(x)\|_\infty \leq C_t\|x\|_\infty$. Then by (25) we have for every $u^0, \ldots, u^t$ and $v^0, \ldots, v^t$

$$|f_t(u^0, \ldots, u^t) - f_t(v^0, \ldots, v^t)| = |g^{t-1}(u^0, \ldots, u^{t-1})u_t - g^{t-1}(v^0, \ldots, v^{t-1})v^t|$$
$$\leq C_t \max_j(\max(|u^j|, |v^j|)\|(u^0, \ldots, u^t) - (v^0, \ldots, v^t)\|_\infty.$$

Thus f_t is Lipschitz continuous on $[-M, M]^{t+1}$, and thus satisfies the first part of Assumption 2.2.

Now motivated by (19) or rather (23) we define $F_t : \mathbb{R}^{t+1} \to \mathbb{R}$ by

$$F_t(y, u^0, \ldots, u^t) = y - b_t g^{t-2} \cdot u^{t-1}$$
$$= y - \eta(t)h^{-1}(t-2)\partial_{xx}\Phi((t-2)\delta, x_t) \cdot u^{t-1}$$

We see that F_t is Lipschitz continuous on $\mathbb{R} \times [-M, M]^{t+1}$ and thus satisfies the second part of Assumption 2.2.

20

10 Some open questions

We now list some questions which remain open. First validating Conjectures 3.2 and 3.6 is of interest. It is conceivable that the first of these conjectures can be approached by analyzing the Parisi measure directly on the Hilbert cube $\mathcal{H}_N$, as opposed to the binary cube $\mathcal{B}_N$. Carrying out the corresponding technical analysis of the associated variational problem could be quite daunting though. Another interesting question left open in this work is establishing the negative result for the binary output $\Pi(V)$ of the AMP scheme (Step 3) directly, as opposed to one for the pen-ultimate state V. Lifting the truncation $[\cdot]_M$ assumption adopted by our class of AMP algorithms is another question which remains open.

Next, it would be interesting to extend the main result of this paper to the p-spin spherical spin glass model and complement the positive result of Subag [Sub18]. We expect that our negative result extends to this model within the same scope of algorithms almost verbatim. Furthermore, by similar arguments one can show that (continuous time) gradient flow started from a uniform at random point fails to reach near minimizers of the spherical p-spin model in $\log(N)$ time with probability tending to 1. That being said, it is important to note that Subag's algorithm is based on an iterative sequence of computations which involve linear projections of gradient and Hessians of $A(u)$ to the linear space orthogonal to u. As such this computational scheme does not formally fit our framework of algorithms. It is conceivable though that the projection step can be approximated well by iterations of the form we consider, say perhaps by imitating the power iteration approach for spectral computations. Perhaps as an easier challenge, one could try to show that Subag's scheme specifically fails to find near ground states in models exhibiting OGP. A related question is whether there exists a connection between the OGP and the algorithmic hardness of the REM model discussed in [ABM18].

Our approach was formulated in terms of bounded (N-independent) number of iterations T. It is easy to see though that the proof method extends without a change to the case $T \leq c \log N$ for small enough constant c. At the same time we believe that AMP scheme is not effective in computing near ground states, regardless of the scale of the number of iterations. Thus an interesting open question is to see whether an AMP scheme achieving near ground states can be designed say when $T = N^{O(1)}$.

Finally, perhaps the most intriguing question which remains open is one regarding the genuine hardness of the problem of finding ground states in models exhibiting the OGP. While formal hardness of problems associated with spin glass models is known, in particular it is shown in [GK18] that computing the partition function of the p-spin models is hard on average even in $p = 2$ regime, these results are established using more "standard" average case hardness proof approaches, and do not take advantage of the intricate solution space topology, such as the one expressed by OGP. At the same time, as of now we have very compelling consistence of the presence of OGP and the apparent hardness of the associated optimization problem in many models. What is lacking, however, is the formal link between the two within a class of algorithms which is broader than AMP. An interesting and challenging conjecture is that the OGP implies formal average case hardness of the underlying optimization problem, perhaps even within the class of all polynomial time algorithms.

Acknowledgements

References

[ABM18]   Louigi Addario-Berry and Pascal Maillard, *The algorithmic hardness threshold for continuous random energy models*, arXiv preprint arXiv:1810.05129 (2018).

[AC15]   Antonio Auffinger and Wei-Kuo Chen, *On properties of Parisi measures*, Probab. Theory Related Fields **161** (2015), no. 3-4, 817–850. MR 3334282

[ACO08]   Dimitris Achlioptas and Amin Coja-Oghlan, *Algorithmic barriers from phase transitions*, Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on, IEEE, 2008, pp. 793–802.

[ACORT11] D. Achlioptas, A. Coja-Oghlan, and F. Ricci-Tersenghi, *On the solution space geometry of random formulas*, Random Structures and Algorithms **38** (2011), 251–268.

[ACZ17]   Antonio Auffinger, Wei-Kuo Chen, and Qiang Zeng, *The sk model is full-step replica symmetry breaking at zero temperature*, arXiv preprint arXiv:1703.06872 (2017).

[AT09]   Robert J Adler and Jonathan E Taylor, *Random fields and geometry*, Springer Science & Business Media, 2009.

[BAGJ18a] Gerard Ben Arous, Reza Gheissari, and Aukosh Jagannath, *Algorithmic thresholds for tensor pca*, arXiv preprint arXiv:1808.00921 (2018).

[BAGJ18b] ———, *Bounding flows for spherical spin glass dynamics*, arXiv preprint arXiv:1808.00929 (2018).

[BAJ18]   Gérard Ben Arous and Aukosh Jagannath, *Spectral gap estimates in mean field spin glasses*, Comm. Math. Phys. **361** (2018), no. 1, 1–52.

[BLM$^+$15] Mohsen Bayati, Marc Lelarge, Andrea Montanari, et al., *Universality in polytope phase transitions and message passing algorithms*, The Annals of Applied Probability **25** (2015), no. 2, 753–822.

[BM11]   Mohsen Bayati and Andrea Montanari, *The dynamics of message passing on dense graphs, with applications to compressed sensing*, IEEE Transactions on Information Theory **57** (2011), no. 2, 764–785.

[BMN17]   Raphael Berthier, Andrea Montanari, and Phan-Minh Nguyen, *State evolution for approximate message passing with non-separable functions*, arXiv preprint arXiv:1708.03950 (2017).

[Bol14]   Erwin Bolthausen, *An iterative construction of solutions of the tap equations for the sherrington–kirkpatrick model*, Communications in Mathematical Physics **325** (2014), no. 1, 333–366.

[CGP$^+$19] Wei-Kuo Chen, David Gamarnik, Dmitry Panchenko, Mustazee Rahman, et al., *Suboptimality of local algorithms for a class of max-cut problems*, The Annals of Probability **47** (2019), no. 3, 1587–1618.

[CHL18]   Wei-Kuo Chen, Madeline Handschy, and Gilad Lerman, *On the energy landscape of the mixed even p-spin model*, Probability Theory and Related Fields, **171** (2018) no.1-2, 53-95

[COE11]   A. Coja-Oghlan and C. Efthymiou, *On independent sets in random graphs*, Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2011, pp. 136–144.

[COHH17]   Amin Coja-Oghlan, Amir Haqshenas, and Samuel Hetterich, *Walksat stalls well below satisfiability*, SIAM Journal on Discrete Mathematics **31** (2017), no. 2, 1160–1173.

[CS17]   Wei-Kuo Chen and Arnab Sen, *Parisi formula, disorder chaos and fluctuation for the ground state energy in the spherical mixed p-spin models*, Comm. Math. Phys. **350** (2017), no. 1, 129–173. MR 3606472

[DI17]   Gamarnik David and Zadik Ilias, *High dimensional regression with binary coefficients. estimating squared error and a phase transtition*, Conference on Learning Theory, 2017, pp. 948–953.

[DMM09]   David L Donoho, Arian Maleki, and Andrea Montanari, *Message-passing algorithms for compressed sensing*, Proceedings of the National Academy of Sciences **106** (2009), no. 45, 18914–18919.

[GJS19]   David Gamarnik, Aukosh Jagannath, and Subhabrata Sen, *The overlap gap property in principal submatrix recovery*, arXiv preprint arXiv:1908.09959 (2019).

[GK18]   David Gamarnik and Eren Kizildag, *Computing the partition function of the sherrington-kirkpatrick model is hard on average*, arXiv preprint arXiv:1810.05907 (2018).

[GL18]   David Gamarnik and Quan Li, *Finding a large submatrix of a gaussian random matrix*, The Annals of Statistics **46** (2018), no. 6A, 2511–2561.

[GS17a]   David Gamarnik and Madhu Sudan, *Limits of local algorithms over sparse random graphs*, Annals of Probability **45** (2017), 2353–2376.

[GS17b]   ——, *Performance of sequential local algorithms for the random nae-k-sat problem*, SIAM Journal on Computing **46** (2017), no. 2, 590–619.

[JM13]   Adel Javanmard and Andrea Montanari, *State evolution for general approximate message passing algorithms, with applications to spatial coupling*, Information and Inference: A Journal of the IMA **2** (2013), no. 2, 115–144.

[JT16]   Aukosh Jagannath and Ian Tobasco, *A dynamic programming approach to the parisi functional*, Proceedings of the American Mathematical Society **144** (2016), no. 7, 3135–3150.

[JT17a]   ——, *Low Temperature Asymptotics of Spherical Mean Field Spin Glasses*, Comm. Math. Phys. **352** (2017), no. 3, 979–1017. MR 3631397

[JT17b]   ——, *Some properties of the phase diagram for mixed p-spin glasses*, Probability Theory and Related Fields **167** (2017), no. 3-4, 615–672.

[JT18]   ——, *Bounds on the complexity of replica symmetry breaking for spherical spin glasses*, Proceedings of the American Mathematical Society **146** (2018), no. 7, 3127–3142.

[Kab03]   Yoshiyuki Kabashima, *A cdma multiuser detection algorithm on the basis of belief propagation*, Journal of Physics A: Mathematical and General **36** (2003), no. 43, 11111.

[MMZ05] M. Mézard, T. Mora, and R. Zecchina, *Clustering of solutions in the random satisfiability problem*, Physical Review Letters **94** (2005), no. 19, 197205.

[Mon18] Andrea Montanari, *Optimization of the sherrington-kirkpatrick hamiltonian*, arXiv preprint arXiv:1812.10897 (2018).

[MPV87] M. Mezard, G. Parisi, and M. A. Virasoro, *Spin-glass theory and beyond,* vol 9 of *Lecture Notes in Physics*, World Scientific, Singapore, 1987.

[MPZ02] Marc Mézard, Giorgio Parisi, and Riccardo Zecchina, *Analytic and algorithmic solution of random satisfiability problems*, Science **297** (2002), no. 5582, 812–815.

[Pan13] Dmitry Panchenko, *The sherrington-kirkpatrick model*, Springer Science & Business Media, 2013.

[RV+17] Mustazee Rahman, Balint Virag, et al., *Local algorithms for independent sets are half-optimal*, The Annals of Probability **45** (2017), no. 3, 1543–1577.

[Sub18] Eliran Subag, *Following the ground-states of full-rsb spherical spin glasses*, arXiv preprint arXiv:1812.04588 (2018).

[Tal06] Michel Talagrand, *Free energy of the spherical mean field model*, Probab. Theory Related Fields **134** (2006), no. 3, 339–382. MR 2226885 (2007i:82034)

[TAP77] David J Thouless, Philip W Anderson, and Robert G Palmer, *Solution of 'solvable model of a spin glass'*, Philosophical Magazine **35** (1977), no. 3, 593–601.

[Tes12] Gerald Teschl, *Ordinary differential equations and dynamical systems*, vol. 140, American Mathematical Soc., 2012.

[Ver18] Roman Vershynin, *High-dimensional probability: An introduction with applications in data science*, vol. 47, Cambridge University Press, 2018.