

Random multi-index matching problems

O C Martin, M Mézard and O Rivoire

Laboratoire de Physique Théorique et Modèles Statistiques, Université
Paris-Sud, Bâtiment 100, 91405 Orsay Cedex, France
E-mail: martino@lptms.u-psud.fr, mezard@lptms.u-psud.fr and
rivoire@lptms.u-psud.fr

Received 8 July 2005

Accepted 2 September 2005

Published 20 September 2005

Online at stacks.iop.org/JSTAT/2005/P09006

[doi:10.1088/1742-5468/2005/09/P09006](https://doi.org/10.1088/1742-5468/2005/09/P09006)

Abstract. The multi-index matching problem generalizes the well known matching problem by going from pairs to d -uplets. We use the cavity method from statistical physics to analyse its properties when the costs of the d -uplets are random. At low temperatures we find for $d \geq 3$ a frozen glassy phase with vanishing entropy. We also investigate some properties of small samples by enumerating the lowest cost matchings to compare with our theoretical predictions.

Keywords: cavity and replica method, typical-case computational complexity

ArXiv ePrint: [cond-mat/0507180](https://arxiv.org/abs/cond-mat/0507180)

Contents

1. Introduction	3
2. Multi-index matchings	4
2.1. Definitions	4
2.2. Physical approach	5
2.3. Mathematical approach	6
2.4. Computer science approach	6
3. Scaling and a lower bound	7
3.1. Scaling	7
3.2. Annealed approximation	8
3.3. Statistical physics reformulation	9
4. Replica symmetric solution	9
4.1. From complete to dilute graphs	9
4.2. Cavity method	11
4.3. Integral relations	14
4.4. Zero-temperature limit	15
4.5. Entropy crisis	17
4.6. Stability of the replica symmetric ansatz	18
5. Replica symmetry breaking	20
5.1. General 1RSB ansatz	20
5.2. Frozen 1RSB ansatz	22
5.3. Distances	23
6. Numerical analysis of finite size systems	24
6.1. The branch and bound procedure	25
6.2. Ground state energies	27
6.3. Other ground state properties	29
6.4. Excited states	30
6.5. Low energy entropy	31
7. Conclusion	31
Acknowledgments	33
Appendix A: Population dynamics algorithm	33
Appendix B: Aspects of the branch and bound algorithm	34
References	36

1. Introduction

The statistical properties of random combinatorial optimization problems can be studied from a number of angles, with tools depending on the discipline. Recent years have however witnessed a convergence of interests and techniques across mathematics, computer science and statistical physics. An archetypal example is the matching problem with random edge weights, defined as follows: suppose one has M different jobs and M people to perform them, one person per job, and let c_{ij} be the cost when job i is executed by person j ; the 2-index matching problem consists in assigning jobs to people in such a way as to minimize the total cost. The statistical properties of the optimal matching when the costs c_{ij} are drawn independently from a common distribution were found two decades ago using the replica [1] and the cavity [2] methods. These two non-rigorous statistical physics approaches have recently been used to tackle a number of computationally more difficult problems such as satisfiability and graph colouring, but the 2-index matching problem is set apart by being one of the very few problems for which such predictions have been rigorously confirmed [3].

In this work, we take the statistical physics approach and study the properties of a generalization of the 2-index to multi-index matching problems (MIMPs) where the elementary costs are now associated with d -uplets, representing for example persons, jobs and machines when $d = 3$. Unlike the 2-index matching, d -index matching problems with $d \geq 3$ are NP-hard. We show here that their low lying configurations also have a different, glassy, structure whose description requires replica symmetry to be broken. Remarkably, the replica symmetry breaking scheme differs from the common picture that has emerged from the study of other optimization problems such as the colouring [4] and satisfiability problems [5]. In particular, a naive application of the 1-RSB cavity method at zero temperature [6], which successfully solves these two problems, is here doomed to fail. The reason for this will be traced back to the presence of ‘hard constraints’. By unravelling this specificity, we put forward arguments whose relevance goes beyond matching problems; they indicate when a similar scenario can be expected for other constrained systems. The particularly simple glassy structure that we find is also of interest from the interdisciplinary point of view: in conjunction with the rigorous formalism available for the 2-index case, it places MIMPs in an excellent position for providing the highly desirable mathematical understanding of replica symmetry breaking.

The present paper provides an extensive account of our results on the MIMPs, some of which have already been mentioned in [7]. The paper is organized as follows. We first define precisely multi-index matching problems, and briefly review the past approaches from physics, mathematics and computer science that were developed mainly to address the 2-index case. Then we start our statistical study by establishing the scaling of the minimal cost as a function of the number of variables and by providing a lower bound from an annealed calculation. A large part of the paper is then devoted to presenting our implementation of the cavity method for matching problems, including a detailed discussion of its relations with the rigorous formalism proposed by Aldous; we explain why and how replica symmetry must be broken when $d \geq 3$, in order to account for the presence of a frozen glassy phase. Finally, the last section is dedicated to a numerical analysis of small samples that provides support to the proposed scenario.

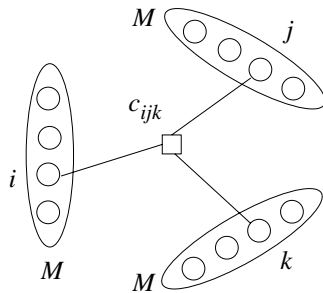


Figure 1. *Tripartite matching problem.* Factor graph representation: the hyperedges, or factor nodes, are represented with squares.

2. Multi-index matchings

2.1. Definitions

Two classes of MIMPs can be distinguished, d -partite matching problems and simple d -matching problems, whose asymptotic properties will be shown to be related. We first start with the d -partite matching problem that corresponds to the version alluded to in the introduction. An instance consists of d sets, A_1, \dots, A_d , of M nodes each, and a cost c_a is associated with every d -uplet $a = \{i_1, \dots, i_d\} \in A_1 \times \dots \times A_d$. Graphically, it is represented by a factor graph as shown in figure 1 with hyperedges (factor nodes) joining exactly one node from each ensemble. A matching \mathcal{M} is a maximal set of disjoint hyperedges, such that each node is associated with one and only one hyperedge of the matching; it can be described by introducing an occupation number $n_a \in \{0, 1\}$ on each hyperedge a , with the correspondence

$$a \in \mathcal{M} \Leftrightarrow n_a = 1. \tag{1}$$

The condition for a set of hyperedges to be a matching can then be written as

$$\forall r = 1, \dots, d, \quad \forall i_r \in A_r, \quad \sum_{a : i_r \in a} n_a = 1. \tag{2}$$

The d -partite matching problem consists in finding the matching with minimal total cost,

$$C_M^{(d)} = \min_{\{n_a\}} \sum_a c_a n_a, \tag{3}$$

with the $\{n_a\}$ subject to the constraints (2). We consider here the random version of the problem, where the costs c_a are independent identically distributed random variables taken from a distribution $\rho(c)$, and we are interested in the typical value of an optimal matching in the $M \rightarrow \infty$ limit. For definiteness, we take for ρ the uniform distribution in $[0, 1]$, but the asymptotic properties of d -matchings depend only on the behaviour of ρ close to $c = 0$, and are identical for all distributions ρ with $\rho(c) \sim 1$ as $c \rightarrow 0$, such as the exponential distribution, $\rho(c) = e^{-c}$. The case $\rho(c) \sim c^r$, $r > 0$, can be treated along the same lines, but gives different quantitative results.

A variant of this set-up is the *simple* d -matching problem, where a unique set of N nodes, with N being a multiple of d , is considered and a cost is associated with each

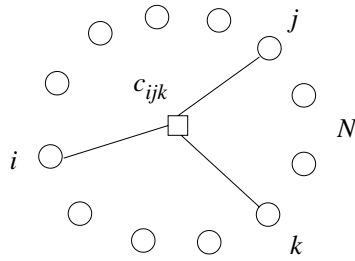


Figure 2. Simple 3-index matching problem. The factor graph representation is similar to the tripartite case.

d -uplet of nodes (see figure 2). The d -partite case can be seen as a particular instance of a simple d -matching problem where the hyperedges joining more than one node of any A_i are given an infinite cost. Simple d -matchings problems are formulated as finding

$$L_N^{(d)} = \min_{\{n_a\}} \sum_a c_a n_a \tag{4}$$

under the constraints

$$\forall i = 1, \dots, N, \quad \sum_{a : i \in a} n_a = 1. \tag{5}$$

Before presenting our analysis of random matching problems by means of an adaptation of the cavity method for finite connectivity statistical physics models, we briefly review past approaches to the subject, with an emphasis on open questions that motivated the present study.

2.2. Physical approach

The 2-index matching problem was the first combinatorial optimization problem to be tackled with the replica method, an analytical method initially developed in the context of spin glasses [8]. In the paper [1], Mézard and Parisi analysed both the simple and bipartite matching problems for cost distributions ρ with $\rho(c) \sim c^r$ as $c \rightarrow 0$. Using replica theory within a replica symmetric ansatz, they derived the minimal total cost; thus, for the bipartite matching with $r = 0$, they predicted $\lim_{M \rightarrow \infty} \mathcal{C}_M^{(2)} = \pi^2/6$; moreover, they obtained the distribution of cost in the optimal matching. Support in favour of their prediction first came from numerical results and from an analytical study of the stability of the replica symmetric solution [9, 10]. This last analysis further yields the leading corrections of order $1/N$ for the value of the minimum matching.

Interestingly, the same results can be reobtained using a variant of the cavity method based on a representation of self-avoiding walks using m -component spins [2]. This alternative formulation, avoiding the bold prescriptions of replica theory, furthermore suggests that, if the costs of the hyperedges connected to a given node are ordered from the lowest to the highest, the probability for the k th hyperedge to be included in the optimal matching is 2^{-k} [11], as first conjectured from a numerical study [12].

2.3. Mathematical approach

Replica theory, while a powerful tool for obtaining analytical formulae, is not a rigorously controlled method, and its predictions have only the status of conjectures within the usual mathematical standards. For the 2-index matching problem with $r = 0$ however, the results mentioned above (value of the optimal matching, distribution of costs and probability of inclusion of the k th hyperedge) have all been confirmed by a rigorous derivation, due to Aldous [3]. His contribution also includes the proof an *asymptotic essential uniqueness* property that mathematically expresses the fact that replica symmetry indeed holds for this problem. The weak convergence approach [13] on which the proof is built is closely related to the cavity method that we will employ, and the relations between the two formalisms will be discussed in section 4.4. Confirmation of the $\zeta(2) = \pi^2/6$ value for the bipartite assignment problem also comes from the recent proofs [14, 15] of a more general conjecture formulated by Parisi [16]; this conjecture states that, for the bipartite matching with exponential distribution of the costs, $\rho(c) = e^{-c}$, the mean optimal matching for *finite* M is $\sum_{k=1}^M k^{-2}$.

These mathematical contributions are part of a more ambitious programme aiming at developing rigorous proofs and possibly a rigorous framework for the replica and cavity methods. Interestingly, Talagrand, one prominent advocate of this programme, devotes the last chapter of his book on the subject [17] to the 2-index matching problem, stressing that, in spite of the major advances mentioned, it remains a particularly challenging issue. Indeed, finite temperature properties have so far resisted mathematical investigations, even in the limit of high temperature, which has been successfully addressed in other spin-glass-like models [17]. We shall comment on the peculiarities of matchings with respect to other constrained systems in section 5.2. It is our hope that our work not only provides new challenging conjectures, but also suggests some hints for solving unanswered pre-existing mathematical questions.

2.4. Computer science approach

If analytical studies of random d -matchings by statistical physicists and mathematicians have been restricted up to now to the $d = 2$ case, d -index matching problems with $d > 2$ have a longer history in the computer science community. d -partite extensions of the bipartite matching problem were introduced in 1968 under the name of *multidimensional assignment problems* [18]; they are also referred to in the literature as *multi-index assignment problems* and, more specifically, as multi-index *axial* assignment problems (to distinguish them from the so-called *planar* versions [19, 20]). MIMPs, as we call them (for multi-index matching problems), have a number of practical applications. The most commonly cited one is for data association in connection with multi-target tracking [21]. Besides major interest as regards real-time air traffic control, such approaches are for instance helpful for tracking elementary particles in high energy physics experiments [22].

From the algorithmic complexity point of view, matching problems also have a pioneering role since the 3-index matching problem was among the first 21 problems to be proved NP-complete [23]. In contrast, polynomial algorithms are known that solve 2-index matching problems [24]. Note that being based on a worst case analysis, NP-hardness is however only a necessary condition for hard typical complexity, which is the issue which interests us here. Due to their intrinsic algorithmic difficulty and to the broad

range of their applications, generalized assignment problems are the subject of numerous studies in the computer science community; we refer the reader to the reviews [19, 20] for additional information and references.

3. Scaling and a lower bound

The first task in studying random optimization problems is to determine the scaling of the optimal cost with the number of variables [25]. Here, we address this issue for the two variants of MIMPs, the multi-partite and simple multi-index matching problems. The scaling is inferred from a heuristic argument, and confirmed by an annealed calculation (first-moment method) yielding a lower bound. This leads us to a statistical physics formulation that encompasses the two versions of MIMPs.

3.1. Scaling

The statistical physics approach of combinatorial optimization problems consists in defining the energy $E(\mathcal{M})$ of each admissible solution, here a d -matching \mathcal{M} , as its total cost, $E(\mathcal{M}) = \sum_{a \in \mathcal{M}} c_a$, and in determining the minimal total cost, identified with the ground state energy, by looking at the zero temperature properties of the system. For d -matchings, the corresponding Hamiltonian

$$\mathcal{H}[\{n_a\}] = \sum_a c_a n_a \quad (6)$$

defines a lattice gas model, where the particles are occupying the hyperedges. The constraints (2) or (5) implement a hard-core interaction between the particles: two ‘neighbouring’ hyperedges are not allowed to be occupied simultaneously. To have a sensible statistical physics model, the ground state has to be extensive, i.e., proportional to M in the d -partite case and to N in the simple case. We propose here a heuristic argument for determining how $\mathbb{E}[C_M^{(d)}]$ and $\mathbb{E}[L_N^{(d)}]$ scale with M and N respectively, where $\mathbb{E}[\cdot]$ represents the average over the different realizations of the costs. The central (local) quantity that monitors the scaling behaviour is the number of hyperedges to which a given node belongs, denoted as $W_\Lambda^{(d)}$ ($\Lambda = M$ or N). Indeed, with the costs uniformly distributed in $[0, 1]$, the lowest costs with which a node can be associated are of order $1/W_\Lambda^{(d)}$ and the optimal matching is expected to scale like $\Lambda/W_\Lambda^{(d)}$. Thus, for d -partite matchings, $W_M^{(d)} = M^{d-1}$ and $\mathbb{E}[C_M^{(d)}] \sim M^{2-d}$, while for simple d -matchings, $W_N^{(d)} = \binom{N-1}{d-1}$ and $\mathbb{E}[L_N^{(d)}] \sim (d-1)!N^{2-d}$. We will therefore be interested in computing the (finite) quantities

$$\begin{aligned} \mathcal{C}^{(d)} &= \lim_{M \rightarrow \infty} M^{d-2} \mathbb{E}[C_M^{(d)}], \\ \mathcal{L}^{(d)} &= \lim_{N \rightarrow \infty} \frac{N^{d-2}}{(d-1)!} \mathbb{E}[L_N^{(d)}]. \end{aligned} \quad (7)$$

The factor $(d-1)!$ in the second definition is meant to reflect the different number of hyperedges to which a given node can connect, in the d -partite and simple versions (this difference is absent when $d = 2$). With this convention we will find the equality $\mathcal{C}^{(d)} = d\mathcal{L}^{(d)}$, where the remaining d factor merely comes from the fact that the total number of nodes is N for simple d -matchings, but is dM for d -partite matchings.

3.2. Annealed approximation

When energies are extensive in the size N of the system, the equilibrium properties of a statistical physics model are entirely encoded in the partition function, $Z_N(\beta) = \sum_{\mathcal{M}} e^{-\beta E(\mathcal{M})}$, or, equivalently, in its logarithm, the free energy $F_N(\beta) \equiv -\log[Z_N(\beta)]/\beta$. The free energy depends on the realization of the elementary costs, but it is expected to be a self-averaging quantity, i.e., such that the free-energy density $f(\beta) = \lim_{N \rightarrow \infty} F_N(\beta)/N$ exists and is independent of the sample. The self-averaging property is proved for $d = 2$ [26], and we assume here that it holds for $d \geq 3$ as well. The value of the optimal matching is given by the ground state energy, obtained as $\lim_{\beta \rightarrow \infty} f(\beta)$, where the free energy is calculated by performing a *quenched* average of the partition function, $\mathbb{E}[\ln Z]$, with $\mathbb{E}[\cdot]$ referring to the average with respect to the realization of the elementary costs.

A much simpler calculation is the *annealed* average, $\ln \mathbb{E}[Z]$. Due to the concavity of the logarithm, it yields a lower bound on the correct quenched free energy, $f_{\text{an}}(\beta) \equiv -\ln \mathbb{E}[Z]/(N\beta) \leq -\mathbb{E}[\ln Z]/(N\beta) \equiv f(\beta)$. In fact, since the entropy $s(\beta) = \beta^2 \partial_\beta f(\beta)$ is necessarily positive for a system with discrete degrees of freedom, the free energy $f(\beta)$ must be an increasing function, and a tighter lower bound can be inferred for the ground state energy [25],

$$\lim_{\beta \rightarrow \infty} f(\beta) \geq \sup_{\beta > 0} f_{\text{an}}(\beta). \quad (8)$$

These considerations are made under the hypothesis that the energies, or equivalently the temperature β , are correctly scaled with N , so that $\lim_{\beta \rightarrow \infty} f(\beta)$ is indeed finite. Reciprocally, requiring the annealed free energy to be extensive provides us with the appropriate scaling of β . For d -index matching problems, we have

$$\mathbb{E}[Z] = \mathbb{E} \left[\sum_{\{n_a\}} \exp \left(-\beta \sum_a c_a n_a \right) \right] = (\#\mathcal{M}) \mathbb{E}[e^{-\beta c_a}]^{\#\{a \in \mathcal{M}\}} \quad (9)$$

where $\#\mathcal{M}$ denotes the total number of possible matchings and $\#\{a \in \mathcal{M}\}$ the number of hyperedges contained in a given matching. For d -partite matchings, $\#\mathcal{M} = (M!)^{d-1}$ and $\#\{a \in \mathcal{M}\} = M$. To enforce the correct scaling of the free energy, we anticipate a rescaling in temperature of the form $\beta = M^\alpha \hat{\beta}$, yielding

$$\ln \mathbb{E}[Z] = [d - 1 - \alpha] M \ln M - [\ln \hat{\beta} + d - 1] M + o(M). \quad (10)$$

An extensive annealed free energy is therefore obtained by taking $\alpha = d - 1$, in which case

$$f_{\text{an}}^{(d\text{-part})}(\hat{\beta}) = \frac{\ln \hat{\beta} + d - 1}{\hat{\beta}}. \quad (11)$$

The scaling $1/(M\hat{\beta}) \sim M^{2-d}$ that we obtain corresponds to one introduced in equation (7). For simple d -index matchings, $\#\mathcal{M} = N!/[(N/d)!(d!)^{N/d}]$ and $\#\{a \in \mathcal{M}\} = N/d$. Rescaling the temperature as $\beta = N^{d-1} \hat{\beta}$, we get

$$\ln \mathbb{E}[Z] = -[\ln \hat{\beta} + d - 1 - \ln(d-1)!] N/d + o(N). \quad (12)$$

To make contact with the d -partite case however, we adopt a slightly different scaling, $\beta = N^{d-1} \tilde{\beta}/(d-1)!$, so that

$$f_{\text{an}}^{(\text{simple})}(\tilde{\beta}) = \frac{\ln \tilde{\beta} + d - 1}{d\tilde{\beta}} = \frac{1}{d} f_{\text{an}}^{(d\text{-part})}(\hat{\beta} = \tilde{\beta}). \quad (13)$$

This annealed calculation illustrates the correspondence between the d -partite and simple d -matchings stated in the previous section. Apart for the trivial factor d , corresponding to the relation $N = dM$, the equality is obtained by normalizing $\hat{\beta}$ and $\tilde{\beta}$ differently, thereby accounting for the difference in the number of hyperedges a given node locally sees (extra factor $(d - 1)!$ in equation (7)). The annealed free energy is a concave function with a maximum for $\hat{\beta}_d^* = e^{2-d}$ so we get lower bounds $\mathcal{C}^{(d)} \geq e^{d-2}$ and $\mathcal{L}^{(d)} \geq e^{d-2}/d$.

3.3. Statistical physics reformulation

From now on, we will cease distinguishing between d -partite and simple d -matchings, and consider a unique statistical physics model that describes the two problems in a common framework. Our approach is indeed based on the cavity method [27] for which only the local properties at the level of each node are relevant, and we have seen that by making the appropriate scalings of β , we can match the local properties of both models. The Hamiltonian that we consider is

$$\mathcal{H}[\{n_a\}] = \sum_a \xi_a n_a, \quad (14)$$

with the $\xi_a \equiv M^{d-1}c_a$ uniformly distributed in $[0, M^{d-1}]$ for the d -partite case and $\xi_a \equiv N^{d-1}c_a/(d-1)!$ in $[0, N^{d-1}/(d-1)!]$ for the simple case. The (inverse) temperature, denoted by β for simplicity, will correspond to $\hat{\beta}$ for the d -partite case and $\tilde{\beta}$ for the simple case. The only remaining difference kept is the factor d between the two free energies, accounting for the relation $N = dM$. Unless explicitly stated, the formulae to be given hold for the simple version; to get the d -partite counterparts, one has consequently to multiply the intensive quantities by d .

4. Replica symmetric solution

The approach we adopt to treat the d -matching problems is the cavity method recently developed to solve statistical physics models defined on finite connectivity graphs [27]. This section explains the formalism of the replica symmetric solution for general d . While the correctness of the replica symmetric approach is a mathematical fact when $d = 2$, we show that it leads to some inconsistency when $d = 3$, requiring replica symmetry to be broken.

4.1. From complete to dilute graphs

The hypergraph on which an instance of the simple d -matching problem is defined is *complete*, in the sense that every possible hyperedge arises once and is given a random cost. However, the factor nodes with the smallest elementary costs are more likely to belong to the optimal matching; for instance the probability that the k th most costly hyperedge originating from a given node will be included in the optimal 2-matching is 2^{-k} [3]. This suggests that hyperedges with large costs can be ignored while retaining most of the structure relevant to the determination of the optimal matching. Eliminating hyperedges results in a *diluted* hypergraph, where each node is connected to only a restricted number of hyperedges. From this point of view, in spite of being defined on a complete graph, random matchings are effectively closer to statistical physics models defined on finite connectivity random graphs. In fact, such a feature already transpired

from the initial replica treatment [1] of 2-index matchings where all the multi-overlaps $Q_{a_1\dots a_p}$ were required, and not only the two-replica overlaps $Q_{a_1 a_2}$, like in usual Curie–Weiss mean field models of disordered systems [8].

To exploit the underlying diluted structure, one possible method is to introduce a cut-off C , suppress all nodes with rescaled cost $\xi_a > C$, solve the matching problem on the diluted hypergraph and finally send $C \rightarrow \infty$. The hypergraph obtained by this procedure is Poissonian: if ξ_1, ξ_2, \dots are the costs ordered in increasing sequence of the hyperedges connected to a given node, the probability for the connectivity to be k is

$$\begin{aligned}
 p_k &= \text{Prob}[\xi_1 < \dots < \xi_k < C < \xi_{k+1} < \dots] \\
 &= \binom{W_\Lambda^{(d)}}{k} \left(\frac{C}{W_\Lambda^{(d)}}\right)^k \left(1 - \frac{C}{W_\Lambda^{(d)}}\right)^{W_\Lambda^{(d)} - k} \rightarrow \frac{C^k}{k!} e^{-C},
 \end{aligned}
 \tag{15}$$

with $W_\Lambda^{(d)}$ giving the number of hyperedges to which a node is connected, as in section 3.2. Diluting the complete graph has a major drawback however: the diluted hypergraph typically does not allow any matching at all, since for instance there is always a finite probability e^{-C} that a given node is isolated.

To circumvent this problem, we come back to the model on the complete graph and start by weakening the constraints, allowing a node not to belong to a matching, at the expense of paying an extra cost. In more physical terms, we view a matching as the close-packing limit of a lattice gas model whose particles are subject to hard-core interactions: particles can occupy the hyperedges but two hyperedges connected through a node cannot both admit a particle. We introduce a *grand-canonical* Hamiltonian

$$\mathcal{H}_\mu[\{n_a\}] = \sum_a \xi_a n_a - d\mu \sum_a n_a = \sum_a (\xi_a - d\mu) n_a
 \tag{16}$$

where $d\mu$ is a chemical potential per hyperedge (μ per node). In the limit $\mu \rightarrow \infty$, the maximum number of hyperedges is occupied by a particle and we recover the matching problem. For finite μ however, the constraints reflecting the hard-core repulsion are

$$\forall i, \quad \sum_{a : i \in a} n_a \leq 1,
 \tag{17}$$

to be compared with the hard constraints of equation (5), recovered only in the $\mu \rightarrow \infty$ limit. Each value of μ defines an optimization problem whose minimum energy E_μ corresponds to a zero-temperature limit $\beta \rightarrow \infty$. The solution of the matching problem thus appears as the result of a double limit, $\beta \rightarrow \infty$ and $\mu \rightarrow \infty$. The point is that a diluted structure is now naturally associated with the system at finite μ . Indeed since \mathcal{H}_μ is minimized by taking $n_a = 0$ whenever $\xi_a > d\mu$, the ground state is unaffected if all hyperedges a with $\xi_a > d\mu$ are suppressed, yielding the Poissonian hypergraph considered above with $C = d\mu$. This construction allows us to formulate the initial MIMP as the limit $\mu \rightarrow \infty$ of optimization problems defined on Poissonian graphs with increasing mean connectivity $d\mu$. We will give in section 4.4 an alternative construction based on regular graphs.

J. Stat. Mech. (2005) P09006

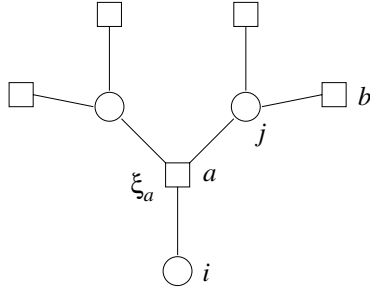


Figure 3. Local structure of a Poissonian hypergraph. When the node i is removed, it leaves a rooted tree with root a .

4.2. Cavity method

The problem at finite μ defined on a Poissonian hypergraph can be studied by means of the cavity method as developed for finite connectivity graphs [27]; one of the main advantages of this method over the replica method [1] or previous versions of the cavity method [2] is that it allows a practical investigation of replica symmetry breaking (RSB). Since we are interested in the ground state properties, the cavity method directly at zero temperature seems particularly well suited [28]. However, it will turn out to be necessary to get the finite temperature equations as well, and we therefore work at finite β , postponing the discussion of the $\beta \rightarrow \infty$ limit to the next section.

In strong analogy with Aldous’s framework (see section 4.4), the RS cavity method associates the diluted hypergraph with an *infinite tree* or, stated differently, a tree with *self-consistent boundary conditions*. The starting point is however finite rooted trees, that is trees with a singularized node i called the root. Consider for instance the part of a tree represented in figure 3: given a hyperedge a and one of its connected nodes i (relation noted $i \in a$), we call $Z^{(a \rightarrow i)}$ the partition function of the system defined on the rooted tree with root a resulting from the removal of i . To express it in terms of the partition functions $Z^{(b \rightarrow j)}$ where j refers to the nodes, connected to a , but distinct from i (noted $j \in a - i$), we decompose $Z^{(a \rightarrow i)}$ as $Z^{(a \rightarrow i)} = Z_0^{(a \rightarrow i)} + Z_1^{(a \rightarrow i)}$, where $Z_0^{(a \rightarrow i)}$ and $Z_1^{(a \rightarrow i)}$ are the conditional partition functions where the root a is either constrained to be empty or occupied by a particle. As an intermediate stage in the recursion, we also introduce $Y_0^{(j \rightarrow a)}$ and $Y_1^{(j \rightarrow a)}$, which are defined similarly to $Z_0^{(a \rightarrow i)}$ and $Z_1^{(a \rightarrow i)}$, but for rooted trees whose root is the node j , in the absence of the hyperedge a : the index 1 means that j is already matched and the index 0 that it is not. With the notation of figure 3, we have the relation

$$\begin{aligned}
 Z_0^{(a \rightarrow i)} &= \prod_{j \in a - i} \left(Y_0^{(j \rightarrow a)} + Y_1^{(j \rightarrow a)} \right), \\
 Z_1^{(a \rightarrow i)} &= e^{-\beta(\xi_a - d\mu)} \prod_{j \in a - i} Y_0^{(j \rightarrow a)}, \\
 Y_0^{(j \rightarrow a)} &= \prod_{b \in j - a} Z_0^{(b \rightarrow j)}, \\
 Y_1^{(j \rightarrow a)} &= \sum_{b \in j - a} Z_1^{(b \rightarrow j)} \prod_{c \in j - \{a, b\}} Z_0^{(c \rightarrow j)}.
 \end{aligned}
 \tag{18}$$

These formulae have simple interpretations: for instance, the first line means that when a is empty, the neighbouring nodes $j \in a - i$ can be equally matched or not with upstream hyperedges, while the second line means that when a is occupied, it generates a cost $\xi_a - d\mu$ and requires the nodes $j \in a - i$ to not be matched. From the conditional partition functions, we define the cavity fields

$$\begin{aligned} e^{\beta x^{(j \rightarrow a)}} &\equiv e^{\beta \mu} \frac{Y_0^{(j \rightarrow a)}}{Y_0^{(j \rightarrow a)} + Y_1^{(j \rightarrow a)}}, \\ e^{\beta u^{(a \rightarrow i)}} &\equiv e^{\beta(\xi_a - \mu)} \frac{Z_1^{(a \rightarrow i)}}{Z_0^{(a \rightarrow i)}}. \end{aligned} \quad (19)$$

These definitions are made to ensure a proper scaling when $\mu \rightarrow \infty$ and recover quantities used in previous studies for $d = 2$. Note however that for finite β , it is more natural to introduce $\psi^{(j \rightarrow a)} \equiv \exp[\beta(x^{(j \rightarrow a)} - \mu)]$ interpreted as the probability that node j is not matched in the absence of a (or equivalently that j is associated with a in a matching); this alternative notation will turn out to be particularly convenient when discussing the freezing phenomenon, in section 5.2. On a given rooted tree, it follows from equation (18) that the fields attached to the different oriented edges are related by the following *message passing rules*:

$$\begin{aligned} u^{(a \rightarrow i)} &= \sum_{j \in a - i} x^{(j \rightarrow a)}, \\ x^{(j \rightarrow a)} &= -\frac{1}{\beta} \ln \left(e^{-\beta \mu} + \sum_{b \in j - a} \exp(-\beta(\xi_a - u^{(b \rightarrow j)})) \right). \end{aligned} \quad (20)$$

The limit of infinite rooted trees is taken implicitly by considering the stationary distribution $\mathcal{P}(x)$ that is assumed to result from the repeated iteration of the message passing relations. By definition $\mathcal{P}(x)$ is a distribution of cavity fields over the different oriented edges that satisfies the following self-consistent equation, called the RS cavity equation:

$$\mathcal{P}(x^{(0)}) = \mathbb{E}_{k, \xi} \int \prod_{a=1}^k \prod_{j_a=1}^{d-1} dx^{(j_a)} \mathcal{P}(x^{(j_a)}) \delta(x^{(0)} - \hat{x}^{(k, \xi)}[\{x^{(j_a)}\}]) \quad (21)$$

where the function $\hat{x}^{(k, \xi)}$ is defined according to equation (20) as

$$\hat{x}^{(k, \xi)}[\{x^{(j_a)}\}] \equiv -\frac{1}{\beta} \ln \left(e^{-\beta \mu} + \sum_{a=1}^k \exp \left(-\beta \left(\xi - \sum_{j_a=1}^{d-1} x^{(j_a)} \right) \right) \right), \quad (22)$$

and the expectation $\mathbb{E}_{k, \xi}$ expresses the average over the disorder, which includes both an average over the connectivity k and over the rescaled costs ξ :

$$\mathbb{E}_{k, \xi}[F^{(k)}(\{\xi_a\})] \equiv \sum_{k=0}^{\infty} \frac{(d\mu)^k e^{-d\mu}}{k!} \prod_{a=1}^k \left(\frac{1}{d\mu} \int_0^{d\mu} d\xi_a \right) F^{(k)}(\xi_1, \dots, \xi_k). \quad (23)$$

The RS cavity equation (21) can be solved by a population dynamics algorithm, whose principle is presented in appendix A; the resulting distribution $\mathcal{P}(x)$ for $d = 3$ and

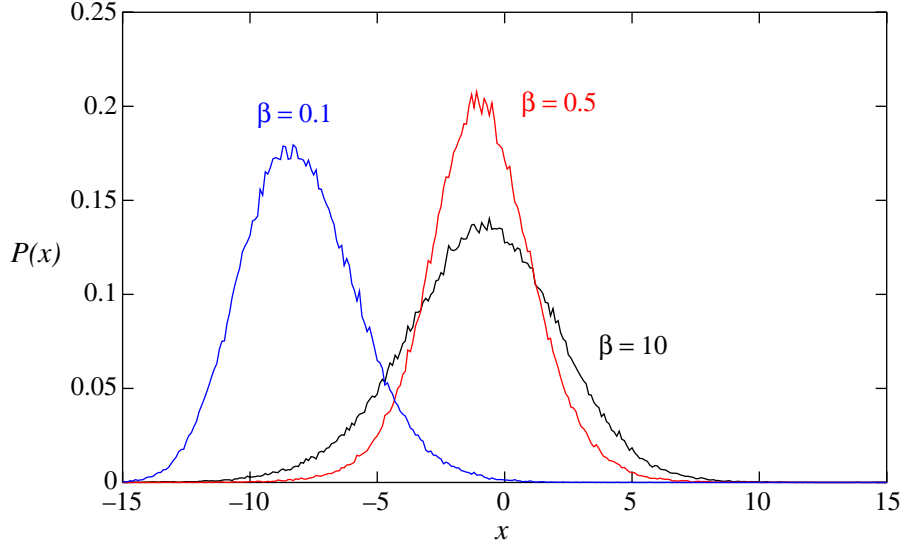


Figure 4. Distribution $\mathcal{P}(x)$ of cavity fields for 3-index matchings at different temperatures β in the replica symmetric approximation. These distributions are obtained by population dynamics, using algorithm P described in appendix A, with parameters $C = 60$, $\mathcal{N}_{\text{iter}} = 1$ and $\mathcal{N}_{\text{pop}} = 200\,000$.

different β is shown in figure 4. $\mathcal{P}(x)$ contains all the information on the equilibrium properties and, in particular, allows one to compute the free-energy density. It can be derived from the Bethe approximation which produces on a given hypergraph the formula

$$f(\beta) = \frac{1}{N} \left[\sum_i \Delta F^{(i+a \in i)}(\beta) - \sum_a (\ell_a - 1) \Delta F^{(a)}(\beta) \right] \quad (24)$$

where ℓ_a is the degree of hyperedge a , which here is $\ell_a = d$ independently of a . The shifts $\Delta F^{(i+a \in i)}(\beta)$ and $\Delta F^{(a)}(\beta)$ correspond respectively to the free-energy shift induced by the addition of a node i together with its connected hyperedges $a \in i$, and to the free-energy shift induced by the addition of hyperedge a . They are given by

$$\begin{aligned} \exp(-\beta \Delta F^{(i+a \in i)}(\beta)) &= \frac{Y_0^{(i)} + Y_1^{(i)}}{\prod_{a \in i} \prod_{j \in a-i} (Y_0^{(j)} + Y_1^{(j)})} \\ &= e^{-\beta \mu} + \sum_{a \in i} \exp \left(-\beta \left(\xi_a - \sum_{j \in a-i} x^{(j \rightarrow a)} \right) \right), \quad (25) \\ \exp(-\beta \Delta F^{(a)}(\beta)) &= \frac{Z_0^{(a)} + Z_1^{(a)}}{\prod_{j \in a-i} (Y_0^{(j)} + Y_1^{(j)})} = 1 + \exp \left(-\beta \left(\xi_a - \sum_{j \in a} x^{(j \rightarrow a)} \right) \right), \end{aligned}$$

where we introduced the analogues of the partition functions for rooted trees, but for the complete trees:

$$\begin{aligned}
 Z_0^{(a)} &= \prod_{j \in a} \left(Y_0^{(j \rightarrow a)} + Y_1^{(j \rightarrow a)} \right), \\
 Z_1^{(a)} &= e^{-\beta(\xi_a - d\mu)} \prod_{j \in a} Y_0^{(j \rightarrow a)}, \\
 Y_0^{(j)} &= \prod_{b \in j} Z_0^{(b \rightarrow j)}, \\
 Y_1^{(j)} &= \sum_{b \in j} Z_1^{(b \rightarrow j)} \prod_{c \in j-b} Z_0^{(c \rightarrow j)}.
 \end{aligned}
 \tag{26}$$

Physically, $Z_1^{(a)} / (Z_0^{(a)} + Z_1^{(a)})$ gives the probability for the hyperedge a to be included in the matching. By averaging over the realizations of the disorder, since the mean number of hyperedges per nodes is μ , we get

$$f_{\text{RS}}(\beta) = \mathbb{E}[\Delta F^{(i+a \in i)}(\beta)] - (d-1)\mu \mathbb{E}[\Delta F^{(a)}(\beta)]
 \tag{27}$$

with explicitly

$$\begin{aligned}
 \mathbb{E}[\Delta F^{(i+a \in i)}(\beta)] &= -\frac{1}{\beta} \mathbb{E}_{k,\xi} \int \prod_{a=1}^k \prod_{j_a=1}^{d-1} dx^{(j_a)} \mathcal{P}(x^{(j_a)}) \\
 &\quad \times \ln \left(e^{-\beta\mu} + \sum_{a=1}^k \exp \left(-\beta \left(\xi_a - \sum_{j_a=1}^{d-1} x^{(j_a)} \right) \right) \right), \\
 \mathbb{E}[\Delta F^{(a)}(\beta)] &= -\frac{1}{\beta} \mathbb{E}_{\xi} \int \prod_{j=1}^d dx^{(j)} \mathcal{P}(x^{(j)}) \ln \left(1 + \exp \left(-\beta \left(\xi_a - \sum_{j=1}^d x_j \right) \right) \right).
 \end{aligned}
 \tag{28}$$

4.3. Integral relations

The $\mu \rightarrow \infty$ limit can be taken explicitly. The corresponding equations generalize the formulae established by Mézard and Parisi in their first treatment of the 2-index matching problem [1]. For general d , they are

$$\begin{aligned}
 G(l) &= \frac{1}{\beta} \int_{-\infty}^{+\infty} \prod_{j=1}^{d-1} dy_j e^{-G(y_j)} B_d \left(l + \sum_{j=1}^{d-1} y_j \right), \\
 B_d(x) &\equiv \sum_{p=1}^{\infty} \frac{(-1)^{p-1} p^{d-2} e^{px}}{(p!)^d}.
 \end{aligned}
 \tag{29}$$

Given $G(l)$, the energy $\epsilon(\beta)$ and entropy $s(\beta)$ are

$$\begin{aligned}
 \epsilon(\beta) &= \frac{1}{\beta d} \int_{-\infty}^{+\infty} dl G(l) e^{-G(l)}, \\
 s(\beta) &= \int_{-\infty}^{+\infty} dl \left[e^{-e^l} - e^{-G(l)} \right] - \frac{d-2}{d} \int_{-\infty}^{+\infty} dl G(l) e^{-G(l)},
 \end{aligned}
 \tag{30}$$

and the free energy is obtained as $f(\beta) = \epsilon(\beta) - s(\beta)/\beta$. The relation between the function $G(l)$ and the order parameter $\mathcal{P}(x)$ is, up to a change of variable, a Laplace transform:

$$e^{-G(l)} = \int_{-\infty}^{+\infty} dx \mathcal{P}(x) e^{-e^l - \beta x}. \tag{31}$$

From the practical point of view of numerically solving the cavity equations, the finite μ cavity equations are however easier to handle than these compact formulae.

4.4. Zero-temperature limit

With a view to an extension of the mathematical approach from 2-index to d -index matchings with $d > 2$, it is interesting to discuss in some detail the relations between our equations and those used by Aldous in his rigorous study of the 2-index matching problem [3]. Aldous’s formalism is obtained from our RS cavity equations by taking the zero-temperature, $\beta \rightarrow \infty$ limit. When $\beta \rightarrow \infty$, equations (20) become

$$\begin{aligned} u^{(a \rightarrow i)} &= \sum_{j \in a-i} x^{(j \rightarrow a)}, \\ x^{(j \rightarrow a)} &= \min_{b \in j-a} (\xi_b - u^{(b \rightarrow j)}). \end{aligned} \tag{32}$$

Taking $\mu = \infty$ and $d = 2$ leads to the *recursive distributional equation* [29],

$$x^{(a)} = \min_b (\xi_b - x^{(b)}) \tag{33}$$

on which Aldous’s work is based [3]. A difference is however that its costs ξ_b derive from a Poisson point process (the uniform distribution does not make sense when $\mu = \infty$). This Poisson process can nonetheless be related to our formalism by implementing a variant of the cut-off procedure. Consider selecting at each step of the cavity recursion the k parents of smallest costs, k being now fixed. Then the successive k costs are distributed according to a Poisson process with rate one. Nonetheless, while the cavity recursion is perfectly well defined, the corresponding system on a given hypergraph does not make sense: a hyperedge may belong to the list of the hyperedges with the k th smallest costs for one of its nodes but not for another one. This is why we introduced the version with a cut-off on the costs, which constitutes for finite μ a perfectly sensible statistical physics model. From a purely formal point of view the version with cut-off on the number of connected clauses works as well, and provides an alternative formulation for numerically solving the cavity equations (see appendix A for the details and figure 9 for an illustration).

The cavity fields at zero temperature have an interpretation in terms of differences in ground state energies. The cavity field $x^{(j \rightarrow a)}$ corresponds to the extra cost of a particle on node j with respect to no particle, in the absence of hyperedge a , and the cavity bias $u^{(a \rightarrow i)}$ to the cost of connecting the node i to the hyperedge a . Note that these quantities are actually well defined only if μ is kept finite; otherwise a particle cannot be removed or added without destroying the perfect matching, i.e., without leaving the space of admissible configurations. Similarly, the total fields on the complete graph are

$$\begin{aligned} U^{(a)} &= \sum_{j \in a} x^{(j \rightarrow a)}, \\ X^{(i)} &= \min_{b \in i} (\xi_b - u^{(b \rightarrow i)}). \end{aligned} \tag{34}$$

From the interpretation given, it appears that the hyperedges which indeed participate in the optimal matching are those which achieve the minima, i.e., the solution is given by

$$n_a = \delta_{a,a^*}, \quad a^* = \arg \min_a (\xi_a - u^{(a \rightarrow i)}). \quad (35)$$

Since this has to hold for all $i \in a$, the question arises of whether this prescription effectively defines a matching, i.e., whether $\arg \min_a (\xi_a - u^{(a \rightarrow i)}) = \arg \min_a (\xi_a - u^{(a \rightarrow j)})$ for all $i, j \in a$. A positive answer is obtained by generalizing to $d > 2$ the *inclusion criterion* invoked by Aldous when $d = 2$, which states

$$a^* = \arg \min_a (\xi_a - u^{(a \rightarrow i)}) = 1 \iff \xi_a \leq u^{(a \rightarrow i)} + x^{(i \rightarrow a)}. \quad (36)$$

The independence on i is then a consequence of the identity $u^{(a \rightarrow i)} + x^{(i \rightarrow a)} = U^{(a)}$. The proof of the inclusion criterion itself is straightforward with the present notation: if $a^* = \arg \min_a (\xi_a - u^{(a \rightarrow i)})$,

$$\xi_{a^*} - u^{(a^* \rightarrow i)} = \min_{b \in i} (\xi_b - u^{(b \rightarrow i)}) \leq \min_{b \in i - a^*} (\xi_b - u^{(b \rightarrow i)}) = x^{(i \rightarrow a^*)}. \quad (37)$$

Reciprocally, if $a \neq \arg \min_b (\xi_b - u^{(b \rightarrow i)})$,

$$\xi_a - u^{(a \rightarrow i)} \geq \min_{b \in i} (\xi_b - u^{(b \rightarrow i)}) = \min_{b \in i - a} (\xi_b - u^{(b \rightarrow i)}) = x^{(i \rightarrow a)}. \quad (38)$$

As an alternative to the Bethe formula, the value of the optimal matching can be obtained by inferring $\langle \xi_{a^*} \rangle$ from the distribution of the fields $\mathcal{P}(x)$. Thanks to the inclusion criterion, we have

$$\begin{aligned} \mathcal{L}_{\text{RS}}^{(d)} &= \frac{1}{d} \langle \xi_{a^*} \rangle = \frac{1}{d} \int_0^\infty d\xi \xi \text{Prob}(U > \xi) \\ &= \frac{1}{d} \int_0^\infty d\xi \int \prod_{j=1}^d dx_j \mathcal{P}(x_j) \xi \theta \left(\sum_{j=1}^d x_j - \xi \right) \end{aligned} \quad (39)$$

where the factor d corresponds to the number of nodes per hyperedge and θ represents the Heaviside function, $\theta(x) = 1$ if $x > 0$ and 0 otherwise. The RS cavity equations at zero temperature can also be written in terms of closed integral relations that generalize known equalities for the $d = 2$ case:

$$\tilde{G}(x) = \int_{\sum_j t_j > -x} \prod_{j=1}^{d-1} dt_j \tilde{G}'(t_j) e^{-\tilde{G}(t_j)} \left(x + \sum_{j=1}^{d-1} t_j \right), \quad (40)$$

$$\mathcal{L}_{\text{RS}}^{(d)} = \frac{1}{2d} \int_{\sum_j x_j > 0} \prod_{j=1}^d dx_j \tilde{G}'(x_j) e^{-\tilde{G}(x_j)} \left(\sum_{j=1}^d x_j \right)^2. \quad (41)$$

The distribution $\tilde{G}(x)$ is related to the RS distribution \mathcal{P} of the cavity fields by

$$\tilde{G}(x) = -\ln \int_x^\infty dt \mathcal{P}(t), \quad (42)$$

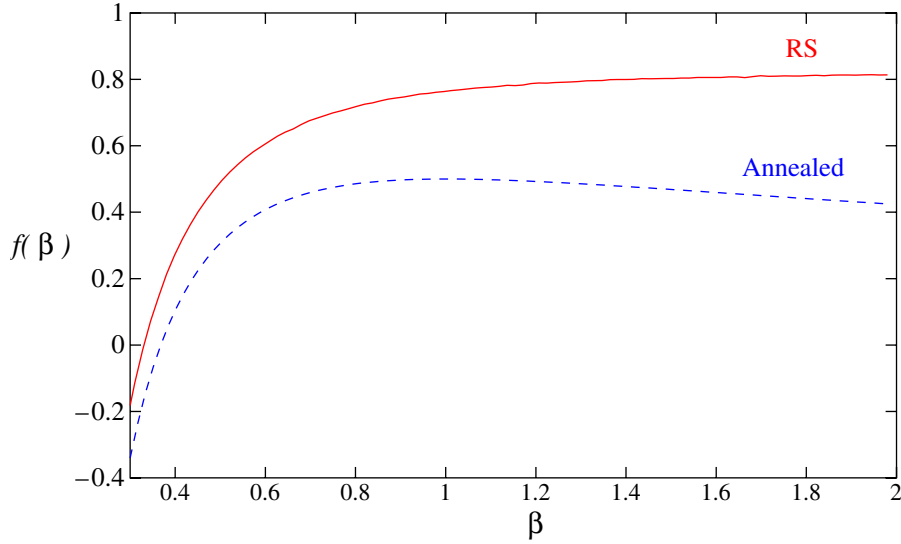


Figure 5. Annealed and replica symmetric free energies for the simple 2-index matching problem. While the annealed curve has a maximum at 1, the RS curve is strictly monotone. Its $\beta \rightarrow \infty$ limit corresponds to the exactly known value $\pi^2/12 \simeq 0.82$.

and can be obtained from the finite temperature order parameter $G(l) = G_\beta(l)$ given in equation (29) via

$$\tilde{G}(x) = \lim_{\beta \rightarrow \infty} G_\beta(\beta^{1/(d-1)}x). \tag{43}$$

Comparing with the predictions of the cavity method, $\mathcal{L}_{\text{RS}}^{(d)} = \mathbb{E}[\Delta\epsilon^{(i+a\epsilon i)}] - (d-1)\mu\mathbb{E}[\Delta\epsilon^{(a)}]$, we obtain a consistency condition that the RS distribution must satisfy:

$$\mathbb{E}[x] = \frac{2-d}{2d} \mathbb{E} \left[\left(\sum_{j=1}^d x_j \right)^2 \theta \left(\sum_{j=1}^d x_j \right) \right] \tag{44}$$

where the average $\mathbb{E}[\cdot]$ is here taken with respect to \mathcal{P} . This formula is indeed numerically verified with a good precision, in agreement with the equivalence between the two approaches. As a corollary, it shows that $\mathbb{E}[x] < 0$ unless $d = 2$ where $\mathbb{E}[x] = 0$ (we recall that in this case one has in fact an explicit formula [1], $\mathcal{P}(x) = 1/[4 \cosh^2(x/2)]$). Finally, we note that for $d > 2$, the RS energy at zero temperature is only dependent on the mean of $\mathcal{P}(x)$, $\mathcal{L}_{\text{RS}}^{(d)} = -\mathbb{E}[x]/(d-2)$. However as shown in the following, the RS approach yields incorrect predictions when $d \geq 3$.

4.5. Entropy crisis

Using the population dynamics algorithm described in appendix A, we obtain for the RS free energy $f_{\text{RS}}(\beta)$ the curves displayed in figure 5 ($d = 2$) and 6 ($d = 3$). For $d = 2$, the free energy is an increasing function of β with limit $f_{\text{RS}}(\beta = \infty) = \pi^2/12 \simeq 0.82$ corresponding to the cost of a minimal 2-index matching. The free energy obtained

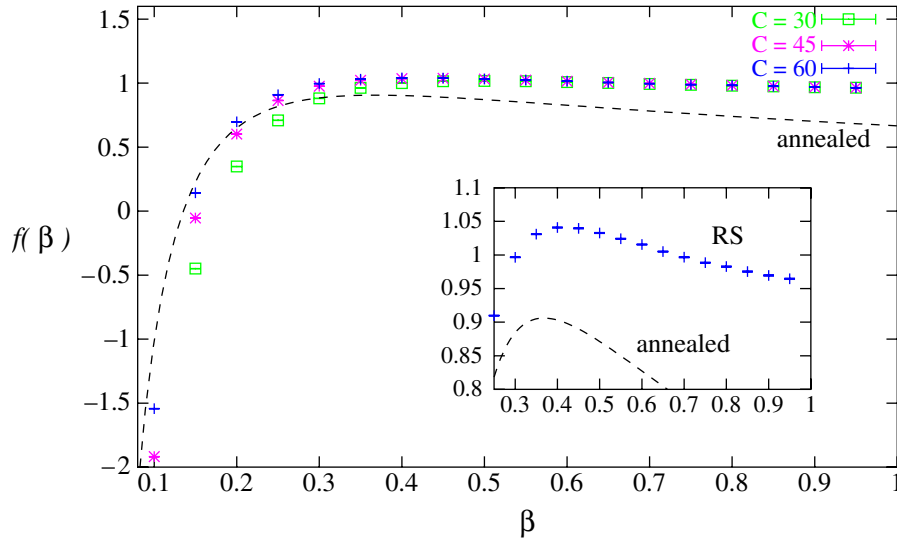


Figure 6. Replica symmetric free energy $f_{\text{RS}}(\beta)$ for the 3-index matching problem with different cut-offs $C = d\mu = 30, 45, 60$ as a function of the inverse temperature β . This curve has been obtained using algorithm P of appendix A with parameters $\mathcal{N}_{\text{pop}} = 20\,000$ and $\mathcal{N}_{\text{iter}} = 20\,000$; for comparison, the annealed free energy is also represented with a dashed line. In the inset is a zoom of the data with $C = 60$ more clearly displaying the non-physical decrease of $f_{\text{RS}}(\beta)$ for $\beta > \beta_s \simeq 0.41$.

for $d = 3$ is qualitatively different, as it displays a maximum at a finite temperature $\beta_s \simeq 0.41$ (see figure 7). This *entropy crisis* reflects an inconsistency of the RS approach¹. If one assumes that the RS approximation holds at high temperature in some range of temperature (a non-trivial statement), a phase transition must occur at some $\beta_c \leq \beta_s$.

4.6. Stability of the replica symmetric ansatz

Replica symmetry fails to correctly describe the low temperature properties of many frustrated systems [8]. A necessary requirement for its validity is that it be stable. Here we show that when $d = 3$ the RS solution is unstable below a strictly positive temperature, that is for $\beta > \beta_i$. Even if the breakdown of the RS hypothesis was already inferred above from the negative value of the RS entropy, studying the stability is instructive since the relative positions of β_i and β_s will establish the discontinuous nature of the phase transition. In [9], Mézard and Parisi used the replica method to prove that the RS ansatz is stable when $d = 2$ [1]; their approach is however quite complicated (see [11] for a recent re-examination of their analysis), and to tackle the $d = 3$ case, we adopt a simpler approach based on the cavity method [30]. Physically, it amounts to computing the non-linear susceptibility χ_2 and checking that it does not diverge [31]. Picking a hyperedge labelled 0 at random, this susceptibility is written as

$$\chi_2 = \sum_a \langle n_0 n_a \rangle_c^2 \simeq \sum_{r=0}^{\infty} [C(d-1)]^r \mathbb{E}[\langle n_0 n_r \rangle_c^2] \quad (45)$$

¹ Additional evidence that the naive generalization of the formulae used when $d = 2$ fails for $d \geq 3$ is found when $d = 4$ (extending presumably to all $d \geq 4$) where the RS prediction for $\mathcal{L}_{\text{RS}}^{(d)}$ violates the annealed bound.

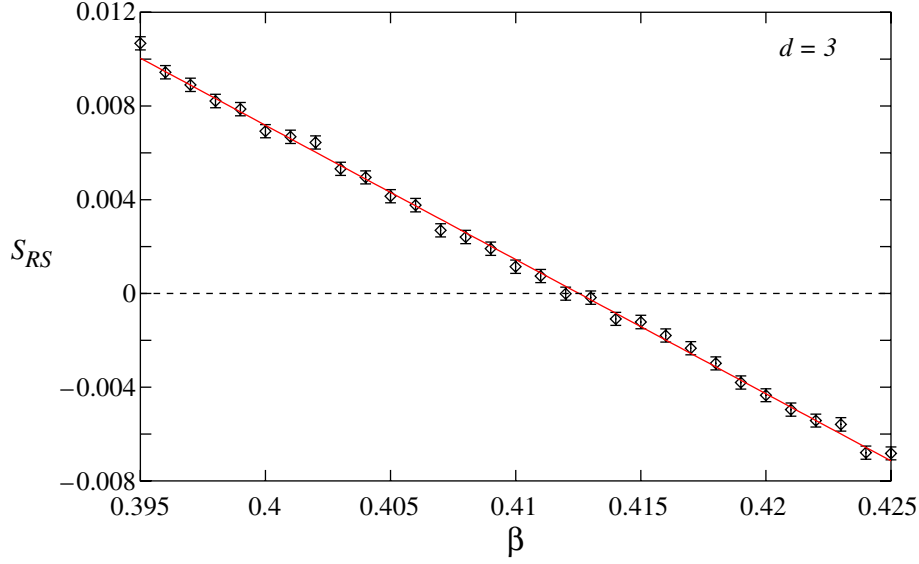


Figure 7. RS entropy of the 3-index matching problem. The data are from population dynamics, using algorithm R presented in appendix A, with $K = 50$, $\mathcal{N}_{\text{pop}} = 50\,000$ and $\mathcal{N}_{\text{iter}} = 5000$. The line is a linear regression. The RS entropy is found to vanish at $\beta_c = 0.412 \pm 0.001$.

where $\langle \cdot \rangle$ denotes the thermal average and $\mathbb{E}[\cdot]$ the spatial average over the disorder. Using the fluctuation-dissipation relation, the averaged squared correlation function $\mathbb{E}[\langle n_0 n_r \rangle_c^2]$ of two hyperedges separated by distance r can be expressed in terms of the cavity fields as [31]

$$\mathbb{E}[\langle n_0 n_r \rangle_c^2] \sim \mathbb{E} \left[\prod_{i=1}^r \left(\frac{\partial \hat{x}^{(k,\xi)}(x_{i_1}, \dots, x_{i_{(d-1)k}})}{\partial x_{i_1}} \right)^2 \right] \quad (r \rightarrow \infty), \quad (46)$$

where the averaging $\mathbb{E}[\cdot]$ is performed with respect to the distribution of the disorder (k, ξ) and to the distribution $\mathcal{P}(x)$ of the cavity fields, except for the x_{i_1} with $i > 1$ which are fixed by $x_{(i+1)_1} = \hat{x}^{(k,\xi)}(x_{i_1}, \dots, x_{i_{(d-1)k}})$. To determine whether the series in equation (45) converges or not, we compute

$$\ln \mu_r = r \ln[C(d-1)] + \ln \mathbb{E} \left[\prod_{i=1}^r \left(\frac{\partial \hat{x}^{(k,\xi)}(x_{i_1}, \dots, x_{i_{(d-1)k}})}{\partial x_{i_1}} \right)^2 \right] \quad (47)$$

by using cavity fields from the population dynamics, and check whether $\lim_{r \rightarrow \infty} (\ln \mu_r)/r < 0$ or not. The numerical results are limited to small values of r , but as shown in figure 8 they are sufficient for us to conclude unambiguously that an instability shows up for 3-index matchings at $\beta_i \simeq 0.6$, thus confirming the incorrectness of the RS ansatz for describing the $\beta = \infty$ limit (the same procedure with $d = 2$ consistently finds no instability). In addition, since the instability takes place only after the entropy crisis, $\beta_i > \beta_s$, we conclude from this analysis that the phase transition, located at $\beta_c \leq \beta_s$, must be *discontinuous* as a function of the order parameter.

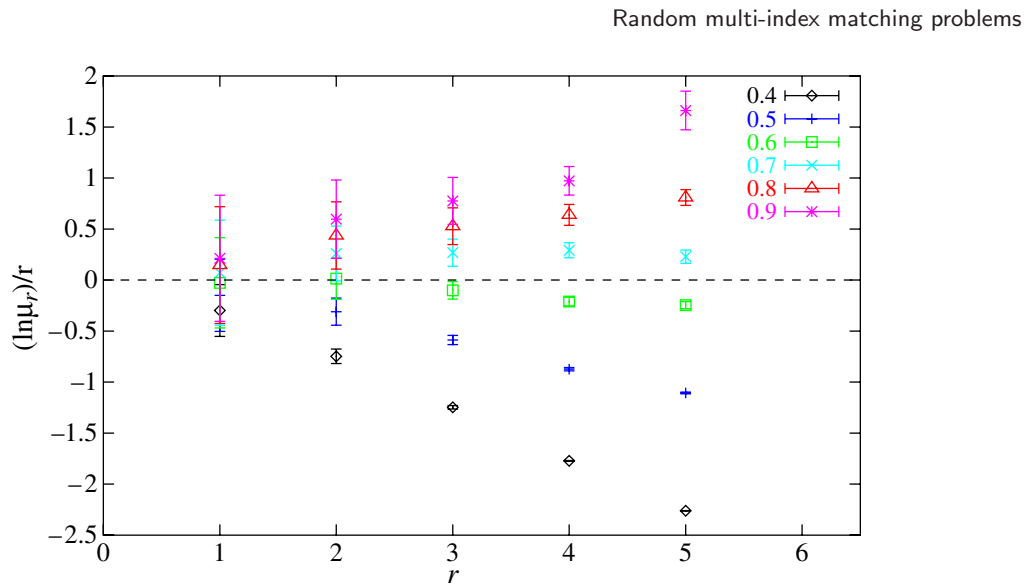


Figure 8. Stability analysis of the RS solution for the 3-index matching problem at finite temperature. $(\ln \mu_r)/r$ is plotted versus r for different temperatures β (from algorithm P given in appendix A with $C = 36$, $\mathcal{N}_{\text{pop}} = 20\,000$ and $\mathcal{N}_{\text{iter}} = 10^9$). The RS solution is stable if the slope of $(\ln \mu_r)/r$ is negative (see the text), which is found to be the case for $\beta < \beta_i \simeq 0.6$.

5. Replica symmetry breaking

The inconsistencies of the RS ansatz indicate that replica symmetry must be broken in the low temperature phase. This feature is present in many other NP-hard combinatorial optimization problems and is commonly overcome by adopting a one-step replica symmetry breaking (1RSB), which, in most favourable cases, turns out to be exact.

5.1. General 1RSB ansatz

As formulated by Aldous with the *essential uniqueness property* [3], replica symmetry in matching problems means that *quasi-solutions*, that is low energy configurations (LECs), all share most of their hyperedges. In contrast, replica symmetry breaking (RSB) refers to a situation where LECs arise, which, while being close in cost to the optimal solution, are far apart in the configurational space (the measure of distances is the overlap between two matchings, i.e., the fraction of common hyperedges; see section 5.3). One-step replica symmetry breaking (1RSB) is a particular scheme of RSB where the structure of the set of LECs can be described with only two characteristic distances, d_0 and $d_1 < d_0$. For it to be correct, two LECs taken at random (with the Gibbs probability measure when working at finite β) must be typically found either at distance d_0 or d_1 . In the replica jargon, close by LECs (at the short distance d_1) are said to belong to the same *state* (or cluster). At the level of 1RSB, it is assumed that the number $\mathcal{N}_N(f)$ of states with a given free energy f grows exponentially with N and is characterized by a *complexity* $\Sigma(f)$ defined by $\Sigma(f) = \lim_{N \rightarrow \infty} [\ln \mathcal{N}_N(f)]/N$.

The 1RSB cavity method derives this ‘entropy of states’ by a Legendre transformation method mimicking the derivation of entropy from the free energy in canonical statistical

mechanics [32]. The object generalizing the free energy is the replica potential $\phi(\beta, m)$; the parameter m is the Lagrange multiplier fixing the free energy of the relevant states, in the same way that the temperature β selects the energy of equilibrium configurations in the canonical ensemble. The replica potential is defined as

$$e^{-N\beta m\phi(\beta, m)} \equiv \sum_{\alpha} e^{-N\beta m f_{\alpha}}, \quad (48)$$

where the sum is over the states α , and f_{α} denotes the free energy of a system whose configurations are restricted to α . To obtain the relevant states for the equilibrium properties, replica theory prescribes choosing the m in $[0, 1]$ that maximizes $\phi(\beta, m)$ [8], so that the equilibrium free energy is given by

$$f_{\text{1RSB}}(\beta) = \max_{0 \leq m \leq 1} \phi(\beta, m). \quad (49)$$

Calculating $\phi(\beta, m)$ requires introducing as order parameter a distribution $\mathcal{Q}[Q^{(j \rightarrow a)}]$ over the oriented edges ($j \rightarrow a$) of distributions $Q^{(j \rightarrow a)}(x)$ of the cavity fields, taken over the different states α [27]. The 1RSB cavity equations for the order parameter read

$$\begin{aligned} \mathcal{Q}[Q^{(0)}] &= \mathbb{E}_{k, \xi} \int \prod_{a=1}^k \prod_{j_a=1}^{d-1} \mathcal{D}Q^{(j_a)} \mathcal{Q}[Q^{(j_a)}] \delta \left[Q^{(0)} - \hat{Q}^{(k, \xi)}[\{Q^{(j_a)}\}] \right], \\ \hat{Q}^{(k, \xi)}[\{Q^{(j_a)}\}](x^{(0)}) &= \frac{1}{Z} \int \prod_{a=1}^k \prod_{j_a=1}^{d-1} dx^{(j_a)} Q^{(j_a)}(x^{(j_a)}) \delta \left(x^{(0)} - \hat{x}^{(k, \xi)}(\{x^{(j_a)}\}) \right) \\ &\quad \times \exp \left(-\beta m \Delta \hat{F}_n^{(k, \xi)}(\{x^{(j_a)}\}) \right), \end{aligned} \quad (50)$$

where $\hat{x}^{(k, \xi)}$ is given by equation (22) and the reweighting term is

$$\exp \left(-\beta m \Delta \hat{F}_n^{(k, \xi)}(\{x^{(j_a)}\}) \right) = e^{-\beta \mu} + \sum_{a=1}^k \exp \left(-\beta \left(\xi_a - \sum_{j_a=1}^{d-1} x^{(j_a)} \right) \right). \quad (51)$$

The latter corresponds to the shift of free energy due to the addition of the new node. Its presence ensures that the different states described by the $Q^{(j \rightarrow a)}(x)$ do indeed all have the same free energy, in spite of the fact that the addition of a node inevitably introduces a free-energy shift. The distribution $\mathcal{Q}[Q]$ determines the replica potential $\phi(\beta, m)$ whose explicit expression is

$$\phi(\beta, m) = \mathbb{E}[\Phi^{(i+a \in i)}(\beta, m)] - (d-1)\mu \mathbb{E}[\Phi^{(a)}(\beta, m)] \quad (52)$$

with

$$\begin{aligned} \mathbb{E}[\Phi^{(i+a \in i)}(\beta, m)] &= -\frac{1}{\beta} \mathbb{E}_{k, \xi} \int \prod_{a=1}^k \prod_{j_a=1}^{d-1} \mathcal{D}Q^{(j_a)} \mathcal{Q}[Q^{(j_a)}] \\ &\quad \times \ln \left[\int \prod_{a=1}^k \prod_{j_a=1}^{d-1} dx^{(j_a)} Q^{(j_a)}(x^{(j_a)}) \exp \left(-m\beta \Delta \hat{F}_n^{(k, \xi)}(\{x^{(j_a)}\}) \right) \right] \end{aligned} \quad (53)$$

and

$$\mathbb{E}[\Phi^{(a)}(\beta, m)] = -\frac{1}{\beta} \mathbb{E}_{\xi} \int \prod_{j=1}^d \mathcal{D}Q^{(j)} \mathcal{Q}[Q^{(j)}] \times \ln \left[\int \prod_{j=1}^d dx^{(j)} Q^{(j)}(x^{(j)}) \left(1 + \exp \left(-\beta \left(\xi - \sum_{j=1}^d x^{(j)} \right) \right) \right)^m \right]. \quad (54)$$

The 1RSB equations can in principle be numerically solved via a population dynamics algorithm [27]. However, our efforts in this direction failed to yield a sensible order parameter because the fields were found to diverge as μ was increased: the reason for this behaviour is elucidated below.

5.2. Frozen 1RSB ansatz

Although rarely explicitly mentioned, there exists a replica symmetry breaking ansatz somewhat intermediate between the RS and general 1RSB as just described. The *frozen 1RSB ansatz*, which will be argued to apply to matchings, is a particular realization of the 1RSB scheme where states are made of single configurations (or, more generally, of a non-exponential number of configurations). In such a case, all the information can be extracted from the RS quantities, provided that they are adequately reinterpreted. Consider for instance the definition given by equation (48) in the special case where states α have no internal entropy, i.e., $f_{\alpha} = \epsilon_{\alpha}$. We thus have

$$e^{-N\beta m \phi(\beta, m)} \equiv \sum_{\alpha} e^{-N\beta m f_{\alpha}} = \sum_{\alpha} e^{-N\beta m \epsilon_{\alpha}} \equiv e^{-N\beta m f_{\text{RS}}(\beta m)} \quad (55)$$

where the last equality holds because of the very definition of a RS free energy. The replica potential ϕ can therefore be expressed in terms of the RS free energy only,

$$\phi(\beta, m) = f_{\text{RS}}(\beta m). \quad (56)$$

Following the prescriptions of replica theory, the quenched free energy is obtained by maximizing $\phi(\beta, m)$ over $m \in [0, 1]$. Being a concave function, the RS free energy can have at most one maximum. If β_s denotes the location of this maximum (with perhaps $\beta_s = \infty$, like for 2-index matchings), we obtain that $f_{\text{1RSB}}(\beta) = \phi(1, \beta) = f_{\text{RS}}(\beta)$ for $\beta < \beta_s$ and $f_{\text{1RSB}}(\beta) = \phi(\beta_s/\beta, \beta) = f_{\text{RS}}(\beta_s)$ for $\beta > \beta_s$. In other words, starting from the assumption that the content of states is trivial, the frozen ansatz predicts a complete freezing of the system at the point β_s where the RS entropy becomes zero: for $\beta > \beta_s$, the system is trapped in a single configuration and its free energy stays constant when the temperature is further decreased (β increased).

This scenario is already known to apply to a few models of disordered systems, including the random energy model (REM) [33], the directed polymer on disordered trees [34], the binary perceptron [35], multi-layer neural networks with binary weights [36] and the XOR-SAT problem on its core [37] (with a particular case being error-correcting codes of the Gallager type [38]); it also applies to p -spin models, q -state Potts models or K -SAT models in the limit where p, q or $K \rightarrow \infty$ [33, 39, 40]. Our intention is here both to add the matchings to this list, and to clarify the conditions under which such a scenario may apply. At this stage, we can already state the following necessary conditions (all satisfied by d -index matchings with $d \geq 3$):

- (i) the RS entropy must become negative at a finite β_s ;
- (ii) the RS solution must be stable up to (at least) β_s ;
- (iii) no discontinuous 1RSB transition must be detected before β_s .

In addition to these properties, the consistency of the frozen ansatz requires the model to have particular kinds of constraints, called *hard constraints*. Elucidating this point requires a more refined description of the relation between the frozen 1RSB order parameter and the RS order parameter. First remember that in the RS picture at finite temperature β , one has a spatial distribution $\mathcal{P}(x^{(j \rightarrow a)})$ of cavity fields, where, following equation (19), $\psi_{\text{RS}}^{(j \rightarrow a)} \equiv \exp[\beta(x^{(j \rightarrow a)} - \mu)]$ is interpreted as giving the probability under the Boltzmann measure that node j is not matched given that the hyperedge a is absent. For a general 1RSB problem, the order parameter is instead $Q[Q^{(j \rightarrow a)}(x^{(j \rightarrow a)})]$ where $\psi^{(j \rightarrow a)} \equiv \exp[\beta(x^{(j \rightarrow a)} - \mu)]$ is again a thermal probability, but now restricted to a particular state taken from the distribution over states $Q^{(j \rightarrow a)}$. In this context, a RS system, characterized by a single state, has $Q^{(j \rightarrow a)}(\psi^{(j \rightarrow a)}) = \delta(\psi^{(j \rightarrow a)} - \psi_{\text{RS}}^{(j \rightarrow a)})$. For a system in a frozen glassy phase instead, the thermal averages inside each state are trivial since there is a single frozen configuration, $\psi^{(j \rightarrow a)} = 0$ or 1 meaning that a particle is present or absent with probability one. Therefore, the relation with the RS order parameter has the form

$$Q^{(j \rightarrow a)}(\psi^{(j \rightarrow a)}) = \psi_{\text{RS}}^{(j \rightarrow a)} \delta(\psi^{(j \rightarrow a)}) + (1 - \psi_{\text{RS}}^{(j \rightarrow a)}) \delta(\psi^{(j \rightarrow a)} - 1). \quad (57)$$

Plugging this expression into the general 1RSB cavity equation, it is found that such an ansatz is consistent only if the system satisfies the condition that in the cavity recursion, the variable on a node is completely determined by the values of the variables on the neighbouring nodes. Such is the case with matchings when $\mu = \infty$ where a particle is to be assigned to a hyperedge if and only if none of the neighbouring edges are occupied. This is however not the case in all constraint problems. Consider for instance the 3-colouring problem where each node is assigned one of three colours with the constraint that its colour must differ from its neighbours: in the case where all the neighbours have the same colour, the choice is left for the node between the two other colours. When a variable is fixed by the value of its neighbours in the cavity recursion, we say that the system has *hard constraints*; hard constraints can be shown [41] to indeed be present in the binary perceptron and in the XOR-SAT model on its core, models where the frozen ansatz applies too. Finally, we note that in the presence of hard constraints, the cavity fields $\psi^{(j \rightarrow a)}$ take at the 1RSB level values 0 and 1 only, which are associated with $x^{(j \rightarrow a)} = \mu$ and $-\infty$. This explains the divergences observed when trying to implement the 1RSB population algorithm at zero temperature with $\mu \rightarrow \infty$.

5.3. Distances

As mentioned in section 5.1, a 1RSB glassy system is generally described by two distances, d_0 , corresponding to the typical distance between two states, and d_1 , corresponding to the typical distance between two configurations inside a common state. In the case of a frozen 1RSB glassy phase, one has however $d_1 = 0$ and the structure of low energy configurations (LECs) is characterized by only one distance, d_0 . If $\langle n_a \rangle$ denotes the mean occupancy of a particular hyperedge a , with the average $\langle \cdot \rangle$ taken over the LECs, the probability for a to

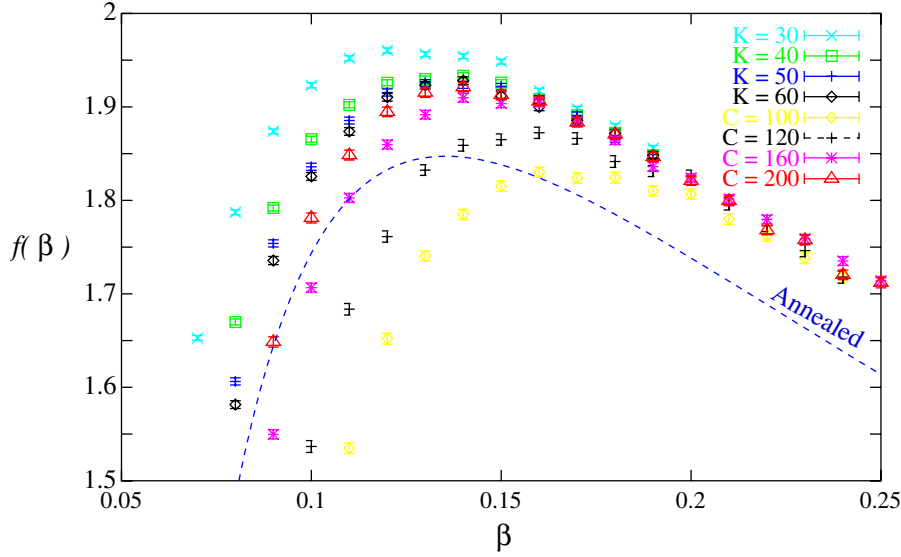


Figure 9. RS free energies for $d = 4$ as obtained from the two versions of the population dynamics algorithm described in appendix A (here with $\mathcal{N}_{\text{pop}} = 10\,000$ and $\mathcal{N}_{\text{iter}} = 1000$). Note that the approximation based on Poissonian graphs (mean connectivities $C = 100, 120, 160, 200$) approaches the solution from below, while the approximation based on regular graphs (fixed connectivities $K = 30, 40, 50, 60$) approaches it from above. As expected, the two limits $C \rightarrow \infty$ and $K \rightarrow \infty$ are found to match.

belong to two different LECs is given by $\langle n_a \rangle^2$. Averaging over the different hyperedges, it defines the overlap

$$q = \mathbb{E}[\langle n_a \rangle^2], \quad (58)$$

which is directly related to the typical distance between LECs through $d_0 = 1 - q$. As argued before, for a system in a frozen glassy phase the distribution of energies of the LECs is described by the thermal average at β_c in the RS approximation, so that

$$\langle n_a \rangle = \frac{Y_1^{(a)}}{Y_0^{(a)} + Y_1^{(a)}} = \frac{1}{1 + \exp(-\beta_c(\xi_a - \sum_{i \in a} x^{(i \rightarrow a)}))}. \quad (59)$$

Averaging over the disorder therefore yields

$$q = \mathbb{E}[\langle n_a \rangle_{\beta_c}^2] = \mathbb{E}_{\xi_a} \int \prod_{j=1}^d dx^{(j)} \mathcal{P}(x^{(j)}) \left(1 + \exp \left(-\beta_c \left(\xi_a - \sum_{i \in a} x^{(j)} \right) \right) \right)^{-2}. \quad (60)$$

The overlap $q(\beta)$ is represented for all values of β in figure 10 when $d = 3$; given the value of β_c obtained before, we get $q = q(\beta_c) = 0.321 \pm 0.002$.

6. Numerical analysis of finite size systems

The theoretical analysis provided concerned the $M \rightarrow \infty$ limit. How is that limit reached, and in particular is the convergence exponentially fast in M or is it algebraic? To answer

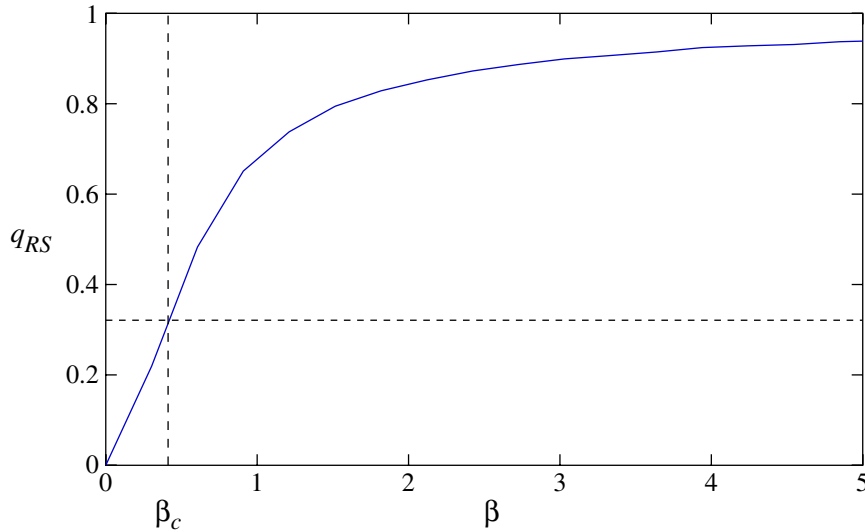


Figure 10. Overlap $q(\beta) = \mathbb{E}[\langle n_a \rangle^2]$ in the 3-index matching as given by equation (60). In particular $q(\beta_c) = 0.321 \pm 0.002$ describes the typical overlap between two low energy matchings, that is the fraction of hyperedges generically shared.

such questions, we consider in this section the properties of d -partite matchings when M is finite; in the absence of other tools, we do this numerically. It should be clear that the most challenging questions concern the low temperature phase of our system; because of that, we will focus on the optimum matching and low lying excitations. Even though such a numerical approach requires sampling the disorder (random instances) and extracting for instance distributions with inevitable statistical uncertainties, it will give evidence that our frozen 1RSB ansatz is correct; the comparison between the two approaches is summarized in table 1. The numerical analysis will also provide some statistical properties of finite size systems that are of interest on their own.

6.1. The branch and bound procedure

When M is very small, it is possible to enumerate all $[M!]^{d-1}$ d -partite matchings of a given sample. Not surprisingly, this becomes unwieldy even when M reaches 10, forcing us to choose an alternative approach. Since it is the *low energy* matchings that are of greatest interest, we have developed a branch and bound algorithm that computes the p lowest energy matchings, for any given p . Some technical aspects of the algorithm are presented in appendix B, but the essential elements are as follows.

We represent a matching via a list of M hyperedges, one for each of the M sites of the first set (recall that there are d sets, each of M sites). Such a representation includes also some non-legal matchings as some of the sites in the second or higher sets could belong to more than one hyperedge; if a matching is not legal, it is discarded. This representation can be mapped onto a rooted tree: each level of the tree is associated with one of the sites of the first set, while a segment (branch) emerging from a node corresponds to a choice of hyperedge that contains the site of that node's level. The root node is associated with

Table 1. This table sums up the different numerical values for the optimal cost $\mathcal{C}^{(d)}$, critical temperature β_c and typical overlap q of d -partite matching problems with $d = 2, 3, 4$. Exact values are given for $d = 2$ while for $d = 3, 4$ we confront the predictions of the cavity method with the results obtained by analysing instances of small size.

	$d = 2$		$d = 3$		$d = 4$	
	Exact	Cav.	Num.	Cav.	Num.	
$\mathcal{C}^{(d)}$	$\pi^2/6$	3.126 ± 0.002	3.09 ± 0.03	7.703 ± 0.002	7.22 ± 0.08	
β_c	∞	0.412 ± 0.001	0.405 ± 0.01	0.135 ± 0.002	0.140 ± 0.01	
q	1	0.321 ± 0.002	0.325 ± 0.005	0.088 ± 0.002	0.080 ± 0.005	

the first site, the nodes of the next level are associated with the second site etc. This tree is regular, each node having M^{d-1} outgoing segments as there are that many hyperedges containing a given site of the first set. Furthermore, it has $M + 1$ levels: there is one level for each site of the first set while the last level consists of leaves rather than of nodes; each leaf corresponds to a candidate matching specified by the list of hyperedges obtained when going from the tree's root to that leaf. This list may correspond to a legal matching or not, but each matching appears exactly once as a leaf. (In fact, there are $M^{M(d-1)}$ leaves while there are only $[M!]^{d-1}$ legal matchings.)

The principle of the branch and bound algorithm is to find those leaves which satisfy the desired criterion (the energy must be less than or equal to that of the p th-lowest energy matching) by exploiting a pruning procedure, thereby avoiding having to explore all leaves. To begin our pruned search, we produce p distinct legal matchings and put them into a list \mathcal{L} ; the largest energy of the matchings in this list is an upper bound E_{UB} on the p th energy level for our system. Then we start at the level of the tree's root and consider all of its segments; for each choice of segment, the search problem corresponds to finding matchings on a smaller system with one less site in each of the d sets; the search can thus be implemented recursively. Suppose we have done k recursions; the sub-problem is associated with the node on our tree that is obtained by following the choices of hyperedges in the recursive construction. This node corresponds to a partial matching in which the first k sites of the first set have each been assigned a hyperedge. An important property is that all hyperedges have positive energies; then we know that any matching that is compatible with the current partial matching has an energy greater than it, thereby providing a lower bound on all the leaf energies obtainable from the current node. If that lower bound is greater than E_{UB} , then the sub-tree rooted on the current node can be pruned (discarded from the search); otherwise, one iterates the recursion (that is, one performs branching on the different choices of the hyperedge to include at the present level) and k goes to $k + 1$. When this process leads to a leaf that corresponds to a legal matching, we compute the energy E of this matching. If $E < E_{\text{UB}}$, we insert that matching into our list \mathcal{L} and remove its worst element so that it always has p elements; we also update E_{UB} which by definition is the largest energy of the matchings in \mathcal{L} ; in contrast, if $E > E_{\text{UB}}$, we discard the matching (leaf). After a finite number of branchings and prunings, the algorithm has explored all choices for the segments emerging from the tree's root and one is done. The best p matchings are then in the list \mathcal{L} .

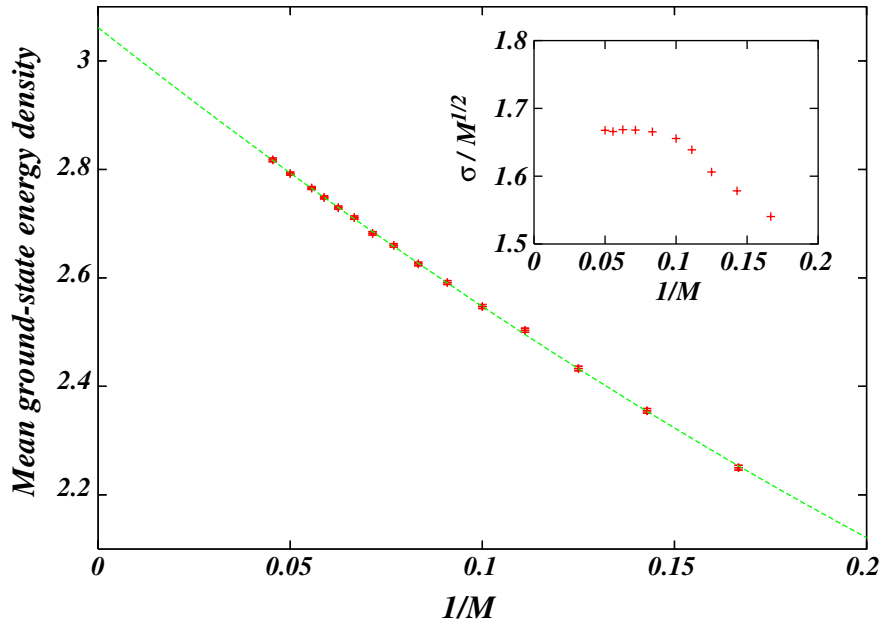


Figure 11. Mean ground state energy density as a function of $1/M$ at $d = 3$. The line is the quadratic fit using $M \geq 10$ data. Inset: the rescaled standard deviation of the ground state energy, suggesting a central limit theorem behaviour.

The algorithm without pruning requires $O(M^{M(d-1)})$ operations; with pruning and the different optimizations sketched in appendix B, the number of operations grows roughly by a constant factor when M is increased by 1; in particular, for the random instances studied here and $d = 3$, this factor is about 2.2.

6.2. Ground state energies

We generated a large number of random samples (disorder instances with the hyperedge costs taken to be independent uniformly distributed random variables in $[0, 1]$) and for each sample determined its ground state. We used several random number generators to check that our results were robust. Because of the exponential growth of the computation time with M , in practice we were limited to relatively modest values of M . For the results presented here and involving only ground states, at $d = 3$ we used 10 000 samples for $M = 20$ and $M = 22$, while for the smaller values of M we used 20 000 samples. We also performed runs at $d = 4$ but with lower statistics because the algorithm becomes less efficient as d increases; in fact, we were limited to $M \leq 14$ for that case and had only 5000 samples for each M .

Let us first focus on the behaviour of the ground state energy. For each sample, we determine with our branch and bound algorithm the ground state energy density $e_0 \equiv E_0/M \equiv M^{d-2}C_M^{(d)}$ (cf equation (7)); then we can analyse its mean in our ensemble or consider other properties of its distribution.

In figure 11 we show how the mean ground state energy density $\mathbb{E}[e_0]$ changes as one increases M . The behaviour is roughly linear in $1/M$, but by eye one can definitely see some curvature. Because of this, linear fits do not give good values of χ^2 unless the

$M < 10$ data are ignored; for instance, keeping only the $M \geq 10$ data, the linear fit gives 3.040(3) as the limiting value with $\chi^2 = 3.6$ for 9 degrees of freedom, while if we use all the data we obtain 3.021(3) with $\chi^2 = 32$ for 14 degrees of freedom. We have also tried corrections of the type $\ln(M)/M$ but this did not work well. Thus we proceed by considering quadratic fits. In that case, the resulting $M = \infty$ intercept does not depend much on whether one uses all or just the highest values of M . In particular, for all the data, we get the limiting value 3.046(5) with $\chi^2 = 9.6$ for 13 degrees of freedom, while using the $M \geq 10$ data only one has 3.06(1) with $\chi^2 = 2.3$ for 8 degrees of freedom. (In all these estimates, the error bars quoted are statistical only, as obtained from the statistical fluctuations.) We have also considered power fits, namely $\mathbb{E}[e_0] = a + b/M^c$. Fitting all the data gives the limiting value 3.08(1) with $\chi^2 = 7.2$ for 13 degrees of freedom while keeping only the $M \geq 10$ data leads to 3.09(3) with $\chi^2 = 2.3$ for 8 degrees of freedom (in both cases, the exponent c is close to 0.88). Since these χ^2 are similar to those of the quadratic fits, we see that the systematic errors are not negligible and are at least of the same order as the statistical errors; because of these effects, the agreement with the theoretical value of 3.126 can be considered rather good.

We studied similarly the case $d = 4$. The data again have positive curvature when plotted as a function of $1/M$, but since we have fewer statistics and a much smaller range of M , much less precision can be obtained for the large M limit. For the linear fit ($M \geq 9$) we get a limiting value of 6.75(3) with $\chi^2 = 4.7$ for 4 degrees of freedom. For the quadratic fit ($M \geq 9$ again), we get 7.22(8) with $\chi^2 = 0.37$ for 3 degrees of freedom. Finally, for the power fit we get 10.2(9) with $\chi^2 = 1.0$ for 5 degrees of freedom; the exponent is $c = 0.3$ which is small and leads to a large upturn for $M > 100$; clearly that regime is far beyond our reach and suggests that the power fit is probably inappropriate as non-robust (note for instance that the uncertainty on the limiting value is far higher here than for the other fits). The different estimates show that uncertainties arising from systematic effects (M too small) are severe; instead of the 1% precision we had at $d = 3$, we have a precision of at best 10% at $d = 4$ (compare to the theoretical prediction of 7.703). The conclusion is that numerics do not teach us much for the case $d = 4$ and so hereafter we shall concentrate on the different properties arising when $d = 3$.

One of the expectations for the d -index matching problem is that the free energy is self-averaging. Although at present there is no proof of such a property, there is no reason to expect otherwise; here we are limited by the numerical approach to ground states, but in that framework we can determine empirically the *distribution* of energies in the ensemble of random instances. Figure 12 displays the probability distribution of the (extensive) ground state energy E_0 for several values of M ($d = 3$). If as expected, the ground state energy is self-averaging, the relative width of these distributions should go to zero. We have thus measured the first few moments of these distributions. In the inset of figure 11, we have plotted the standard deviation σ of the ground state energy divided by \sqrt{M} as a function of $1/M$. Self-averaging corresponds to having $\sigma/M \rightarrow 0$; from the inset we see that $\sigma/M^{1/2}$ goes to a constant at large M so self-averaging holds and the convergence of the distribution is compatible with a central limit theorem type of behaviour; such a scaling arises from sums of not too dependent random variables and leads to a Gaussian limiting shape. To confirm this, we have looked at higher moments: we find that the skewness and kurtosis of the distributions do indeed decrease, in line with a central limit theorem type of convergence.

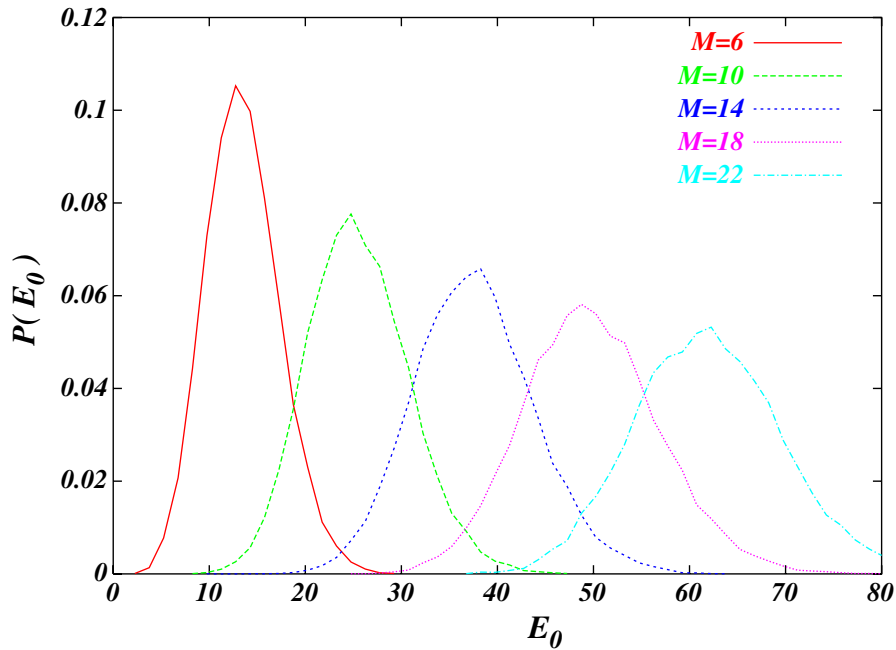


Figure 12. Distribution of the extensive ground state energy for increasing M values (from left to right) at $d = 3$.

Having a limiting Gaussian distribution for E_0 is not a consequence of the frozen 1RSB pattern of replica symmetry breaking since in the random energy model the distribution of E_0 follows a Gumbel distribution; furthermore, in that case the fluctuations in E_0 are $O(1)$ whereas in the matching problem they are $O(\sqrt{M})$. To see why such large fluctuations are ‘natural’, consider instead of E_0 the quantity \mathcal{E}_0 obtained by adding the lengths ℓ_i of the shortest hyperedges containing each site i of the first set. This quantity arises in a greedy algorithm (which does not necessarily generate a legal matching) and clearly one has $E_0 \leq \mathcal{E}_0$. The central limit theorem applies to \mathcal{E}_0 , so it will have a standard deviation that grows as \sqrt{M} and its distribution will become Gaussian at large M . The actual ground state energy E_0 is obtained by allowing hyperedge lengths that are slightly larger than the ℓ_i , but this should not suppress the large fluctuations nor prevent the central limit theorem scaling.

6.3. Other ground state properties

As discussed at the beginning of this paper, one expects the hyperedge containing a given site in the ground state matching to be one of the shortest possible ones. To investigate this issue quantitatively, let us order all the hyperedges containing a given site, going from the shortest to the longest hyperedge. The ‘order’ of a hyperedge is then 1 if it is the shortest, 2 if it is the next shortest etc. The orders arising in the ground state should be dominated by the lowest ones, 1, 2, 3, ... Consider thus the frequencies with which these orders arise; in figure 13 we show the behaviour of these frequencies for increasing M in the case $d = 3$. We see that there is a limiting histogram at large M , and that the lowest orders do indeed dominate. Furthermore, we see that for large k the probability

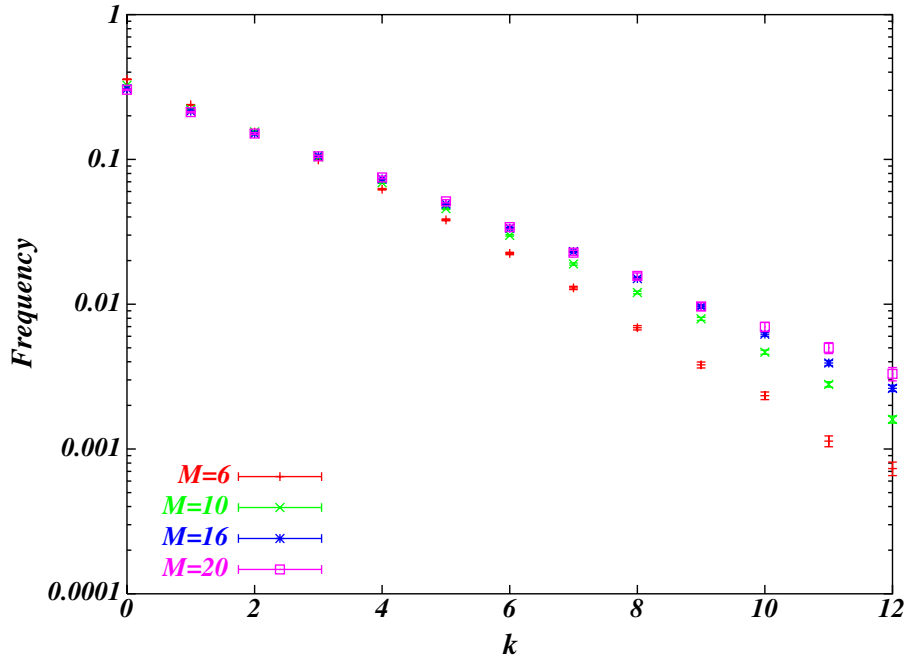


Figure 13. Histogram of the occupation probabilities in the ground state of the hyperedges as a function of their order k ($d = 3$). (The order is 1 for the lowest value among those hyperedges containing a given site, 2 for the next lowest value etc.) At large k , these frequencies approach an exponential law.

of occupation of an edge tends to decrease exponentially with k (the data are displayed on a semi-log plot). Note that in the standard matching ($d = 2$) problem, the decrease goes as $1/2^k$ exactly, while for our $d = 3$ case, the exponential decay is only asymptotic; furthermore, we have found no simple expression giving the decay rate of this exponential.

6.4. Excited states

Let us consider now states above the ground state. Define the excitation energy or ‘gap’ as $E_1 - E_0$ where E_0 is the extensive ground state energy and E_1 that of the next lowest energy state. In figure 14 we show that this random variable has a limiting distribution so that $E_1 - E_0 = O(1)$ in the large M limit, just as happens in the random energy model. Furthermore, the distribution is very well fitted by an exponential (cf the curve shown in the figure).

Following our theoretical conclusions obtained earlier, consider now the overlap between the ground state and the first excited state. In our frozen 1RSB picture, these matchings are expected to have a fixed (self-averaging) overlap when M grows. In figure 15 we show the probability distribution of such overlaps for increasing M . We see that there is a local peak at large overlap that shifts toward $q = 1$ but which simultaneously decays. The bulk of the overlaps however arise around $q = 0.3$ and when M increases we see that the corresponding peak gets both higher and more narrow. Overall, the behaviour is compatible with a convergence toward a Dirac peak near $q = 0.32$, to be compared with the theoretical prediction $q_c = 0.321$.

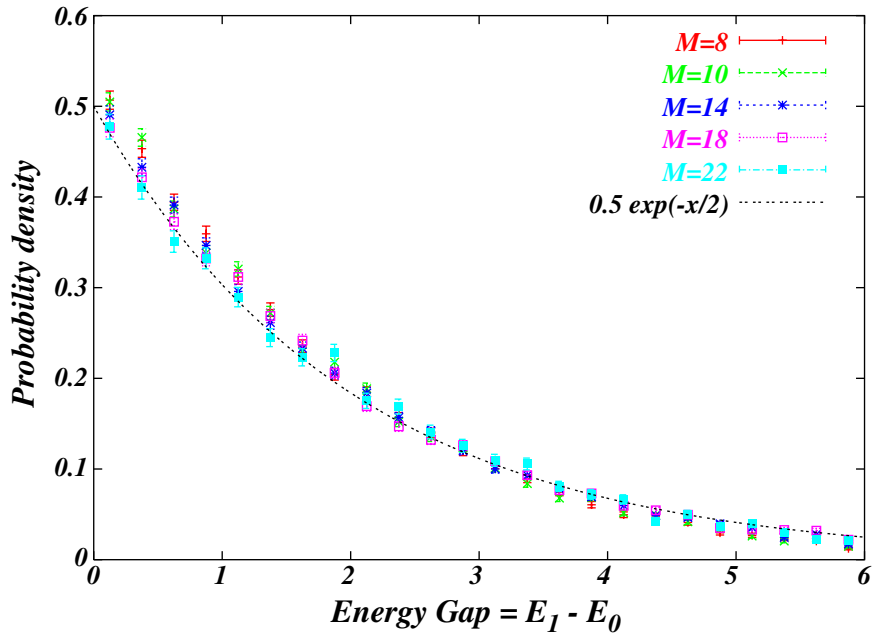


Figure 14. Probability density of the gap, $E_1 - E_0$, that is the energy difference between the first excited state and the ground state (extensive) energies in the case $d = 3$. The curve is a pure exponential to guide the eye.

6.5. Low energy entropy

Finally, consider the *density* of energy levels. In the case of the random energy model, this density becomes self-averaging when the excitation energy grows. We have thus computed the disorder averaged density of levels as a function of the excitation energy, $E - E_0$. This is a measure of the exponential of the microcanonical entropy; within the frozen 1RSB scenario, it gives the critical temperature via $\rho(E - E_0) \sim \exp[\beta_c(E - E_0)]$. In figure 16 we display our numerical estimate of ρ and see that it is very nearly a pure exponential. From the slope on the semi-log plot we extract $\beta_c \approx 0.405$; this value should be compared to the theoretical prediction of 0.412; the agreement is reasonable but not perfect. To get better agreement, we believe it would be necessary to go to larger M and also to go further in the self-averaging regime, i.e., to consider larger $E - E_0$ which numerically is an arduous task.

7. Conclusion

We presented an analysis of multi-index matching problems (MIMPs) based on an adaptation of the cavity method for finite connectivity systems. For the well known two-index matching problem, our approach provides an alternative derivation of results previously obtained using the replica and cavity methods. With respect to these older studies, the present one has the advantages of being closer to the mathematical framework developed by Aldous, and of allowing replica symmetry breaking effects to be incorporated in a tractable manner. Exploiting this latter possibility, we predict the value of the asymptotic minimal cost to be given for d -index matching problems by $\mathcal{L}^{(d)} = \epsilon_{\text{RS}}(\beta_s)$

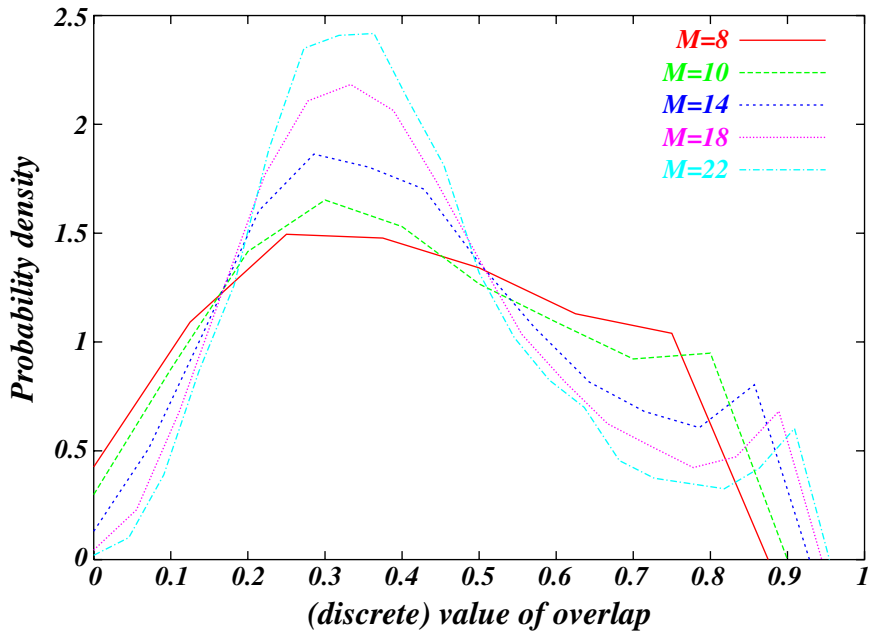


Figure 15. Probability density of the overlap q between the ground state and the first excited state for increasing M ($d = 3$).

with ϵ_{RS} obtained from equations (29) and (30) and β_s satisfying $s_{RS}(\beta_s) = 0$. Formally, this $d \geq 3$ conjecture differs from the case $d = 2$ (where it is a theorem) in that $\beta_s = \infty$ when $d = 2$, while $\beta_s < \infty$ when $d \geq 3$. The distinction between 2-index and d -index matching problems with $d \geq 3$ arises clearly from our analytical and numerical analysis: in the first case all low cost matchings share most of their hyperedges, while in the second case they differ from each other by a finite fraction of their hyperedges. In mathematical terms, the essential uniqueness property does not hold when $d \geq 3$ or in physical terms replica symmetry must be broken. Extending Aldous's framework to rigorously account for this fact and providing a proof of our conjecture for $d \geq 3$ seems to us a particularly interesting mathematical challenge.

From a physical perspective, the qualitative difference between 2-index and d -index matchings problems with $d \geq 3$ hinges on the presence at low temperature of a glassy phase. This is similar to the difference that has been found between the 2-SAT and 2-colouring problems, which are polynomial, and the K -SAT and q -colouring problems with $K \geq 3$ and $q \geq 3$, which are NP-complete. At variance with SAT or colouring problems, the nature of the glassy phase of MIMPs is however simpler, as it is made of isolated configurations instead of separate clusters of many configurations. We termed this phase a 'frozen 1RSB glassy phase' and attributed it to the nature of the constraints, called hard constraints. As a technical consequence of this distinctive feature, a particular frozen 1RSB ansatz has to be implemented. Such an ansatz has repeatedly been used in the literature as a convenient (but rarely justified) substitute for the more complicated general 1RSB ansatz; our discussion on the role of hard constraints provides a clarification of its conditions of validity which we believe is of general interest for the investigation of glassy phases in other systems. Finally, past studies of SAT and colouring problems have

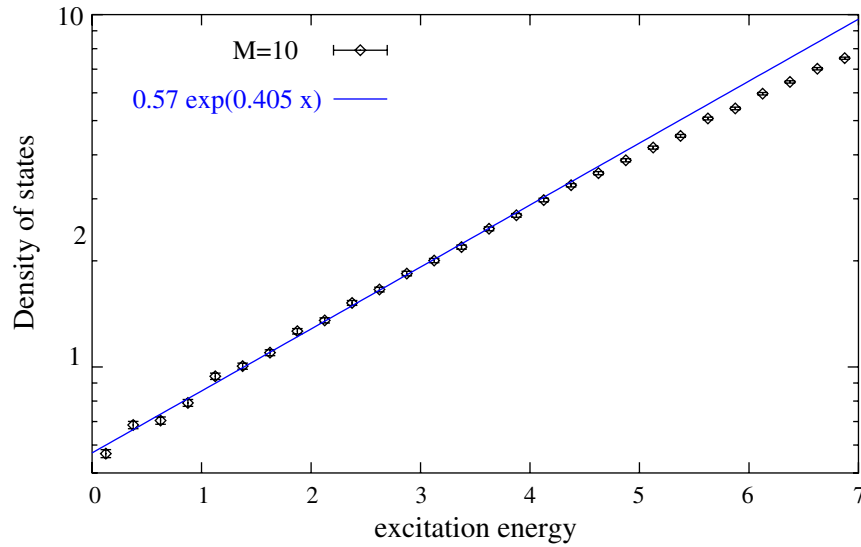


Figure 16. Density of energy levels, measured from the ground state energy. At low energies (before finite size effects dominate), this density grows exponentially as $\exp[\beta_c(E - E_0)]$, thereby giving the model's critical temperature. Shown is the case $d = 3$ for $M = 10$.

shown how the cavity framework could serve as a basis for developing efficient algorithms for analysing single instances and finding low energy configurations [4, 5]: as matchings present a number of differences from the problems studied so far, it would be particularly interesting to investigate whether such an approach can also be successfully implemented with MIMPs.

Acknowledgments

This work was supported in part by the European Community's Human Potential Programme under contracts HPRN-CT-2002-00307 (DYGLAGEMEM) and HPRN-CT-2002-00319 (STIPCO) as well as by the Community's EVERGROW Integrated Project.

Appendix A: Population dynamics algorithm

Here we give a short description of the population dynamics algorithm we used to solve the RS cavity equations. We implemented two different versions, corresponding to the two different cut-off procedures mentioned in the text, associated either with Poissonian (algorithm P) or with regular graphs (algorithm R). In addition to the inputs d and β , the algorithm has essentially three parameters: the mean degree of the nodes, C (algorithm P) or K (algorithm R), the size of the population, \mathcal{N}_{pop} , and the number of iterations $\mathcal{N}_{\text{iter}}$. The common structure of the two algorithms is the following:

- Initialize with random values a population of cavity fields $x[i]$, $i = 1, \dots, \mathcal{N}_{\text{pop}}$.
- Do $\mathcal{N}_{\text{trans}} = 100$ times: Update().
- Do $\mathcal{N}_{\text{iter}}$ times: Update() and Measure().

The first loop allows the system to equilibrate toward the stationary distribution. The subroutine Update() depends on the cut-off procedure and can be schematically described as follows:

Do \mathcal{N}_{pop} times:

- Draw k either at random with a Poissonian distribution of mean C (algorithm P), or take $k = K$ (algorithm R).
- Draw costs $\{\xi_a\}_{a=1,\dots,k}$ either independently with the uniform distribution in $[0, C]$ (algorithm P), or according to a Poisson process with rate 1 (algorithm R).
- Draw at random $k(d-1)$ members of the population and use them together with the ξ_a to compute a new field x_0 according to equation (22).
- Draw at random one member of the population and replace its cavity field value with x_0 .

The subroutine Measure() is implemented similarly and computes the free energy according to equations (27) and (28). The final output for the free energy is obtained by averaging over the $\mathcal{N}_{\text{iter}}$ iterations, while the fluctuations across iterations are used to check convergence. The algorithm must be run for increasing values of C or K to extrapolate the $C \rightarrow \infty$ (algorithm P) or $K \rightarrow \infty$ (algorithm R) limit, requiring one to consider larger and larger population sizes \mathcal{N}_{pop} to obtain reliable results. Taking this limit is however facilitated by the numerical observation that the Poissonian approximation (algorithm P) approaches the solution from below while the regular approximation (algorithm R) approaches it from above; this is illustrated in figure 9 with $d = 4$. We refer the reader to the captions of the various figures for typical choices of the parameters C , \mathcal{N}_{pop} and $\mathcal{N}_{\text{iter}}$. The numerical results we obtained for $d = 2$ are consistent with the exact solution, $\beta_c = \infty$ and $f_{\text{RS}}(\beta_c) = \pi^2/12$, and are the following for $d = 3, 4$:

$$\begin{aligned} d = 3: \quad \beta_c &= 0.412 \pm 0.001, & f_{\text{RS}}(\beta_c) &= 1.042 \pm 0.0003, \\ d = 4: \quad \beta_c &= 0.135 \pm 0.002, & f_{\text{RS}}(\beta_c) &= 1.925 \pm 0.0006. \end{aligned} \tag{A.1}$$

The free-energy densities are given here for simple matching problems and their counterparts for d -partite matchings are obtained by multiplying the values by d .

We have also implemented the generalization of this algorithm to solve the 1RSB cavity equations (50) and used it to check that no discontinuous transition occurs prior to the entropy crisis (see [27] for algorithmic details).

Appendix B: Aspects of the branch and bound algorithm

Our objective is to solve d -partite matching problems at sufficiently large M so that an extrapolation to the $M \rightarrow \infty$ limit can be performed without too much uncertainty. For many problems (satisfiability, colouring etc), one prefers an easily implementable algorithm such as one in the class of ‘heuristic’ algorithms; in such approaches one performs a fast search for the ground state but no guarantee is provided that the global optimum will be found. Examples of these algorithms are simulated annealing and variable depth local search. Heuristic algorithms typically attempt to move towards regions of lower energy by searching in the neighbourhood of a current configuration. However, since the search is local, such an approach is bound to break down for problems in

which the frozen 1RSB scenario applies. This fact pushed us towards the development of an ‘exact’ algorithm capable of delivering a certificate of optimality of the proposed ground state. Amongst exact algorithms, enumeration can be discarded because it is much too slow; ‘branch and bound’ gets around this problem through pruning of the enumeration/search. There are also other possible methods such as ‘branch and cut’, but these require an in depth understanding of polytopes and rely on separation procedures which have not yet been developed for MIMP. Note that in all exact methods, the key to efficiency is having good bounds; fortunately MIMPs are relatively well adapted to such a strategy.

We already discussed in the main text our choice of representation of matchings and partial matchings. Given a partial matching of the first k sites of the first set, we have to solve a MIMP with $M - k$ sites and so the algorithm can be implemented recursively. Since at each node we need to consider all of its possible branchings (naively, there are $(M - k)^{d-1}$ of these), it is useful to order these branchings according to the length of the corresponding hyperedges, going from short to long. Rather than recompute these orderings dynamically every time the partial matching changes, we do it once and for all at the initialization of the program. This allows for speed but it must be compensated by a rapid determination of whether a given hyperedge is allowed; for that we use a data structure which tells us for each site of each set whether it is matched (belongs to one of the occupied hyperedges). This structure is updated whenever a partial matching is extended or reduced.

The pruning of the search must be as stringent as possible, and this depends on the quality of the bounds. Our simplest bound B_1 is just the current partial matching’s energy E_k : if that energy is higher than E_{UB} (the upper bound E_{UB} as defined in section 6.1), then the whole sub-tree below the current node can be pruned. A better bound is B_2 , obtained by adding to E_k the sum over each remaining unmatched site of the first set of the shortest hyperedge containing that site. This sum can be precomputed and tabulated. A still better bound is B_3 obtained as B_2 but where now one takes for each site the shortest hyperedge that is compatible with the current partial matching. This bound cannot be predefined once and for all and is slow to compute. Since we have found it to be useful for pruning, we have optimized its determination by noticing that it can be tabulated and modified incrementally: every time the partial matching is extended (a hyperedge is added), we perform the search for the compatible hyperedges of each unmatched site starting from the index (order) previously found to be compatible. When backtracking, one has to remove a hyperedge and there we simply go back to the tables we had at that level: in effect, we maintain efficiency if we assign tables at each level and follow their updating one step at a time.

The rates of pruning are very different for the three bounds, and we found that a good strategy (for balancing pruning rate and computation time) was to apply the three bounds successively: if the first one does not prune, one goes on to the second one and so forth. To speed up the computation further, we found it useful to implement the recursivity of the program in a limited mode only: the data structures are set up once and for all at initialization time, the hyperedges are ordered once and for all too and then the recursion is used mainly to go through the branchings and to maintain the tables. Efficiency is gained as no reorganization of the instance (hyperedge weights) is performed, and in particular no ‘smaller matching problem’ is ever defined explicitly.

References

- [1] Mézard M and Parisi G, 1985 *J. Physique* **46** L771
- [2] Mézard M and Parisi G, 1986 *Europhys. Lett.* **2** 913
- [3] Aldous D J, 2001 *Random Struct. Algorithms* **18** 381
- [4] Mulet R, Pagnani A, Weigt M and Zecchina R, 2002 *Phys. Rev. Lett.* **89** 268701
- [5] Mézard M and Zecchina R, 2002 *Phys. Rev. E* **66** 056126
- [6] Mézard M, Parisi G and Zecchina R, 2002 *Science* **297** 812
- [7] Martin O C, Mézard M and Rivoire O, 2004 *Phys. Rev. Lett.* **93** 217205
- [8] Mézard M, Parisi G and Virasoro M A, 1987 *Spin-Glass Theory and Beyond (Lecture Notes in Physics vol 9)* (Singapore: World Scientific)
- [9] Mézard M and Parisi G, 1987 *J. Physique* **48** 1451
- [10] Parisi G and Ratiéville M, 2002 *Eur. J. Phys. B* **29** 4
- [11] Parisi G and Ratiéville M, 2001 *Eur. J. Phys. B* **22** 229
- [12] Houdayer J, Boutet de Monvel J H and Martin O C, 1998 *Eur. Phys. J. B* **6** 383
- [13] Aldous D and Steele M, 2003 *Encyclopedia of Mathematical Sciences* vol 110, ed H Kesten (Berlin: Springer) pp 1–72
- [14] Linusson S and Wastlund J, 2003 *Preprint* math.CO/0303214
- [15] Nair B P C and Sharma M, 2003 available at <http://www.stanford.edu/~balaji/rap.html>
- [16] Parisi G, 1998 *Preprint* cond-mat/9801176
- [17] Talagrand M, 2003 *Spin Glasses: A challenge for Mathematicians* (Berlin: Springer)
- [18] Pierskalla W P, 1968 *Oper. Res.* **16** 422
- [19] Spieksma F C R, 2000 *Nonlinear Assignment Problems, Algorithms and Applications* ed L Pitsoulis and P Pardalos (Dordrecht: Kluwer) pp 1–12
- [20] Burkard R E, 2002 *Discrete Appl. Math.* **123** 257
- [21] Poore A B, 1994 *Comput. Opt. Appl.* **3** 27
- [22] Puztaszeri J, Rensing P E and Liebling T M, 1995 *J. Global Optim.* **16** 422
- [23] Karp R, 1972 *Complexity of Computer Computations* ed R Miller and J Thatcher (New York: Plenum) pp 85–103
- [24] Papadimitriou C H and Steiglitz K, 1982 *Combinatorial Optimization: Algorithms and Complexity* (Englewood Cliffs, NJ: Prentice-Hall)
- [25] Vannimenus J and Mézard M, 1984 *J. Physique Lett.* **45** L1145
- [26] Aldous D J, 1990 *Random Struct. Algorithms* **1** 383
- [27] Mézard M and Parisi G, 2001 *Eur. Phys. J. B* **20** 217
- [28] Mézard M and Parisi G, 2003 *J. Stat. Phys.* **111** 1
- [29] Aldous D J and Bandyopadhyay A, 2004 *Preprint* math.PR/0401388
- [30] Montanari A and Ricci-Tersenghi F, 2003 *Eur. Phys. J. B* **33** 339
- [31] Rivoire O, Biroli G, Martin O C and Mézard M, 2004 *Eur. Phys. J. B* **37** 55
- [32] Monasson R, 1995 *Phys. Rev. Lett.* **75** 2847
- [33] Derrida B, 1980 *Phys. Rev. Lett.* **45** 79
- [34] Derrida B and Spohn H, 1988 *J. Stat. Phys.* **51** 817
- [35] Krauth W and Mézard M, 1989 *Europhys. Lett.* **8** 213
- [36] Barkai E and Kanter I, 1991 *Europhys. Lett.* **14** 107
- [37] Mézard M, Ricci-Tersenghi F and Zecchina R, 2003 *J. Stat. Phys.* **111** 505
- [38] Montanari A, 2001 *Eur. Phys. J. B* **23** 121
- [39] Gross D J, Kanter I and Sompolinsky H, 1985 *Phys. Rev. Lett.* **55** 304
- [40] Mertens S, Mézard M and Zecchina R, 2003 *Preprint* cs.CC/0309020
- [41] Rivoire O, 2005 *PhD Thesis* Université Paris-Sud